Physics-Informed Signed Distance Fields for Flood Arrival Time Prediction and Evacuation Routing

Yameng Guo
yameng.guo@ugent.be
dept. Business Informatics and Operations Management
Ghent University
Gent, Belgium

Seppe vanden Broucke
seppe.vandenbroucke@ugent.be
dept. Business Informatics and Operations Management
Ghent University
Gent, Belgium
Research Centre for Information Systems Engineering
KU Leuven
Leuven, Belgium

***Abstract*—Accurate flood prediction with explicit arrival times is critical for effective evacuation planning. Traditional physics-based models provide reliable predictions but are computationally expensive, while existing data-driven approaches typically predict binary flood extent or water depth without explicit timing information. We propose a novel framework that represents flood evolution as a temporal Signed Distance Field (SDF), enabling efficient prediction of both flood boundaries and arrival times. Our approach combines spatial-temporal deep learning with physics-informed constraints in a variational optimization framework. The predicted SDF naturally supports time-dependent evacuation routing by providing continuous distance-to-flood information and explicit arrival time fields. The proposed idea will be compared with grid-based deep learning approaches as well as physical constraints approaches regarding real-time evacuation planning and arrival time accuracy.**

Keywords-Signed distance function; Flood prediction; Evacuation routing.

I. INTRODUCTION

Flooding is one of the most devastating natural disasters, affecting millions worldwide annually. Effective evacuation requires knowledge of not only *where* flooding will occur, but critically *when* each location will be inundated. Traditional approaches face a fundamental trade-off between computation time and accuracy performance. Physics-based models (e.g., Hydrologic Engineering Center's River Analysis System (HEC-RAS) [1], LISFLOOD-FP [2]) provide reliable predictions by solving shallow water equations but require hours of computation, while data-driven methods [3] offer speed but typically predict only binary flood extent or depth fields without explicit timing.

To address this, we propose representing flood evolution as a Signed Distance Field (SDF) [4][5], where $\varphi(x, y, t)$ encodes the signed distance to the flood boundary at each location and time. Compared to traditional ML grid-based methods and physical modeling, SDF provides advantages in both geometric encoding and physics constraint formulation. The direct and continuous encoding of flood extent enables natural extraction of arrival times from level set evolution and helps the model learn continuous boundaries rather than discrete water depths. Additionally, SDF uses information more efficiently for sparse flooding patterns. For instance, for zero values, depth grids will waste the capacity on zeros, while the distance field utilizes zero as meaningful input indicating the flood boundary. Moreover, SDF enables simpler, geometrically cleaner physical constraints that allow constraining boundary motion without modeling all fluid dynamics.

In summary, this work aims to introduce the following key contributions:

- Apply temporal SDF evolution to flood prediction with explicit arrival time modeling.
- Use a physics-informed variational optimization framework to ensure realistic propagation.
- Enable end-to-end system from prediction to evacuation routing with safety guarantees.
- Implement comprehensive evaluation against both physics-based and data-driven baselines.

To the best of our knowledge, this is the first work applying temporal SDF evolution to flood prediction.

The paper is organized as follows: In Section III, we discuss the methodology of using SDF representations, the model architecture, and the design of the physics-aware loss function. In Section IV, we describe the experimental setup, including the dataset, baselines, and selected evaluation metrics. The initial results are presented in Section V, and we conclude in Section VI.

II. RELATED WORK

SDFs have been widely used in computer vision [6], 3D graphics [7], and robotic systems [8], where they play a critical role in geometry processing, mapping, and planning.

Beyond these domains, SDFs have also been applied in structural engineering tasks such as topology optimization [5], and have been extended to incorporate physical awareness in scenarios such as fluid simulation [9].

Despite their success in various physics-aware applications, SDFs remain largely unexplored in the context of flood modeling. Current flood modeling approaches are predominantly divided into two categories: physics-based models and machine learning or deep learning models [10][11], typically leveraging multisource data such as rainfall, Digital Elevation Models (DEMs), and land surface characteristics. Although the methodological spectrum continues to expand, the underlying data representation and prediction targets remain largely unchanged. That is, the study area is typically discretized into grid cells, and the flood state is predicted for each cell.

In this context, SDF-based representations offer a promising alternative by encoding geometric and spatial information in a continuous and physically meaningful manner. Unlike discrete grid-based representations, SDFs naturally capture boundary information and spatial relationships, which are critical for accurately modeling flood dynamics such as water propagation and interaction with terrain. Furthermore, SDFs can be seamlessly integrated with machine learning and deep learning models, providing rich geometric priors that may improve generalization and robustness.

