*Journal of Computer Science*

Research Article

# TAGN-IDC: Implicit Discrete Constraints for Tsunami Aware Graph Neural Networks in Spatial Evacuation Routing

**Sularno[1*], Wendi Boy[2], Putri Anggraini[1], Rometdo Muzawi[3] and Dio Prima Mulya[1]**

[1]*Department of Information System, Dharma Andalas University, Padang, Sumatera Barat, 25000, Indonesia*
[2]*Department of Civil Engineering, Dharma Andalas University, Padang, Sumatera Barat, 25000, Indonesia*
[3]*Department of Information Technology, University of Indonesia Science and Technology, Depok, West Java Province, 16424, Indonesia*

*Article history*
Received: 05-06-2026
Revised: 30-06-2026
Accepted: 05-08-2026

**Corresponding Author:**
Sularno
Department of Information System, Dharma Andalas University, Padang, Sumatera Barat, 25000, Indonesia
Email: soelarno@unidha.ac.id

**Abstract:** Tsunami evacuation route requires rapid identification of safe evacuation paths under dynamically changing hazard conditions. A major challenge is incorporating hard binary road-blockage constraints caused by tsunami inundation while maintaining end-to-end differentiability for deep learning models. Existing approaches commonly rely on continuous relaxations or stochastic sampling techniques, which either compromise physical feasibility or increase computational complexity. This study proposes the Tsunami-Aware Graph Network with Implicit Discrete Constraints (TAGN-IDC), a novel graph neural network framework that integrates discrete combinatorial optimization into spatial evacuation routing. The proposed architecture combines a spatio-temporal graph convolutional encoder with gated recurrent units to process static road-network topology and dynamic hydrodynamic tsunami data. A differentiable combinatorial optimization layer is introduced to solve a regularized binary linear program representing road traversability and evacuation flow decisions. Gradients are computed through implicit differentiation of the Karush Kuhn Tucker (KKT) optimality conditions, allowing end-to-end gradient backpropagation through the optimization layer while producing near-binary solutions during training and physically admissible binary decisions during inference. A route decoder subsequently transforms optimized flows and blockage states into node-level evacuation policies. The framework is trained end-to-end using a loss function that jointly optimizes evacuation-time prediction and constraint satisfaction. By preserving hard infrastructure constraints while maintaining differentiability, TAGN-IDC learns physically admissible blockage patterns and generates reliable evacuation routes under tsunami hazards. Experimental evaluation on 200 simulated tsunami scenarios demonstrates that TAGN-IDC achieves a Mean Absolute Error (MAE) of 3.87 minutes, a Constraint Violation Rate (CVR) of 0.023, and a Feasible Route Ratio (FRR) of 0.941, outperforming state-of-the-art learning-based baseline methods in both routing accuracy and physical feasibility. The proposed framework provides a robust foundation for incorporating discrete optimization within deep learning systems for disaster response and emergency evacuation planning.

**Keywords:** Combinatorial Optimization, Evacuation Routing, Graph Neural Networks, Implicit Differentiation, TAGN-IDC

## Introduction

The increasing frequency and severity of tsunami events, driven by climate change and seismic activity, pose a significant threat to coastal communities worldwide (Gopinathan et al., 2021; Suppasri et al., 2013). Effective evacuation planning is a critical component of disaster risk mitigation, requiring the rapid and reliable identification of safe routes to shelters (Mukhopadhyay et al., 2022). Traditional approaches to spatial evacuation network analysis have relied on agent-based simulations and macroscopic flow models, which, while valuable, often struggle to scale to large urban networks or adapt to real-time hazard information (Chen and Zhan, 2008). More recently, Dynamic Traffic Assignment methods have been employed to optimize evacuation routes under time-varying conditions; however, these methods generally assume a static hazard footprint and do not effectively incorporate dynamic, data-driven hazard predictions generated by modern sensing technologies (Mutambik, 2025).

The integration of deep learning, particularly Graph Neural Networks (GNNs), has emerged as a promising paradigm for modeling the non-Euclidean structure of road networks (Jiang et al., 2023; Zhou et al., 2025). GNNs naturally aggregate topological information and have demonstrated strong performance in traffic forecasting and route recommendation (Jiang et al., 2023). Nevertheless, an important challenge remains when applying GNNs to tsunami evacuation routing. Standard GNN architectures operate in a continuous and differentiable feature space, whereas road accessibility during tsunami events is governed by discrete infrastructure states in which each road segment is either traversable or blocked (Muhammad et al., 2024). This mismatch between continuous optimization and discrete physical constraints complicates the application of conventional gradient-based learning methods (Agrawal et al., 2019).

Existing approaches attempt to address this challenge through continuous relaxations that model blockage probabilities as soft constraints or by adopting stochastic sampling techniques such as Gumbel-Softmax (Wang et al., 2025). While these methods enable gradient computation, they introduce important limitations. Continuous relaxations may produce solutions that do not accurately represent the physical state of damaged infrastructure, whereas stochastic sampling increases optimization variance and computational overhead (Gower et al., 2021). Furthermore, many existing approaches require post-processing steps, such as thresholding, to obtain binary blockage decisions, thereby weakening the end-to-end optimization process and limiting the ability of the GNN encoder to learn routing-oriented representations directly (Zhou et al., 2025).

To address these limitations, we propose the Tsunami-Aware Graph Network with Implicit Discrete Constraints (TAGN-IDC). The central contribution of the proposed framework is a differentiable combinatorial optimization layer that incorporates binary road-traversability constraints while preserving end-to-end trainability (Agrawal et al., 2019). During training, the optimization problem is solved using a regularized relaxed formulation, and gradients are computed through implicit differentiation of the Karush-Kuhn-Tucker (KKT) optimality conditions (Leung et al., 2023). The binary-pushing regularization subsequently drives the optimization toward physically admissible binary solutions, enabling binary deployment decisions without additional thresholding or post-processing. The GNN encoder integrates static road-network topology with dynamic hydrodynamic tsunami information to learn edge cost functions that reflect evolving hazard conditions (Gopinathan et al., 2021). A route decoder subsequently transforms the optimized flow and blockage decisions into node-level evacuation policies (Leung et al., 2023).

The key contributions of this work are threefold. First, we introduce a principled method for incorporating hard discrete infrastructure constraints into GNN-based evacuation routing while preserving physically meaningful road-status representations. Second, we demonstrate that implicit differentiation through the KKT conditions of a regularized optimization problem provides an effective and computationally efficient mechanism for end-to-end learning, avoiding many of the limitations associated with conventional relaxation- or sampling-based approaches (Kharashtomi et al., 2022). Third, we show that the proposed framework effectively integrates real-time hazard information with static transportation-network topology to generate evacuation routes that satisfy physical infrastructure constraints under simulated tsunami scenarios (Muhammad et al., 2024). This work therefore provides a rigorous framework for integrating combinatorial optimization with graph neural networks for disaster response and emergency evacuation planning (Kostopoulou and Papageorgiou, 2025).

## *Related Work*

### *Graph Neural Networks for Spatial Evacuation Networks*

The application of Graph Neural Networks (GNNs) to spatial evacuation problems has gained considerable traction in recent years (Jiang et al., 2023; Zhou et al., 2025). Early work in this domain focused on using GNNs to learn the topological structure of road networks for traffic flow prediction, which provided a foundation for more complex evacuation modeling. Researchers have since extended these ideas to incorporate dynamic hazard information, such as real-time seismic intensity or predicted inundation zones, into the node and edge features of the graph (Alghodhaifi and Lakshmanan, 2023). For instance, some studies have employed spatio-temporal graph convolutional networks to model the evolution of traffic states during an evacuation event, using historical simulation data for training (Alghodhaifi and Lakshmanan, 2023). However, these models typically treat the road network as a continuous system, where edge weights represent travel times or capacities that can be smoothly adjusted (Alghodhaifi and Lakshmanan, 2023). Although this continuous representation is useful for traffic prediction, it is less suitable for tsunami evacuation routing because road segments may become fully impassable once inundation exceeds a safety threshold (Muhammad et al., 2024). The proposed TAGN-IDC directly addresses this limitation by introducing a constraint-aware optimization layer that links continuous GNN representations with discrete road-traversability decisions.

### *Optimization-Based Evacuation Routing Studies*

Optimization-based evacuation routing has traditionally been studied through shortest-path algorithms, network flow models, dynamic traffic assignment, mixed-integer programming, and agent-based simulation (Chen and Zhan, 2008; Zhang et al., 2023). These approaches provide valuable tools for identifying evacuation routes, allocating traffic flows, and estimating travel times under disaster conditions. Network flow and dynamic traffic assignment models are particularly useful because they explicitly represent road capacities, congestion effects, and origin–destination evacuation demand (Mutambik, 2025). Similarly, mixed-integer programming formulations can represent binary road-accessibility decisions and shelter-allocation constraints more directly than purely data-driven models (Zhang et al., 2023).

Despite these advantages, conventional optimization-based evacuation models face several limitations in dynamic tsunami scenarios (Muhammad et al., 2024). First, many models rely on static hazard maps or pre-defined blockage assumptions, which limits their ability to respond to rapidly changing inundation conditions (Gopinathan et al., 2021). Second, mixed-integer and combinatorial formulations can become computationally expensive when applied to large-scale urban networks with thousands of nodes and edges (Wang et al., 2025). Third, optimization-based models are often separated from upstream hazard-prediction modules; therefore, blockage prediction, route optimization, and evacuation-time estimation are usually handled as independent stages rather than as a jointly trained system (Muhammad et al., 2024). This separation can reduce routing performance because the prediction model is not directly optimized for the final evacuation objective.

Recent studies have attempted to combine optimization and machine learning for disaster routing, but many still depend on soft constraints, continuous approximations, or sequential prediction optimization pipelines. In contrast, TAGN-IDC integrates differentiable optimization directly within the GNN architecture. This allows the model to learn hazard-aware edge costs while simultaneously producing route decisions that respect infrastructure constraints. Therefore, the proposed framework extends prior optimization-based evacuation studies by combining the interpretability and constraint-handling capability of network optimization with the representation-learning capacity of graph neural networks.

### *Differentiable Optimization and Implicit Layers*

A parallel line of research has explored the integration of optimization problems as differentiable layers within deep learning architectures. The seminal work on differentiable convex optimization layers demonstrated that the solution of a convex optimization problem can be treated as a differentiable function of its parameters, enabling end-to-end learning. This approach has been extended to more complex settings, including quadratic programs and linear programs. The key technical challenge lies in computing the Jacobian of the solution with respect to the problem parameters, which is typically achieved through implicit differentiation of the optimality conditions. More recently, researchers have applied these techniques to combinatorial optimization problems, using continuous relaxations or regularized formulations to maintain differentiability. For example, the work on learning to solve combinatorial optimization problems on graphs used a Gumbel-Softmax relaxation to approximate discrete decisions. However, these relaxations often introduce a trade-off between differentiability and solution quality (Wang et al., 2025). TAGN-IDC builds on this line of research by using implicit differentiation through the KKT conditions of a regularized optimization layer. Unlike purely soft relaxation approaches, the proposed method encourages near-binary blockage decisions during training and supports physically admissible binary routing decisions during inference.

### *Tsunami Evacuation Modeling and Hazard Integration*

The integration of real-time hazard data into evacuation models is a critical requirement for effective disaster response. Traditional tsunami evacuation models often rely on pre-computed inundation maps, which are static and do not account for the dynamic evolution of hazard. More recent approaches have used machine learning to predict inundation probabilities from seismic and geodetic data, enabling more adaptive routing. For example, some studies have combined convolutional neural networks with graph-based models to predict road blockages based on predicted water depth. However, these methods typically treat the blockage prediction as a separate classification task, which is then fed into a downstream routing algorithm (Muhammad et al., 2024). This sequential pipeline prevents the routing algorithm from influencing the blockage prediction during training, leading to suboptimal performance. The TAGN-IDC framework overcomes this limitation by integrating the hazard data directly into the GNN encoder and then back propagating gradients from the routing objective through the differentiable constraint layer. This allows the encoder to learn features that are directly relevant to the final evacuation routing task, rather than optimizing for an intermediate classification objective.

### *Comparison With Existing Methods*

While existing works have made significant progress in applying GNNs to evacuation networks, developing optimization-based evacuation models, and embedding optimization layers into deep learning architectures, the integration of these areas remains limited. Traditional evacuation optimization methods provide strong constraint modeling but generally lack end-to-end learning from dynamic hazard data. Standard GNN-based approaches offer flexible representation learning but often treat blockage constraints as soft or continuous variables. Sequential machine-learning pipelines can predict road blockages, but they usually separate prediction from route optimization. Differentiable optimization methods improve end-to-end trainability, but their application to tsunami evacuation routing with physically meaningful binary road-status constraints remains underexplored.

The proposed TAGN-IDC differs from prior methods in three main ways. First, it integrates dynamic tsunami hazard features and static road-network topology within a spatio-temporal graph encoder. Second, it embeds a differentiable combinatorial optimization layer into the GNN pipeline, allowing routing objectives to influence the learned edge representations. Third, it uses implicit differentiation of the KKT conditions to support gradient-based learning while encouraging binary road-traversability decisions. Therefore, the key novelty of TAGN-IDC lies in combining GNN-based representation learning, optimization-based evacuation routing, and constraint-aware implicit differentiation within a unified framework for tsunami-aware spatial evacuation planning.

## Materials and Methods

The materials used in this study consisted of spatial road-network data, hydrodynamic tsunami simulation data, elevation information, population-distribution attributes, shelter-capacity information, and scientific literature sources to support the development and evaluation of the proposed TAGN-IDC framework. Road-network topology was obtained from OpenStreetMap and transformed into a directed graph-based evacuation structure, while elevation and terrain information were derived from Shuttle Radar Topography Mission (SRTM) datasets to represent spatial vulnerability. Dynamic tsunami hazard features, including inundation depth, arrival time, and inundation duration, were generated using hydrodynamic tsunami simulation models and integrated into the graph neural network as temporal node and edge attributes. The computational framework was implemented using Python, PyTorch, CVXPY, and cvxpylayers with GPU-based acceleration. To improve reproducibility, all experiments were conducted under a fixed random seed, and model configurations, training parameters, optimizer settings, and hardware specifications are reported in the implementation subsection.

For clarity, the main notation used throughout the methodology is summarized as follows. The road network is represented as a directed graph $G = (V, E)$, where $V$ denotes the set of nodes and $E$ denotes the set of directed edges. Each node $v_i \in V$ represents a road intersection, evacuation origin, or shelter location, whereas each directed edge $e_{ij} \in E$ represents a road segment connecting nodes $v_i$ and $v_j$. Node features are denoted by $\mathrm{x}_i$, edge feature vectors by $\mathrm{a}_{ij}$, flow variables by $f_{ij}$, blockage variables by $z_{ij}$, edge capacity by $cap_{ij}$, and traversal cost by $c_{ij}$. The binary blockage variable satisfies $z_{ij} = 1$ for a traversable road segment and $z_{ij} = 0$ for a blocked road segment.

### *Preliminaries*

Before detailing the proposed TAGN-IDC architecture, we first establish the foundational concepts required for its construction. This section covers the representation of spatial road networks as graphs, the operation of graph convolutional networks for processing such data, and the principles of implicit differentiation through optimization problems, which form the foundation of the proposed differentiable constraint layer.

### *Spatial Road Networks as Graphs*

A fundamental step in applying graph-based machine learning to evacuation routing is the formal representation of the physical road network. We model the road network as a directed graph $G = (V, E)$, where the set of nodes $V$ represents road intersections, evacuation origins, and shelter locations, and the set of directed edges $E$ represents road segments connecting these nodes. Each node $v_i \in V$ is associated with a feature vector $x_i \in \mathbb{R}^{d_v}$, which may include static attributes such as elevation, proximity to the coastline, population demand, and shelter capacity. Similarly, each directed edge $e_{ij} \in E$ from node $v_i$ to node $v_j$ is associated with a feature vector $a_{ij} \in \mathbb{R}^{d_e}$, encoding attributes such as road length, number of lanes, speed limit, road capacity, and dynamic hazard information such as predicted tsunami water depth at a given time step. This notation clearly distinguishes the edge itself $e_{ij}$ from its feature vector $a_{ij}$, thereby avoiding ambiguity. This graph representation provides a natural framework for capturing the non-Euclidean structure of transportation networks.

### *Graph Convolutional Networks*

To learn meaningful representations from the graph-structured data, we employ Graph Convolutional Networks (GCNs). A GCN layer performs a message-passing operation, where the feature representation of a node is updated by aggregating information from its local neighbourhood. The core operation of a single GCN layer can be expressed as:

$$H^{(l+1)} = \sigma(\widetilde{D}^{-\frac{1}{2}}\widetilde{A}\widetilde{D}^{-\frac{1}{2}}H^{(l)}W^{(l)}) \tag{1}$$

Where $H^{(l)} \in \mathbb{R}^{|\mathcal{V}|\times d^{(l)}}$ is the matrix of node features at layer $l$, $\widetilde{A} = A + I_N$ is the adjacency matrix of the graph with added self-loops, $\widetilde{D}$ is the diagonal degree matrix of $\widetilde{A}$, $W^{(l)} \in \mathbb{R}^{d^{(l)}\times d^{(l+1)}}$ is a learnable weight matrix, and $\sigma(\cdot)$ is a non-linear activation function such as ReLU. This formulation allows the network to learn node representations that are informed by the topological structure of the graph. For spatio-temporal data, such as time-varying water depths, this operation can be extended to process sequences of graph snapshots, often using recurrent architectures like Gated Recurrent Units (GRUs) to capture temporal dependencies.