III. METHODOLOGY

A. SDF Representation

We define the flood state at time t as:

$$\varphi(x,y,t) = \begin{cases} d((x,y), \partial\Omega_t) & \text{if } (x,y) \in \Omega_t \\ -d((x,y), \partial\Omega_t) & \text{otherwise} \end{cases} \quad (1)$$

where Ω_t is the flooded region, $\partial\Omega_t$ its boundary, and $d(\cdot,\cdot)$ is Euclidean distance. The zero level set $\{\varphi = 0\}$ represents the flood boundary.

Arrival time at location (x,y) is extracted as:

$$t_{\text{arr}}(x,y) = \min\{t : \varphi(x,y,t) \geq 0\} \quad (2)$$

The representation shown in Figure 1 illustrates the difference between a grid-based field and a signed-distance field. In the grid-based field, each discrete cell has its own target water-depth value, whereas in the SDF representation, the target is a continuous boundary of the water area.

B. Temporal SDF Evolution Model

Our model predicts the SDF trajectory $\{\varphi_t\}_{t=0}^{T}$ given initial conditions φ_0, terrain elevation z, and rainfall forcing $r(t)$.

The proposed hybrid spatial-temporal architecture consists of three components, including a spatial encoder-decoder to extract multi-scale spatial features from $[\varphi_t, z, r_t]$, a temporal model to capture long-range temporal dependencies efficiently and a physics layer to enforce constraints during decoding.

C. Physical-constraints Loss

We formulate training as a constrained optimization problem with a physical-constraints loss applying penalty on both data fitting error and physical feasibility.

$$\mathcal{L} = \underbrace{\|\varphi_{\text{pred}} - \varphi_{\text{obs}}\|^2}_{\text{data fit}} + \lambda_1 \underbrace{\|\partial_t\varphi + v \cdot \nabla\varphi\|^2}_{\text{level set PDE}}$$
$$+ \lambda_2 \underbrace{\|\text{ReLU}(\nabla\varphi \cdot \nabla z)\|^2}_{\text{downhill flow}} + \lambda_3 \underbrace{\|\text{volume}_t - \text{rain}_t\|}_{\text{mass balance}} \quad (3)$$

The level-set term enforces $\partial\varphi/\partial t + v \cdot \nabla\varphi = 0$ where v is flood velocity (estimated from terrain/rainfall). This ensures the SDF evolves consistently with flood propagation physics. The downhill flow term is to enforce that flood propagation does not move uphill, and the last mass term is to ensure flooded water volume match rainfall input over time.

D. Evacuation Routing

Given the predicted SDF trajectory and arrival time field, we solve time-dependent evacuation routing as follows ("safe zone computation"): at current time t_c, location (x,y) is safe if:

$$t_{\text{arr}}(x,y) > t_c + \tau_{\text{buffer}} \quad (4)$$

IV. EXPERIMENTAL DESIGN

A. Datasets

We plan to evaluate our approach on both real and synthetic datasets. The real-life dataset includes historical flood events with satellite-derived inundation maps, Digital Elevation Models (DEMs), and rainfall records. In our initial experiments, we use a UK flood event dataset from FloodCastBench [12].

However, existing real-world datasets are significantly limited by the coarse granularity of rainfall data, the relatively small number of recorded flood events, and the frequent absence of a long-term temporal dimension. To address these limitations, we plan to incorporate synthetic data in future work, particularly for rainfall generation and flood extent simulation. Specifically, we will leverage the open-source package Landlab [13], which supports numerical modeling of Earth surface dynamics, including rainfall generation and water-driven erosion processes.

B. Baselines

Two baseline categories will be included, i.e., physics-based models and grid-based deep learning approaches. Physics-based baselines, such as the industry-standard HEC-RAS 2D and LISFLOOD-FP for large-scale floods, can be included into comparison experiment. For deep learning baselines, we use convolutional models with Long Short-Term Memory (LSTM) kernels to predict water depth at each cell over time.

C. Evaluation Metrics

Performance will be evaluated across three dimensions: flood prediction accuracy, evacuation performance, and computational cost.

For flood prediction, we will use Intersection over Union (IoU) [14][15] to measure the overlap between predicted and ground-truth flood masks. We will also evaluate arrival time accuracy using Root Mean Square Error (RMSE) between predicted and actual inundation times.

For evacuation performance, we measure success rate, i.e., the percentage of evacuations reaching safety before flood arrival. Finally, we record time and memory usage during training and inference to assess computational cost.

V. INITIAL RESULTS

We conducted an initial proof-of-concept study on the UK event from the FloodCastBench data set [12] to assess whether an SDF target already provides practical value before introducing the full physics-informed training setup. Each sample consists of six past flood states together with the DEM height map and a scalar mean-rainfall summary. The model predicts the next six flood states in a direct sequence-to-sequence manner. To keep the comparison controlled, we used

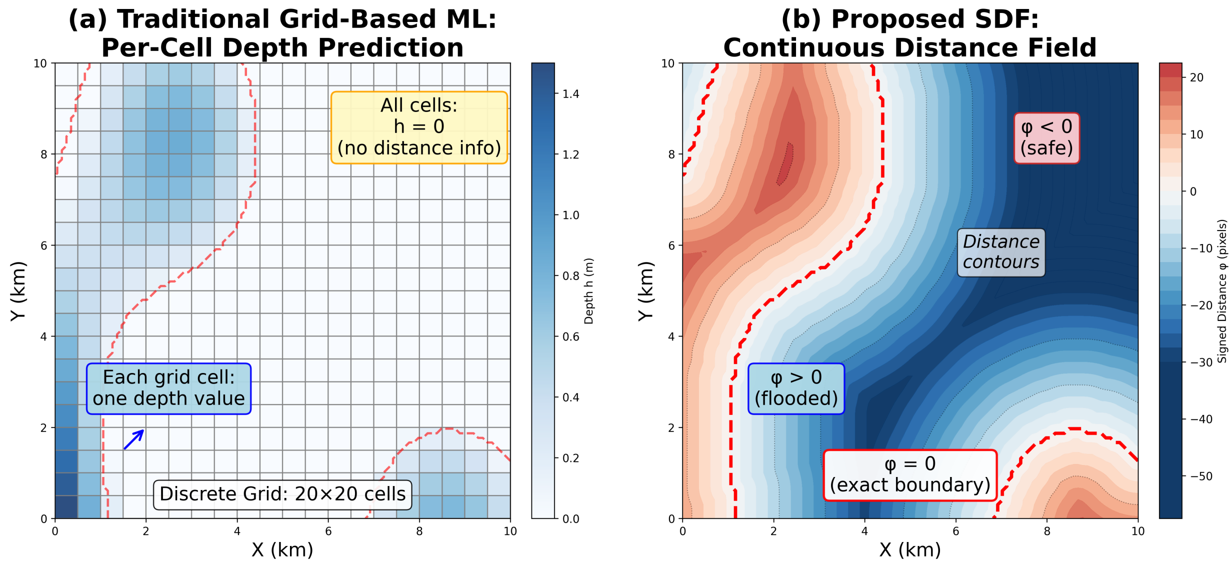

Figure 1. Comparison of flood area representations using a grid-based method and a signed-distance field.

the same small U-Net style backbone across all runs and only varied the prediction target. The resulting baselines compare a signed-distance field with a classic coarse cell-based mask representation.

The targets are organized as follows. For the mask approach, each grid cell is assigned a binary label indicating whether the corresponding area is flooded. For the SDF representation, the target takes continuous values between -1 and 1, where -1 denotes non-flooded areas, 0 represents the SDF boundary, and 1 indicates flooded areas.

The UK event targets were cleaned up and evaluated using a chronological train/validation/test split. We incorporated four complementary evaluation metrics. Flood-extent overlap is measured using Intersection over Union (IoU),

$$\mathrm{IoU} = \frac{|P \cap G|}{|P \cup G|}, \tag{5}$$

and Dice,

$$\mathrm{Dice} = \frac{2|P \cap G|}{|P| + |G|}, \tag{6}$$

where P and G denote the predicted and ground-truth flooded regions. We additionally consider a boundary F1 score to assess whether predicted and reference flood boundaries match within a small spatial tolerance, and an arrival-time mean absolute error

$$\mathrm{F1}_\partial = \frac{2\,\mathrm{Prec}_\partial\,\mathrm{Rec}_\partial}{\mathrm{Prec}_\partial + \mathrm{Rec}_\partial}, \tag{7}$$

where Prec_∂ and Rec_∂ denote boundary precision and recall after allowing a small tolerance around the reference contour, and an arrival-time mean absolute error

$$\mathrm{MAE}_{\mathrm{arr}} = \frac{1}{N}\sum_{i=1}^{N} \left|\hat{t}_{\mathrm{arr},i} - t_{\mathrm{arr},i}\right| \tag{8}$$

to summarize timing differences between predicted and observed first inundation.