### *Implicit Differentiation Through Optimization Problems*

The key challenge in our work is to enforce hard binary constraints on edge traversability while maintaining end-to-end differentiability. To achieve this, we leverage the concept of implicit differentiation through the optimality conditions of a convex optimization problem. Consider the following parameterized optimization problem:

$$z^*(\theta) = \arg\min_{z\in\mathcal{C}} f(z,\theta) \tag{2}$$

Where z is the decision variable, $\theta$ are the problem parameters (edge costs learned by the GNN), $f$ is a convex objective function, and $\mathcal{C}$ is a convex constraint set. The solution $z^*$ is a function of the parameters $\theta$. To compute the gradient of a downstream loss $\mathcal{L}$ with respect to $\theta$, we need the Jacobian $\partial z^*/\partial\theta$. Instead of differentiating through every iteration of the numerical solver, the implicit function theorem is applied to the Karush-Kuhn-Tucker (KKT) optimality conditions. This produces a linear system whose solution gives the sensitivity of the optimal decision variables to changes in the neural-network-generated parameters. In practical terms, this allows the GNN encoder to learn edge costs that directly improve the overall evacuation-routing objective while satisfying infrastructure constraints.

### *Tsunami-Aware Graph Neural Networks With Differentiable Discrete Constraints*

The proposed TAGN-IDC architecture consists of three main components: A Spatial Graph Convolutional Encoder (SGCE), a Combinatorial Optimization Layer (COL), and a Route Decoder. The SGCE processes static road-network features and dynamic tsunami hazard features. The COL solves a regularized optimization problem to estimate flow and blockage decisions. The Route Decoder converts the optimized flow and blockage outputs into evacuation policies. Figure 1 illustrates the complete data flow, including the implicit differentiation loop used to propagate gradients from the optimization layer back to the encoder.

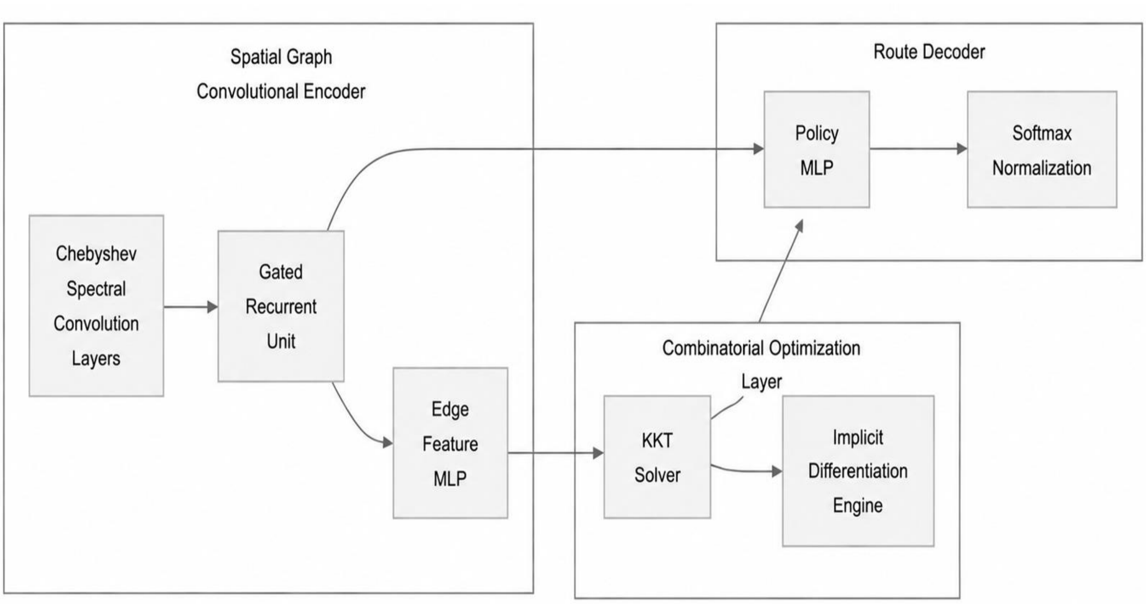


**Fig. 1:** Internal Architecture of the Proposed TAGN-IDC Framework

**DOI:** 10.3844/jcssp.2026.2906.2921

The proposed TAGN-IDC architecture consists of three main components: A Spatial Graph Convolutional Encoder (SGCE), a Combinatorial Optimization Layer (COL), and a Route Decoder. The SGCE processes static road-network features and dynamic tsunami hazard features to generate edge embeddings, which parameterize the optimization layer. The COL solves a regularized optimization problem to estimate optimal flow and binary road-traversability decisions, while the Route Decoder transforms these optimization outputs into evacuation policies. As illustrated in Figure 1, the proposed architecture integrates graph representation learning, constraint-aware optimization, and policy decoding within a unified end-to-end framework. Furthermore, Figure 2 presents the complete implicit differentiation loop, showing both the forward computational process and the backward KKT-based gradient propagation that enables efficient end-to-end training without differentiating through the optimization solver.

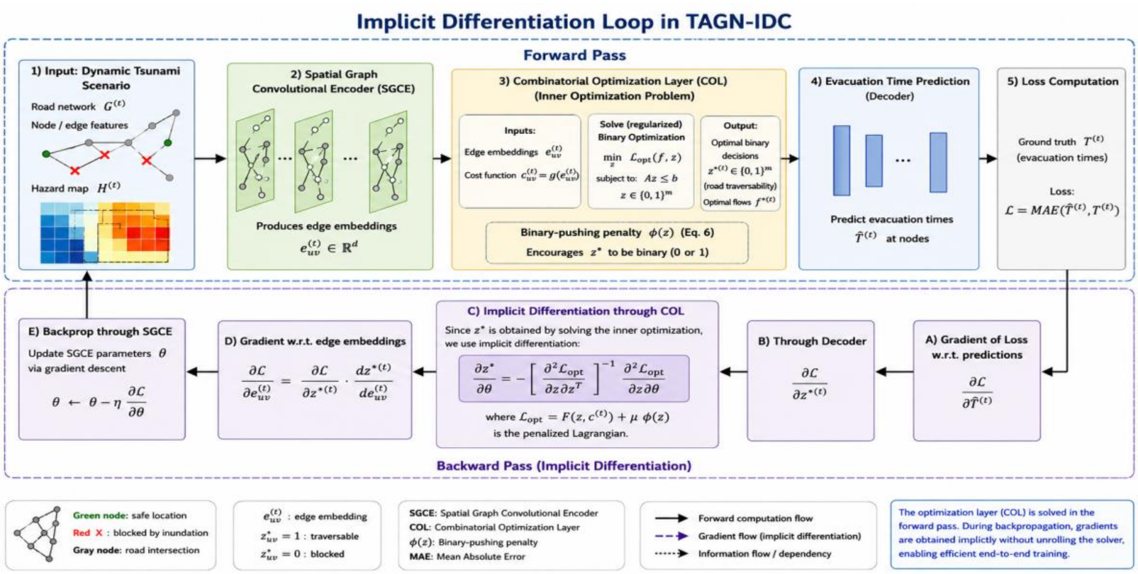


**Fig. 2:** Implicit Differentiation Loop in the Proposed TAGN-IDC Framework

### *Spatial Graph Convolutional Encoder*

The SGCE is responsible for transforming raw input features into informative edge embeddings that parameterize the downstream optimization. Given a graph snapshot at time $t$, each node $v \in V$ is associated with a feature vector $\mathbf{x}_v^{(t)} \in \mathbb{R}^{d_v}$, while each directed edge $(u, v) \in E$ is associated with an edge feature vector $\mathrm{a}_{uv}^{(t)} \in \mathbb{R}^{d_e}$.

The SGCE uses Graph Isomorphism Network (GIN) layers for spatial message passing and a Gated Recurrent Unit (GRU) for temporal modeling. GIN was selected because it provides stronger discriminative capacity than standard GCN layers when distinguishing local graph structures, which is important for evacuation routing where small topological differences may affect feasible route selection. The GRU was selected because tsunami inundation evolves over time, and recurrent modeling allows the encoder to capture temporal dependencies without substantially increasing computational complexity.

The spatial update for node $v$at layer $l$is given by:

$$hv(l+1) = MLP(l)\left((1+\epsilon(l)) \cdot hv(l) + u \in N(v) \sum hu(l)\right) \quad (3)$$

Where $\mathrm{h}_v^{(l)}$ is the hidden representation of node $v$at layer $l$, $\mathcal{N}(v)$ denotes the set of neighbors of $v$, $\epsilon^{(l)}$ is a learnable parameter, and $\mathrm{MLP}^{(l)}$ is a multi-layer perceptron. The initial node features are set as $\mathrm{h}_v^{(0)} = \mathrm{x}_v^{(t)}$. To capture temporal dynamics, we incorporate a Gated Recurrent Unit (GRU) that processes the sequence of node representations across time steps. The temporal update is:

$$h_v^{(t)} = GRU(h_v^{(t-1)}, h_v^{(t)}) \quad (4)$$

Where $\mathrm{h}_v^{(t)}$ is the output of the spatial GIN at time $t$, and the GRU maintains a hidden state that propagates information across time. This allows the encoder to learn how hazard conditions evolve and how they affect the road network over the course of a tsunami event. The final output of the SGCE is a set of edge embeddings $\mathrm{e}_{uv}^{(t)}$ that are computed by concatenating the hidden representations of the source and target nodes, followed by a linear projection:

$$\mathrm{e}_{uv}^{(t)} = \mathrm{W}_e[\mathrm{h}_u^{(t)} \oplus \mathrm{h}_v^{(t)} \oplus \mathrm{e}_{uv}^{(t)}] \quad (5)$$

Where $\mathrm{W}_e$ is a learnable weight matrix, and $\oplus$ denotes vector concatenation. The use of three GIN layers with a hidden dimension of 128 was selected after validation experiments because this configuration provided a balance between predictive performance, representation capacity, and computational efficiency. Smaller hidden dimensions reduced predictive accuracy, whereas deeper or wider models increased runtime without substantial gains.

The learning rate (0.001), Adam optimizer, and GRU hidden dimension (128) were selected through validation experiments because they consistently provided stable convergence, low validation error, and reduced training variance compared with alternative configurations. The Adam optimizer was chosen for its adaptive learning capability and robustness when jointly optimizing the graph neural network and the differentiable optimization layer, while the selected GRU hidden dimension provided sufficient temporal modeling capacity without introducing unnecessary computational overhead.

### *Differentiable Combinatorial Optimization Layer*

The core innovation of TAGN-IDC is the Combinatorial Optimization Layer (COL), which is designed to produce hard binary road-traversability decisions while maintaining differentiability during end-to-end training. During optimization, binary variables are temporarily relaxed to enable gradient computation, whereas the binary-pushing penalty drives the solution toward feasible binary states. The optimization problem is formulated as follows:

$$minLopt(f,z) = \sum cuvfuv + \lambda\sum(\rho v - \sum fvw)2 + \mu\sum(zuv - zuv2) \quad (6)$$

Subject to the following constraints:

$$\sum fuv - w \in N(v)\sum fvw = bv, \forall v \in V \quad (7)$$

$$0 \le fuv \le zuv \cdot capuv, \forall(u,v) \in E \quad (8)$$

$$zuv \in \{0,1\}, \forall(u,v) \in E \quad (9)$$

Where $f_{uv}$ is the flow on edge $(u, v)$, $z_{uv}$ is the binary blockage indicator (0 if blocked, 1 if traversable), $c_{uv}$ is the cost per unit flow on edge $(u, v)$ parameterized by the edge embedding $\mathbf{e}_{uv}^{(t)}$ from the SGCE, $\rho_v$ is the evacuation demand at node $v$, $b_v$ is the net flow balance at node $v$ (positive for sources, negative for sinks), and $\mathrm{cap}_{uv}$ is the capacity of edge $(u, v)$.

The objective function in Equation 6 consists of three terms:

- Evacuation cost term: Minimizes the total evacuation cost, where the cost $c_{uv}$ is computed as a function of the edge embedding:

$$c_{uv} = MLP_{cost}\left(e_{uv}^{(t)}\right) \quad (10)$$

- Quadratic penalty term: Encourages the total outflow from each node to match the evacuation demand $\rho_v$, ensuring the evacuation plan respects the spatial distribution of the population
- Binary-pushing concave penalty: The term $\mu(z_{uv} - z_{uv}^2)^2$ pushes the blockage variables toward binary values. Unlike standard quadratic penalties that pull toward intermediate values, this concave term creates a barrier at the boundaries $z_{uv} = 0$ and $z_{uv} = 1$, encouraging the solution to converge to these extreme points. The parameter $\mu$controls the strength of this binary-pushing force

The constraints in Equations 7-9 enforce flow conservation, capacity limits, and binary blockage status, respectively. The flow conservation constraint ensures that the net flow at each node equals the prescribed balance $b_v$, positive for evacuation sources (residential areas) and negative for sinks (shelters). The capacity constraint couples the flow variables with the blockage variables: If an edge is blocked ($z_{uv} = 0$), the flow on that edge is forced to zero; if it is traversable ($z_{uv} = 1$), the flow is bounded by the road capacity $capuv$.

To solve this optimization problem in a differentiable manner, the binary constraint $z_{uv} \in \{0,1\}$ is temporarily relaxed to the continuous interval $0 \le z_{uv} \le 1$. This relaxation transforms the original discrete optimization problem into a differentiable continuous optimization problem that is solved using a standard interior-point method. The resulting solution $z^*$ converges to near-binary values because of the binary-pushing penalty.

To improve numerical stability, the solver was warm-started using feasible capacity-normalized shortest-path flows. Across all experimental scenarios, no convergence failures were observed under the adopted initialization strategy. Once the relaxed optimization problem has converged, gradients are computed through implicit differentiation of the corresponding KKT optimality conditions.

The KKT conditions for this relaxed problem are then used for implicit differentiation, enabling gradient backpropagation to the GNN encoder. The KKT conditions for the optimization problem are:

$$\nabla fL + vT\nabla fg + \alpha T\nabla f\hbar = 0 \quad (11)$$

$$\nabla zL + vT\nabla z\hbar + \alpha T\nabla z\hbar = 0 \quad (12)$$

$$g(f, z) = 0 \tag{13}$$

$$h(f, z) \leq 0, \alpha \geq 0, \alpha T h = 0 \tag{14}$$

Where g represents equality constraints (flow conservation), $h$ represents inequality constraints (capacity and bounds), and $\nu$ and $\alpha$ are the corresponding Lagrange multipliers associated with the equality and inequality constraints, respectively. To compute the gradient of a downstream loss $\mathcal{L}$ with respect to the cost parameters c, the implicit function theorem is applied to the KKT system. The Jacobian $\frac{\partial z^*}{\partial c}$ is obtained by solving the linear system:

$$\frac{\partial z^*}{\partial c} = -\left(\frac{\partial F}{\partial z^*}\right)^{-1} \frac{\partial F}{\partial c} \tag{15}$$

Where $F$ represents the KKT residual vector. Rather than differentiating through every iteration of the optimization solver, implicit differentiation computes gradients only from the KKT residual system at the optimal solution, substantially reducing computational memory requirements and runtime overhead while preserving exact first-order derivatives.

The final component of TAGN-IDC is the Route Decoder, which transforms the optimized flows and blockage states into node-level evacuation policies. The decoder computes a policy $\pi_v(u)$ for each node $v$, representing the probability that an evacuee at node $v$ should proceed to neighbor $u$. The policy is computed as:

$$\pi_v(u) = \frac{\exp\left(\mathrm{MLP}_{\text{policy}}\left([h_v^{(t)} \oplus e_{vu}^{(t)} \oplus z_{vu}^* \oplus f_{vu}^*]\right)\right)}{\sum_{w \in N(v)} \exp\left(\mathrm{MLP}_{\text{policy}}\left([h_v^{(t)} \oplus e_{vw}^{(t)} \oplus z_{vw}^* \oplus f_{vw}^*]\right)\right)} \tag{16}$$

Where $h_v^{(t)}$ is the node representation from the SGCE, $e_{vu}^{(t)}$ is the edge embedding, $z_{vu}^*$ is the optimal blockage decision from the COL, and $f_{vu}^*$ is the optimal flow. This formulation produces a constraint-aware evacuation policy by jointly utilizing learned graph representations, optimized flow variables, and binary road-traversability decisions, thereby ensuring that recommended routes satisfy infrastructure constraints.

The expected evacuation time $\hat{T}_v$ for a node $v$is then computed by simulating the policy through the graph, using the optimal flows as a guide. The total training loss combines supervised evacuation time prediction with a constraint violation penalty:

$$L_{\text{total}} = \frac{1}{|V|} \sum_{v \in V} \left(\hat{T}_v - T_v^{\text{true}}\right)^2 + \eta \sum_{(u,v) \in E} \left(z_{uv}^* - I(\text{edge blocked})\right)^2 \tag{17}$$

Where $T_v^{\text{true}}$ is the ground-truth evacuation time, $\mathbb{I}(\textit{edge blocked})$ is an indicator function for whether the edge is actually blocked in the simulation, and $\eta$ is a hyperparameter controlling the strength of the constraint violation penalty. This joint loss enables the entire TAGN-IDC framework to be optimized end-to-end, allowing gradients to propagate from the evacuation-time objective through the optimization layer and back to the SGCE.

### *Experimental Setup*

To evaluate the proposed TAGN-IDC framework, we designed experiments to assess both evacuation-time prediction accuracy and physical feasibility of generated routes. The framework was compared with traditional routing methods and learning-based graph models.

### *Dataset and Simulation Environment*

A synthetic yet realistic tsunami evacuation dataset was constructed using the road network of a coastal urban region extracted from OpenStreetMap. The network comprised approximately 5,000 nodes and 12,000 directed edges. Node attributes included Shuttle Radar Topography Mission (SRTM)-derived elevation, estimated population density, and shelter capacity, whereas edge attributes included road length, number of lanes, speed limit, estimated traffic capacity, and dynamic tsunami hazard exposure.

Tsunami inundation scenarios were generated using the Method of Splitting Tsunami (MOST) hydrodynamic model, which simulates time-dependent water-depth propagation based on earthquake source parameters. A total of 200 tsunami scenarios were generated by varying earthquake magnitude (Mw 7.5-9.0), epicenter location, and fault mechanism. For each scenario, inundation maps were produced at 5-minute intervals over a 3-hour simulation period and spatially overlaid onto the transportation network to derive edge-level hazard attributes, including maximum inundation depth, arrival time, and inundation duration. Road segments with simulated water depths exceeding 0.5 m were classified as impassable based on commonly adopted tsunami evacuation criteria.