Under this setup, the SDF approach achieved slightly improved flood-extent performance compared to the coarse-cell mask baseline, as shown in Table I. This suggests that, at least for the current direct-prediction backbone, the SDF target is viable and already competitive as a representation, even though it does not yet yield a decisive aggregate-metric advantage over the mask baseline. This is mainly due to the fact that this event is generally stable in terms of flood-extent and that after an initial on-fall of rain, the flood depths remain relatively stable over time.

TABLE I. INITIAL HELD-OUT RESULTS ON THE CLEANED UK FLOOD EVENT.

Model	IoU	Dice	$\mathrm{F1}_\partial$	$\mathrm{MAE}_{\mathrm{arr}}$
Coarse mask	0.9975	0.9987	0.9998	0.0074
SDF	0.9976	0.9988	0.9988	0.0059

At the same time, the qualitative behavior of the SDF target is encouraging. As can be observed in the comparison of Figure 2, Figure 3 and Figure 4, the SDF representation yields a continuous geometric field that is easier to overlay on top of terrain and more naturally highlights evolving flood structure than a purely binary occupancy map. On the particular frame shown, the SDF approach presents much better evaluation metrics than the averaged results in Table I. The initial experiments therefore support the core motivation of the paper: SDF is a promising geometric representation for flood prediction, even though larger gains will likely depend on stronger temporal models and the full physics-informed formulation proposed in this work.

These initial results refine the direction of the full study. Rather than assuming that SDF will dominate grid-based targets out of the box, the experiments indicate that the most promising benefit of SDF lies in boundary-aware structure, interpretability, and its compatibility with future physics-informed constraints.

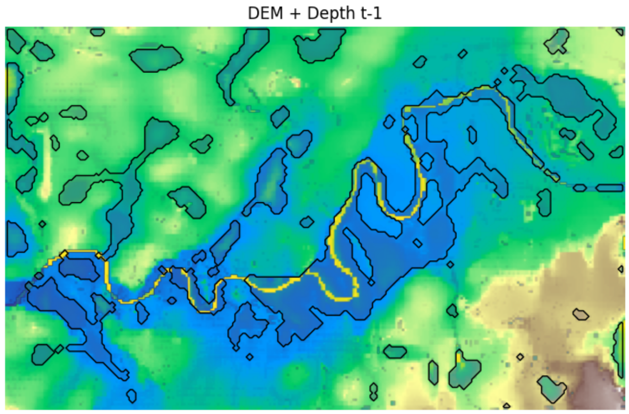

Figure 2. Initial DEM of the cleaned UK flood event.

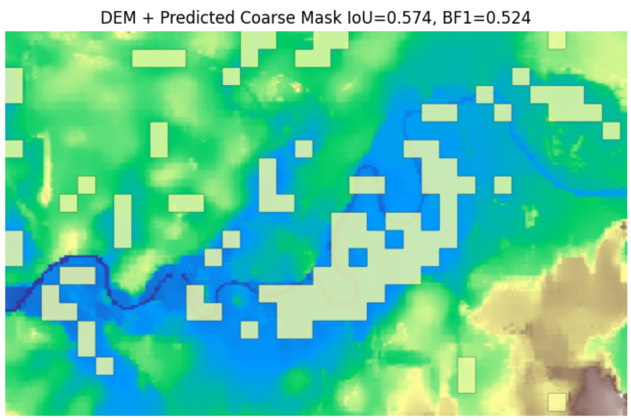

Figure 3. The results of coarse cell-based mask prediction (without explicit flood boundary delineation).

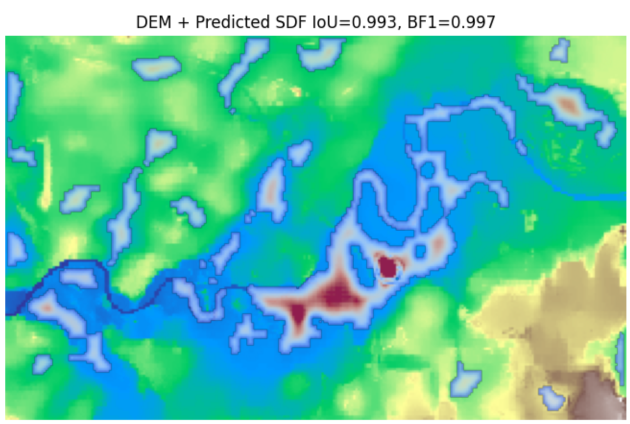

Figure 4. The continuous signed-distance field prediction (with explicit flood boundary delineation).

This makes SDF a strong candidate representation for the next phase of work.