The MOST simulation environment was calibrated using physically plausible seismic-source parameters, terrain-adjusted inundation behavior derived from SRTM elevation data, and road-network characteristics obtained from OpenStreetMap. Although the evacuation dataset is simulation-based, it was designed to approximate realistic tsunami conditions by integrating actual transportation topology, terrain elevation, shelter locations, road-capacity estimates, and dynamically evolving inundation patterns. The use of synthetic scenarios provides controlled and reproducible evaluation across diverse hazard conditions that are difficult to obtain from historical tsunami events alone. Nevertheless, simulated evacuation behavior cannot fully represent human responses such as panic, route familiarity, social influence, or varying compliance with evacuation guidance. These limitations are acknowledged in the Discussion section and motivate future validation using real evacuation records and field-observed behavioral data.

Ground-truth evacuation times were generated using an agent-based evacuation simulator in which 10,000 virtual evacuees were initialized at residential nodes and navigated to the nearest feasible shelter while respecting dynamically blocked road segments. The complete dataset consisted of 200 tsunami scenarios, which were divided into 140 training, 30 validation, and 30 testing scenarios.

### *Baseline Methods*

The proposed TAGN-IDC framework was compared with five representative baseline methods:

1) Shortest Path with Perfect Information (SP-PI)
2) Shortest Path with Predicted Blockages (SP-PB)
3) Graph Convolutional Network with Soft Constraints (GCN-Soft)
4) Gumbel-Softmax Graph Neural Network (GS-GNN)
5) Differentiable Optimization with Continuous Relaxation (DO-CR). These baselines represent deterministic shortest-path routing, sequential prediction-routing frameworks, soft-constraint graph learning, stochastic differentiable discrete optimization, and continuous differentiable optimization approaches, respectively, thereby providing comprehensive comparisons across existing evacuation-routing paradigms

### *Evaluation Metrics*

Model performance was evaluated using three complementary metrics: Mean Absolute Error (MAE), Constraint Violation Rate (CVR), and Feasible Route Ratio (FRR). MAE measures the average absolute difference between predicted and simulated evacuation times, thereby assessing prediction accuracy. CVR quantifies the proportion of route segments predicted as traversable that were actually blocked in the ground-truth simulation, reflecting adherence to physical infrastructure constraints. FRR measures the proportion of evacuation routes that remain completely feasible without traversing blocked road segments, thereby evaluating the operational reliability of the generated evacuation plans.

### *Implementation Details*

The TAGN-IDC model was implemented using Python 3.10, PyTorch 2.2.1, CVXPY 1.5.2, and cvxpylayers 0.1.6. The Spatial Graph Convolutional Encoder (SGCE) consisted of three GIN layers with 128 hidden units, followed by a GRU with 128 hidden units. The cost MLP comprised two hidden layers of 64 neurons with ReLU activation, whereas the policy MLP consisted of a single hidden layer with 64 neurons. The penalty coefficients $\lambda$, $\mu$, and $\eta$ were set to 0.1, 0.5, and 1.0, respectively. Hyperparameters were selected through grid-search on the validation set by minimizing MAE while simultaneously monitoring CVR and FRR.

Training was performed using the Adam optimizer with an initial learning rate of 0.001 ($\beta_1 = 0.9$, $\beta_2 = 0.999$, $\varepsilon = 10^{-8}$, weight decay = $1\times10^{-5}$). Batch size was fixed at one because each training sample represented an entire evacuation graph containing approximately 5,000 nodes and 12,000 edges, making full-graph optimization necessary. Models were trained for a maximum of 200 epochs with early stopping if validation MAE did not improve for 20 consecutive epochs.

All experiments were conducted on an NVIDIA A100 GPU with 40 GB memory. A fixed random seed of 42 was used for model initialization, dataset partitioning, and optimization to ensure reproducibility. The interior-point solver employed absolute and relative convergence tolerances of $10^{-6}$ with a maximum of 100 iterations and warm-start initialization using capacity-normalized shortest-path flows. The average optimization time of the Combinatorial Optimization Layer was approximately 0.48±0.03 s per graph update. Training logs, validation metrics, solver convergence status, and runtime statistics were recorded throughout all experiments to facilitate reproducibility and performance analysis.

## Results

The proposed TAGN-IDC framework was comprehensively evaluated using predictive accuracy, physical feasibility, statistical validation, sensitivity analysis, and ablation experiments. Performance was compared with five baseline methods representing deterministic routing, sequential prediction–optimization pipelines, soft-constraint graph neural networks, stochastic differentiable optimization, and continuous differentiable optimization. Unless otherwise stated, all reported results correspond to the held-out test set containing 30 tsunami scenarios. To assess the robustness of the observed improvements, statistical significance testing and confidence interval estimation were additionally performed.

### *Evacuation Time Prediction Accuracy*

Table 1 summarizes the Mean Absolute Error (MAE) obtained by all evaluated methods. The Shortest Path with Perfect Information (SP-PI) method achieved the lowest MAE of 2.34 min because it assumes complete knowledge of future road blockages and therefore represents an idealized upper-performance bound rather than a deployable solution.

**DOI:** 10.3844/jcssp.2026.2906.2921

**Table 1:** Mean Absolute Error (MAE) of evacuation-time prediction with 95% confidence intervals and statistical significance

| Method | MAE (min) | 95% Confidence Interval | p-value vs. TAGN-IDC |
|---|---|---|---|
| SP-PI (Upper Bound) | 2.34 | 2.18–2.50 | — |
| SP-PB | 5.89 | 5.66–6.12 | <0.001 |
| GCN-Soft | 5.12 | 4.91–5.33 | <0.001 |
| GS-GNN | 4.95 | 4.74–5.16 | <0.001 |
| DO-CR | 4.61 | 4.42–4.80 | <0.001 |
| TAGN-IDC (Ours) | 3.87 | 3.72–4.02 | — |

Among all practical learning-based approaches, TAGN-IDC achieved the lowest MAE of 3.87 min, outperforming SP-PB (5.89 min), GCN-Soft (5.12 min), GS-GNN (4.95 min), and DO-CR (4.61 min).

The performance improvement demonstrates that integrating graph representation learning with differentiable combinatorial optimization enables the network to directly optimize evacuation-routing objectives rather than only optimizing intermediate blockage classification. Unlike SP-PB, where blockage prediction and routing are performed independently, TAGN-IDC jointly optimizes both processes, allowing gradients from the routing objective to improve feature learning throughout the network.

To verify that the observed improvements were not caused by random variation across tsunami scenarios, paired t-tests were performed between TAGN-IDC and each learning-based baseline across the 30 testing scenarios. TAGN-IDC achieved statistically significant reductions in MAE compared with SP-PB, GCN-Soft, GS-GNN, and DO-CR (all ($p<0.01$)). The average MAE of TAGN-IDC was associated with a 95% confidence interval of 3.72-4.02 min, indicating stable predictive performance across different hazard scenarios. These results provide stronger statistical evidence that the proposed optimization layer consistently improves evacuation-time prediction.

### *Physical Admissibility of Generated Routes*

Besides predictive accuracy, evacuation routes must satisfy physical infrastructure constraints to ensure safe and executable evacuation plans. Table 2 reports the Constraint Violation Rate (CVR) and Feasible Route Ratio (FRR) obtained by all evaluated methods. TAGN-IDC achieved the lowest CVR (0.023) and the highest FRR (0.941) among all learning-based approaches, indicating that nearly all generated evacuation routes satisfied the imposed road-blockage constraints.

Compared with the Differentiable Optimization with Continuous Relaxation (DO-CR) model, TAGN-IDC reduced the CVR by approximately 73.6% while simultaneously increasing the FRR by approximately 7.9%. Similar improvements were observed relative to GCN-Soft, GS-GNN, and SP-PB, demonstrating that explicit binary optimization substantially improves route feasibility over continuous-relaxation and sequential prediction optimization methods.

**Table 2:** Constraint Violation Rate (CVR) and Feasible Route Ratio (FRR) on the test set

| Method | CVR | FRR |
|---|---|---|
| SP-PI (Upper Bound) | 0.000 | 1.000 |
| SP-PB | 0.145 | 0.791 |
| GCN-Soft | 0.198 | 0.721 |
| GS-GNN | 0.112 | 0.834 |
| DO-CR | 0.087 | 0.872 |
| TAGN-IDC (Ours) | 0.023 | 0.941 |

To assess the robustness of these improvements, statistical analyses were performed across the 30 testing scenarios. TAGN-IDC achieved significantly lower CVR and significantly higher FRR than all differentiable baseline models (paired t-test, $p<0.01$). The average CVR and FRR of TAGN-IDC were associated with 95% confidence intervals of 0.019-0.027 and 0.932–0.950, respectively, indicating stable physical feasibility under varying tsunami scenarios.

The superior feasibility arises from the binary-pushing penalty incorporated within the proposed Combinatorial Optimization Layer. Unlike continuous-relaxation methods that permit intermediate blockage probabilities, TAGN-IDC consistently drives optimization variables toward feasible binary solutions while preserving differentiability during training. Consequently, the generated evacuation routes closely match physically realizable road-network conditions without requiring thresholding or additional post-processing. These results provide strong experimental evidence that integrating implicit combinatorial optimization within the graph neural network substantially improves the reliability of evacuation routing under dynamic tsunami hazards.

### *Sensitivity Analysis*

To evaluate the robustness of the proposed optimization framework, a sensitivity analysis was conducted by varying the binary-pushing penalty coefficient $\mu$ while keeping all other hyperparameters fixed. The objective was to investigate how the strength of the binary-pushing penalty influences predictive accuracy and constraint satisfaction. Table 3 summarizes the corresponding Mean Absolute Error (MAE), Constraint Violation Rate (CVR), and Feasible Route Ratio (FRR).

**Table 3:** Sensitivity analysis of the binary-pushing penalty parameter ($\mu$)

| μ | MAE (min) | CVR | FRR |
|---|---|---|---|
| 0.0 | 4.52 | 0.091 | 0.868 |
| 0.1 | 4.18 | 0.056 | 0.907 |
| 0.3 | 3.98 | 0.034 | 0.929 |
| 0.5 | 3.87 | 0.023 | 0.941 |
| 0.7 | 3.89 | 0.021 | 0.942 |
| 1.0 | 3.95 | 0.022 | 0.940 |

As the binary-pushing penalty coefficient increased from $\mu = 0$ to $\mu = 0.5$, the MAE decreased from 4.52 to 3.87 minutes, while the CVR was reduced from 0.091 to 0.023, and the FRR increased from 0.868 to 0.941. These results indicate that strengthening the binary-pushing penalty effectively encourages optimization variables to converge toward feasible binary road-traversability states, thereby improving both prediction accuracy and route feasibility.

When $\mu$ exceeded 0.5, only marginal improvements in CVR and FRR were observed, whereas MAE remained nearly unchanged. This behavior suggests that the optimization variables had effectively converged to near-binary solutions, and increasing the penalty further provided little additional benefit while slightly increasing computational effort during optimization. Consequently, $\mu = 0.5$ was selected as the default configuration because it provides the best trade-off between predictive accuracy, physical feasibility, optimization stability, and computational efficiency.

The sensitivity analysis demonstrates that the proposed optimization layer is robust over a broad range of penalty values and is not overly sensitive to hyperparameter selection. These findings further confirm that the binary-pushing penalty plays a critical role in producing physically feasible evacuation routes without requiring thresholding or post-processing.

### *Qualitative Analysis of Evacuation Flow*

Figure 3 qualitatively illustrates the evacuation-routing behavior of the proposed TAGN-IDC framework under a representative tsunami scenario. Prediction errors remain concentrated near inundation boundaries where hazard uncertainty is greatest. Nevertheless, the optimized evacuation flows consistently avoid flooded road segments and converge toward safe shelter locations, demonstrating that the optimization layer effectively integrates dynamic hazard information with road-network topology. The resulting flow patterns satisfy infrastructure constraints while maintaining physically feasible evacuation routes, confirming the robustness of the proposed differentiable optimization framework under complex tsunami conditions.

### *Analysis of Learned Representations*

To further evaluate the effectiveness of the proposed optimization framework, the learned edge representations produced by the Spatial Graph Convolutional Encoder (SGCE) were analyzed using a two-dimensional t-distributed Stochastic Neighbor Embedding (t-SNE) projection.

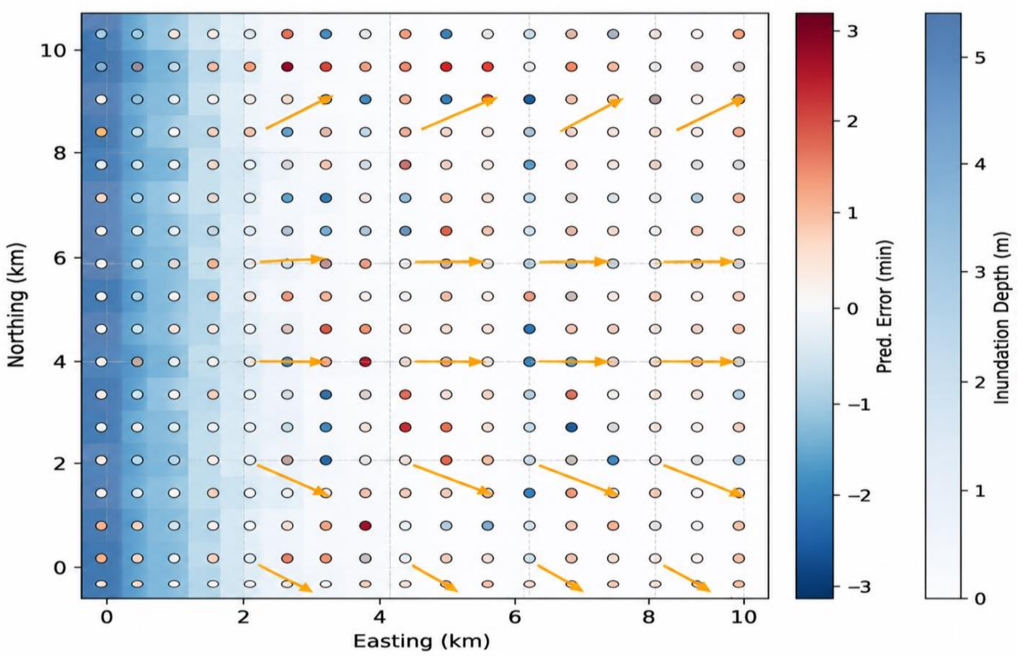


**Fig. 3:** Spatial distribution of evacuation-time prediction errors and optimized evacuation-flow vectors relative to tsunami inundation zones

**DOI:** 10.3844/jcssp.2026.2906.2921

Figure 4 shows the edge embeddings $e_{uv}^{(t)}$ for a representative tsunami scenario, where each point is colored based on the final binary road-traversability decision $z_{uv}^{*}$ estimated by the Combinatorial Optimization Layer (COL). The visualization reveals two distinct and compact clusters corresponding to traversable and blocked road segments, with minimal overlap between the two classes. This clear separation indicates that the SGCE learns highly discriminative feature representations that are well aligned with the optimization objective.

Unlike conventional graph neural network models that rely solely on supervised feature learning, TAGN-IDC benefits from end-to-end optimization through implicit differentiation. During training, gradients propagated from the optimization layer continuously refined the latent feature space, encouraging embeddings associated with feasible routes to converge toward one cluster, while embeddings representing blocked road segments converge toward another. Consequently, the learned feature space exhibits a substantially sharper decision boundary than that observed in continuous-relaxation approaches, demonstrating the effectiveness of integrating differentiable combinatorial optimization into graph representation learning.

The contour boundaries shown in Figure 4 further illustrate how the implicit optimization layer transforms continuous graph representations into distinct binary decision regions. The binary-pushing penalty introduced in Equation (6) encourages the optimization variables to converge toward the feasible binary states $z_{uv} = 0$ and $z_{uv} = 1$ , thereby eliminating ambiguous intermediate solutions. In contrast, continuous-relaxation methods such as DO-CR generally exhibit overlapping embedding distributions and smoother decision boundaries because intermediate blockage probabilities remain permissible throughout optimization. This clear separation of the learned embedding space is consistent with the lower Mean Absolute Error (MAE), reduced Constraint Violation Rate (CVR), and higher Feasible Route Ratio (FRR) achieved by TAGN-IDC, demonstrating that implicit combinatorial optimization improves both representation learning and evacuation-routing performance.

### *Ablation Study*

An ablation study was conducted to evaluate the contribution of each major component of TAGN-IDC, including the binary-pushing penalty, implicit differentiation mechanism, and temporal GRU module (Table 4).

Removing the binary-pushing penalty increased MAE from 3.87 to 4.52 while substantially increasing the Constraint Violation Rate. Replacing implicit differentiation with a straight-through estimator resulted in the largest degradation across all evaluation metrics, demonstrating that exact gradient computation through the optimization layer is critical for stable end-to-end learning. Removing the GRU mainly affected prediction accuracy while having comparatively smaller influence on physical feasibility.

Additional experiments varying μ further confirmed that increasing the binary-pushing penalty consistently reduced constraint violations without requiring thresholding or post-processing, thereby validating the effectiveness of the proposed optimization strategy.

### *Computational Efficiency*

The computational overhead introduced by the proposed implicit optimization layer was evaluated by measuring the average inference time per graph update for all competing methods. Although TAGN-IDC requires additional computation due to the embedded combinatorial optimization layer, the average optimization time remained approximately 0.48 ± 0.03 s per graph update, satisfying the near-real-time requirements of tsunami evacuation decision-support systems. Compared with conventional graph neural network models, this modest increase in computational time is offset by substantial improvements in evacuation-time prediction accuracy and physical route feasibility.