VI. Conclusion and Future Work

We present a novel physics-informed SDF framework for flood prediction that combines geometric representation, temporal evolution modeling, and optimization-based constraints. By explicitly targeting arrival times alongside flood extent, the framework is intended to bridge the gap between fast data-driven methods and reliable physics-based models. Our initial UK case study already indicates that SDF is a viable and visually compelling representation, with performance comparable to a grid-based mask baseline under a matched small U-Net backbone and with clearer geometric structure for qualitative analysis.

At the same time, several parts of the full idea remain to be implemented and tested in the next phase. First, the current results are based on simple direct-prediction baselines rather than the full physics-informed temporal architecture. Second, the physical-constraints loss still needs to be integrated and evaluated systematically, including the level-set evolution term, downhill-flow consistency, and mass-balance constraints. Third, the arrival-time and evacuation-routing components have been formulated conceptually but still need end-to-end empirical validation. Finally, broader experiments are required to test whether the representation advantage of SDF becomes more pronounced under stronger temporal backbones (e.g. autoregressive backbone models), harder forecasting regimes, and more realistic operational settings.

Overall, the present study provides encouraging early evidence for the representation itself and clarifies the most promising next steps for turning the proposed SDF framework into a complete physics-informed flood forecasting and evacuation-support system.

References

[1] G. W. Brunner, "Hec-ras river analysis system. hydraulic reference manual. version 1.0.", *User manual. US Army Corps of Engineers, Hydrologie Engineering Center*, 1995.

[2] P. Bates, M. Trigg, J. Neal, and A. Dabrowa, "Lisflood-fp", *User manual. School of Geographical Sciences, University of Bristol. Bristol, UK*, 2013.

[3] A. Mosavi, P. Öztürk, and K. Chau, "Flood prediction using machine learning models: Literature review", *CoRR*, vol. abs/1908.02781, 2019. arXiv: 1908.02781.

[4] S. Osher and R. Fedkiw, "Signed distance functions", in *Level set methods and dynamic implicit surfaces*, Springer, 2003, pp. 17–22.

[5] O. Jezek et al., "Smooth geometry extraction from SIMP topology optimization: Signed distance function approach with volume preservation", *Adv. Eng. Softw.*, vol. 212, p. 104 071, 2026. DOI: 10.1016/J.ADVENGSOFT.2025.104071.

[6] D. Vicini, S. Speierer, and W. Jakob, "Differentiable signed distance function rendering", *ACM Trans. Graph.*, vol. 41, no. 4, 125:1–125:18, 2022. DOI: 10.1145/3528223.3530139.

[7] J. Shim, C. Kang, and K. Joo, "Diffusion-based signed distance fields for 3d shape generation", in *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, IEEE, 2023, pp. 20 887–20 897. DOI: 10.1109/CVPR52729.2023.02001.

[8] B. Liu, G. Jiang, F. Zhao, and X. Mei, "Collision-free motion generation based on stochastic optimization and composite signed distance field networks of articulated robot", *IEEE Robotics Autom. Lett.*, vol. 8, no. 11, pp. 7082–7089, 2023. DOI: 10.1109/LRA.2023.3311357.

[9] C. Zhang, "Sdfibm: A signed distance field based discrete forcing immersed boundary method in openfoam", *Comput. Phys. Commun.*, vol. 255, p. 107 370, 2020.

[10] D. K. Hakim, R. Gernowo, and A. W. Nirwansyah, “Flood prediction with time series data mining: Systematic review”, *Natural Hazards Research*, vol. 4, no. 2, pp. 194–220, 2024.

[11] F. Yang et al., “Rapid urban flood inundation forecasting using a physics-informed deep learning approach”, *Journal of Hydrology*, vol. 643, p. 131 998, 2024.

[12] Q. Xu, Y. Shi, J. Zhao, and X. X. Zhu, “Floodcastbench: A large-scale dataset and foundation models for flood modeling and forecasting”, *Scientific Data*, vol. 12, no. 1, p. 431, 2025.

[13] K. R. Barnhart et al., “Landlab v2. 0: A software package for earth surface dynamics”, *Earth Surface Dynamics*, vol. 8, no. 2, pp. 379–397, 2020. DOI: 10.5194/esurf-8-379-2020.

[14] H. Rezatofighi et al., “Generalized intersection over union: A metric and a loss for bounding box regression”, in *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, Computer Vision Foundation / IEEE, 2019, pp. 658–666. DOI: 10.1109/CVPR.2019.00075.

[15] Y. Wang et al., “A data-driven approach for flood prediction using grid-based meteorological data”, *Hydrological Processes*, vol. 37, no. 3, e14837, 2023.

 ISBN: 978-1-68558-384-2