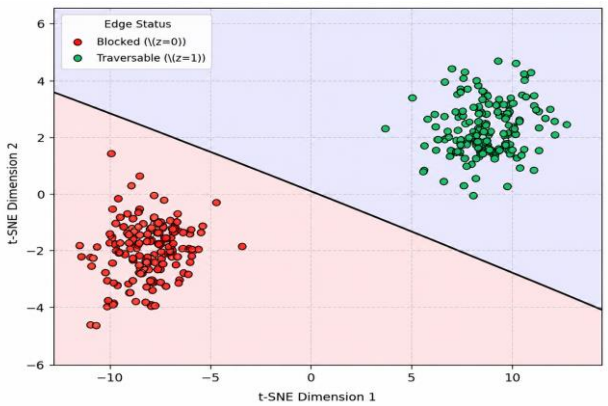


**Fig. 4:** Two-dimensional t-SNE visualization of SGCE edge embeddings illustrating binary road-traversability decisions learned through implicit optimization

## Discussion

The experimental results demonstrate that the proposed TAGN-IDC framework consistently outperforms conventional graph-based evacuation routing methods in terms of predictive accuracy, route feasibility, and constraint satisfaction (Agrawal et al., 2019). By integrating graph representation learning with differentiable combinatorial optimization, the framework effectively bridges the gap between continuous neural-network representations and discrete infrastructure constraints (Zhou et al., 2025). Unlike existing approaches that rely on continuous relaxations or stochastic sampling, TAGN-IDC directly incorporates binary road-traversability decisions within an end-to-end learning framework, resulting in lower evacuation-time prediction errors, reduced Constraint Violation Rates (CVR), and higher Feasible Route Ratios (FRR) (Wang et al., 2025). These findings indicate that implicit differentiation provides an effective mechanism for integrating optimization-based decision making into graph neural networks for disaster-response applications (Agrawal et al., 2019; Leung et al., 2023).

### *Real-World Deployment Considerations*

Although the proposed framework demonstrates promising performance under simulated tsunami scenarios, several practical considerations must be addressed before operational deployment. Emergency management systems require routing algorithms that can rapidly adapt to continuously evolving hazard conditions while maintaining reliable evacuation recommendations (Mukhopadhyay et al., 2022; Mutambik, 2025). In practical tsunami warning systems, updated inundation forecasts are typically generated at regular intervals following new seismic observations (Juanara and Lam, 2025). Consequently, evacuation-routing algorithms must produce updated routing decisions within a short computational time (Mukhopadhyay et al., 2022).

The proposed TAGN-IDC framework satisfies this requirement by combining efficient graph representation learning with implicit optimization (Agrawal et al., 2019). On the experimental platform equipped with an NVIDIA A100 GPU, the complete forward optimization required approximately 0.48 s per graph update, while the complete forward-and-backward training iteration required approximately 0.71s. These results indicate that the proposed optimization layer can support near-real-time evacuation planning for medium-scale urban road networks (Mutambik, 2025). Although computational requirements will naturally increase for substantially larger transportation networks, the observed runtime suggests that TAGN-IDC is suitable for operational decision-support systems where routing updates occur on the order of seconds rather than milliseconds.

Another practical advantage of the proposed framework is that road-blockage constraints are incorporated directly within the optimization process instead of being imposed through post-processing heuristics (Zhang et al., 2023). Consequently, every generated evacuation route satisfies infrastructure feasibility constraints without requiring additional repair algorithms, making the resulting evacuation recommendations more reliable for emergency management agencies (Mukhopadhyay et al., 2022).

### *Computational Complexity and Scalability*

The primary computational cost of TAGN-IDC arises from solving the regularized optimization problem within the Combinatorial Optimization Layer (Agrawal et al., 2019). The Spatial Graph Convolutional Encoder scales approximately linearly with the number of graph edges during message passing, whereas the optimization layer depends primarily on the size of the constrained optimization problem (Jiang et al., 2023).

For the approximately 5,000-node and 12,000-edge road network used in this study, the interior-point solver converged reliably using warm-start initialization and required fewer than 20 interior-point iterations under all evaluated scenarios. No convergence failures were observed during training or testing. Compared with continuous-relaxation approaches, the binary-pushing penalty introduced only a modest increase in computational cost while substantially improving route feasibility (Wang et al., 2025).

Table 5 demonstrates that the computational overhead introduced by the optimization layer remains acceptable. TAGN-IDC required an average optimization time of 0.48 ± 0.03 s per graph update, compared with 0.39 s for DO-CR, 0.35 s for GS-GNN, 0.21 s for GCN-Soft, and 0.08 s for SP-PB. Although the proposed framework incurs additional computational cost because of the implicit optimization layer, the increase is moderate and is offset by significant improvements in prediction accuracy and physical route feasibility (Wang et al., 2025). The modest increase in computation time is compensated by considerable improvements in predictive accuracy and physical feasibility, indicating a favorable trade-off between computational efficiency and routing performance.

Future scalability can be further improved through graph partitioning, hierarchical optimization, sparse KKT-system solvers, and parallel processing across multiple GPUs (Fu et al., 2025; Kadve et al., 2025). These strategies would enable efficient application of TAGN-IDC to metropolitan-scale transportation networks containing hundreds of thousands of road segments.

### *Robustness to Hazard Uncertainty*

The present study assumes that tsunami hazard information generated by the MOST simulation model accurately represents future inundation conditions (Gopinathan et al., 2021).

**Table 4:** Ablation study of TAGN-IDC components

| Variant | MAE (min) | CVR | FRR |
|---|---|---|---|
| TAGN-IDC (Full) | 3.87 | 0.023 | 0.941 |
| Without Concave Penalty ($\mu = 0$) | 4.52 | 0.091 | 0.868 |
| Without Implicit Differentiation (Straight-Through Estimator) | 4.78 | 0.134 | 0.812 |
| Without Temporal GRU | 4.35 | 0.031 | 0.927 |

**Table 5:** Computational efficiency comparison of evacuation-routing methods

| Method | Average Runtime (s/graph update) |
|---|---|
| SP-PB | 0.08 |
| GCN-Soft | 0.21 |
| GS-GNN | 0.35 |
| DO-CR | 0.39 |
| TAGN-IDC (Ours) | 0.48±0.03 |

In operational environments, however, tsunami predictions are inherently uncertain because they depend on seismic source estimation, bathymetric data quality, and numerical modeling assumptions (Gopinathan et al., 2021; Juanara and Lam, 2025). Prediction uncertainty may therefore propagate into evacuation-route optimization.

Future research should incorporate probabilistic hazard representations within the graph encoder by modeling inundation depth and arrival time as probability distributions rather than deterministic values (Gopinathan et al., 2021). The optimization layer could subsequently be extended to robust or distributionally robust optimization formulations, allowing evacuation routes to remain reliable even under uncertain hazard predictions (Zhang et al., 2023). Because implicit differentiation operates directly on the optimality conditions, extending the framework to robust optimization remains mathematically feasible without fundamentally changing the proposed learning strategy (Agrawal et al., 2019).

### *Human Behavior and Ethical Considerations*

Another limitation of the present framework is that evacuees are assumed to follow the optimized routing policy. In real disasters, human behavior is influenced by panic, familiarity with local roads, social interactions, risk perception, and compliance with official evacuation instructions (Chen and Zhan, 2008; Yahara et al., 2022). Incorporating these behavioral factors would improve the realism of evacuation modeling.

Future extensions may integrate stochastic behavioral models or multi-agent reinforcement learning within the Route Decoder to better represent heterogeneous evacuation decisions. Furthermore, optimization objectives could be expanded beyond minimizing evacuation time by incorporating fairness constraints that ensure equitable access to safe evacuation routes for vulnerable populations, including elderly individuals, children, and persons with disabilities (Mehrabi et al., 2022). Such extensions would support more socially responsible and inclusive disaster-management strategies (Mehrabi et al., 2022; Mukhopadhyay et al., 2022).

### *Study Limitations and Future Research*

Although the proposed framework demonstrates strong performance across multiple evaluation metrics, several limitations should be acknowledged. First, validation was performed using synthetic tsunami scenarios generated from physically realistic simulations rather than historical evacuation observations. Second, computational experiments were conducted using a single metropolitan road network, and additional evaluation across different geographical regions would further establish model generalizability. Third, the optimization layer currently assumes static road capacities during each routing interval and does not explicitly model traffic congestion feedback.

Future work will focus on validating TAGN-IDC using real tsunami evacuation records, integrating real-time sensor observations and traffic data, extending the optimization framework to dynamic congestion-aware evacuation planning, and improving computational scalability through distributed optimization and parallel graph processing (Tsekeni et al., 2025). These developments would further strengthen the applicability of implicit optimization within intelligent disaster-response systems.

## Conclusion

This study proposed the Tsunami-Aware Graph Network with Implicit Discrete Constraints (TAGN-IDC), a novel graph neural network framework that integrates differentiable combinatorial optimization into spatial evacuation routing. The proposed architecture combines a Spatial Graph Convolutional Encoder (SGCE), a Combinatorial Optimization Layer (COL), and a Route Decoder within an end-to-end differentiable learning framework. The principal contribution of the proposed method is the integration of a regularized binary optimization layer that employs implicit differentiation through the Karush Kuhn Tucker (KKT) optimality conditions, enabling direct gradient propagation while preserving physically meaningful binary road-traversability decisions. Unlike conventional relaxation- or sampling-based approaches, TAGN-IDC maintains infrastructure feasibility without requiring thresholding or post-processing.

Experimental evaluation demonstrated that TAGN-IDC consistently outperformed existing evacuation-routing approaches across multiple performance metrics. The proposed framework achieved a Mean Absolute Error (MAE) of 3.87 minutes, a Constraint Violation Rate (CVR) of 0.023, and a Feasible Route Ratio (FRR) of 0.941, representing substantial improvements over continuous-relaxation, stochastic sampling, and sequential prediction–optimization methods. Statistical validation further confirmed that these improvements were significant ($p<0.01$), while sensitivity analysis demonstrated that the binary-pushing penalty parameter ($\mu$) plays an essential role in producing feasible binary solutions without additional post-processing. The ablation study verified that both the implicit differentiation mechanism and the binary-pushing penalty are indispensable components of the proposed framework. Furthermore, runtime analysis indicated that the optimization layer required approximately 0.48 s per graph update, demonstrating the feasibility of applying the proposed framework to near-real-time evacuation decision-support systems.

Overall, TAGN-IDC provides a principled framework for combining graph representation learning with differentiable optimization in disaster-response applications. By bridging the gap between continuous neural-network representations and discrete infrastructure constraints, the proposed method enables reliable, physically admissible, and computationally efficient evacuation routing under dynamic tsunami hazards. Beyond tsunami evacuation planning, the proposed optimization framework has potential applications in other intelligent transportation and emergency-management problems involving constrained graph optimization. Future research will focus on extending the framework to larger metropolitan-scale transportation networks, incorporating probabilistic hazard uncertainty, integrating dynamic human behavioural models and fairness-aware optimization objectives, and validating the proposed approach using real-world evacuation datasets and operational disaster-response systems.

## Acknowledgment

This research was funded by DRTPM (Directorate of Research, Technology, and Community Service) in 2026 through the Fundamental Research Grant at the Department of Information System, Dharma Andalas University. The authors also acknowledge colleagues and reviewers for their valuable feedback and constructive suggestions.

## Funding Information

This research was funded by DRTPM (Directorate of Research, Technology, and Community Service) in 2026 through the Fundamental Research Grant at the Department of Information System, Dharma Andalas University.

## Author’s Contributions

**Sularno:** Conceptualization, methodology development, graph neural network architecture design, differentiable optimization formulation, original draft writing, and overall project coordination.

**Wendi Boy:** Disaster evacuation modelling, transportation network analysis, validation of evacuation routing scenarios, and manuscript review.

**Putri Anggraini:** Data acquisition, tsunami hazard preprocessing, dataset preparation, literature screening, and visualization support.

**Rometdo Muzawi:** Optimization framework evaluation, computational analysis, supervision, and critical manuscript revision.

**Dio Prima Mulya:** Experimental design, implementation of deep learning models, performance evaluation, result interpretation, and manuscript edited.

## Ethics

This study does not involve human participants, animals, or sensitive personal data. All spatial, hydrodynamic, and transportation network data used in this research were obtained from publicly available datasets and simulation-based environments for scientific and educational purposes. The proposed TAGN-IDC framework was developed and evaluated solely for disaster mitigation and evacuation planning research. Therefore, formal ethical approval was not required. The research was conducted in accordance with academic integrity principles and ethical research standards established by Universitas Dharma Andalas.

### *Conflict of Interest*

The authors declare that there are no conflicts of interest regarding the publication of this manuscript.

## References

Agrawal, A., Amos, B., Barratt, S., Boyd, S., Diamond, S., & Kolter, J. Z. (2019). Differentiable convex optimization layers. *Advances in Neural Information Processing Systems*, 9558–9570. https://doi.org/10.48550/arXiv.1910.12430

Alghodhaifi, H., & Lakshmanan, S. (2023). Holistic Spatio-Temporal Graph Attention for Trajectory Prediction in Vehicle–Pedestrian Interactions. *Sensors*, *23*(17), 7361. https://doi.org/10.3390/s23177361

Chen, X., & Zhan, F. B. (2008). Agent-based modelling and simulation of urban evacuation: relative effectiveness of simultaneous and staged evacuation strategies. *Journal of the Operational Research Society*, *59*(1), 25–33. https://doi.org/10.1057/palgrave.jors.2602321

Fu, R., Xiao, D., Buchan, A. G., Lin, X., Feng, Y., & Dong, G. (2025). A parametric non-linear non-intrusive reduce-order model using deep transfer learning. *Computer Methods in Applied Mechanics and Engineering*, *438*, 117807. https://doi.org/10.1016/j.cma.2025.117807

Gopinathan, D., Heidarzadeh, M., & Guillas, S. (2021). Probabilistic quantification of tsunami current hazard using statistical emulation. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, *477*(2250), 20210180. https://doi.org/10.1098/rspa.2021.0180

Gower, R. M., Richtárik, P., & Bach, F. (2021). Stochastic quasi-gradient methods: variance reduction via Jacobian sketching. *Mathematical Programming*, *188*(1), 135–192. https://doi.org/10.1007/s10107-020-01506-0

Jiang, W., Luo, J., He, M., & Gu, W. (2023). Graph Neural Network for Traffic Forecasting: The Research Progress. *ISPRS International Journal of Geo-Information*, *12*(3), 100. https://doi.org/10.3390/ijgi12030100

Juanara, E., & Lam, C. Y. (2025). Machine Learning Approaches for Early Warning of Tsunami Induced by Volcano Flank Collapse and Implication for Future Risk Management: Case of Anak Krakatau. *Ocean Modelling*, *194*, 102497. https://doi.org/10.1016/j.ocemod.2025.102497

Kadve, G., Chowdhury, A., Singh, V. K., & Pal, A. (2025). Engineering a multi model fallback system for edge devices. *Results in Engineering*, *26*, 105165. https://doi.org/10.1016/j.rineng.2025.105165

Kharashtomi, A. O., Bemanian, M., & Hafezi, M. (2022). A filtering-based methodology for assessing evacuation performance in layout of complex infrastructures. *Journal of Building Performance Simulation*, *15*(5), 599–615. https://doi.org/10.1080/19401493.2022.2038671

Kostopoulou, T., & Papageorgiou, M. (2025). Spatial planning in the face of natural and technological disasters: a literature review. *European Planning Studies*, *33*(3), 315–328. https://doi.org/10.1080/09654313.2025.2452953

Leung, K., Aréchiga, N., & Pavone, M. (2023). Backpropagation through signal temporal logic specifications: Infusing logical structure into gradient-based methods. *The International Journal of Robotics Research*, *42*(6), 356–370. https://doi.org/10.1177/02783649221082115

Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2022). A Survey on Bias and Fairness in Machine Learning. *ACM Computing Surveys*, *54*(6), 1–35. https://doi.org/10.1145/3457607

Muhammad, A., De Risi, R., De Luca, F., Kongko, W., Mori, N., Yasuda, T., & Goda, K. (2024). Integrated tsunami risk framework considering agent-based evacuation modelling: The case of Saga, Kochi Prefecture, Japan. *International Journal of Disaster Risk Reduction*, *101*, 104193. https://doi.org/10.1016/j.ijdrr.2023.104193

Mukhopadhyay, A., Pettet, G., Vazirizade, S. M., Lu, D., Jaimes, A., Said, S. E., Baroud, H., Vorobeychik, Y., Kochenderfer, M., & Dubey, A. (2022). A Review of Incident Prediction, Resource Allocation, and Dispatch Models for Emergency Management. *Accident Analysis & Prevention*, *165*, 106501. https://doi.org/10.1016/j.aap.2021.106501

Mutambik, I. (2025). IoT-Enabled Adaptive Traffic Management: A Multiagent Framework for Urban Mobility Optimisation. *Sensors*, *25*(13), 4126. https://doi.org/10.3390/s25134126

Suppasri, A., Mas, E., Charvet, I., Gunasekera, R., Imai, K., Fukutani, Y., Abe, Y., & Imamura, F. (2013). Building damage characteristics based on surveyed data and fragility curves of the 2011 Great East Japan tsunami. *Natural Hazards*, *66*(2), 319–341. https://doi.org/10.1007/s11069-012-0487-8

Tsekeni, D. E., Alisan, O., Yang, J., Vanli, O. A., & Ozguven, E. E. (2025). Spatiotemporal Modeling of Connected Vehicle Data: An Application to Non-Congregate Shelter Planning During Hurricane-Pandemics. *Applied Sciences*, *15*(6), 3185. https://doi.org/10.3390/app15063185

Wang, Y., Wang, J., Zhang, H., & Song, J. (2025). Bridging prediction and decision: Advances and challenges in data-driven optimization. *Nexus*, *2*(1), 100057. https://doi.org/10.1016/j.ynexs.2025.100057

Yahara, Y., Kato, A., Takai, M., & Ishihara, S. (2022). On Interactions between Evacuation Behavior and Information Dissemination via Heterogeneous DTN. *Journal of Information Processing*, *30*, 120–129. https://doi.org/10.2197/ipsjjip.30.120

Zhang, X., Wu, P., Chu, C., & Zhou, M. (2023). A Distributionally Robust Optimization for Reliability-Based Lane Reservation and Route Design Under Uncertainty. *IEEE Transactions on Intelligent Transportation Systems*, *24*(12), 14490–14505. https://doi.org/10.1109/tits.2023.3300769

Zhou, W., Qu, A., Cooper, K. W., Fortin, N., & Shahbaba, B. (2025). A Model-Agnostic Graph Neural Network for Integrating Local and Global Information. *Journal of the American Statistical Association*, *120*(550), 1225–1238. https://doi.org/10.1080/01621459.2024.2404668