



## OPEN ACCESS

## EDITED BY

Huthaifa I. Ashqar,  
Arab American University, Palestine

## REVIEWED BY

Pavel Loskot,  
The Zhejiang University-University of Illinois at  
Urbana-Champaign Institute, United States  
Sen Wei,  
Chang'an University, China

## \*CORRESPONDENCE

Ömer Kaya,  
✉ omer.kaya@erzurum.edu.tr

RECEIVED 11 December 2025

REVISED 11 January 2026

ACCEPTED 19 January 2026

PUBLISHED 11 February 2026

## CITATION

Kaya Ö (2026) Design and feasibility study of a smart speed hump for selective urban speed management: risk-informed deployment via conditional generative modelling. *Front. Future Transp.* 7:1765920. doi: 10.3389/ffutr.2026.1765920

## COPYRIGHT

© 2026 Kaya. This is an open-access article distributed under the terms of the [Creative Commons Attribution License \(CC BY\)](#). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

# Design and feasibility study of a smart speed hump for selective urban speed management: risk-informed deployment via conditional generative modelling

Ömer Kaya\*

Transportation Department, Engineering and Architecture Faculty, Erzurum Technical University, Erzurum, Türkiye

Urban speed management in developing countries frequently relies on fixed physical speed humps. While effective for compliance, these devices can reduce comfort for compliant drivers, increase structural loads on heavy vehicles, and complicate winter maintenance operations. This study develops and evaluates a selective speed-management approach centred on an adaptive speed hump concept that remains flush for compliant drivers and actuates only when speeding is likely or detected. To support deployment decisions in data-scarce settings, a conditional generative (parametric) decision-support module is used to generate synthetic speed distributions from roadway and scenario attributes based on sparse observations. Segment-level speed-violation risk is estimated and combined with additional criteria to compute a Speed-Calming Suitability Index (SSI) for site prioritization. A low-cost laboratory prototype with real-time speed detection and a servo-driven movable surface demonstrates selective actuation at a single point. The modelling workflow produces actionable risk and SSI-based prioritization for targeted traffic calming, and the prototype demonstrates the feasibility of selective actuation. Together, these components support risk-informed selection of candidate locations and practical implementation of a selective traffic-calming mechanism. The results suggest that conditional generative modelling can support sustainable mobility by enabling risk-informed deployment of adaptive traffic-calming infrastructure under data scarcity. Here, “generative” denotes distributional speed sampling for risk inference; the implementation is a lightweight parametric conditional generator (mean plus dispersion) rather than GAN/VAE/diffusion-style architectures.

## KEYWORDS

adaptive traffic calming, probabilistic generative modelling, risk mapping, synthetic speed modelling, urban mobility safety

## 1 Introduction

Speed humps used to ensure urban traffic safety remain a common practice worldwide (Šaric et al., 2025). In developing countries like Turkey, such physical speed reducers are frequently preferred due to their low cost and widespread availability (Afukaar, 2003). However, due to their fixed and static structure, traditional speed humps have various problems that negatively affect both driving comfort and road safety (Alshabibi, 2025).

These problems can pose serious risks not only to drivers and passengers, but also to the operation of emergency vehicles and the sustainability of road infrastructure (Kosakowska, 2022).

Indeed, a study conducted in Turkey found that only one of 10 different types of speed humps met comfort standards, while the remaining ones caused significant jolts and discomfort for drivers (Yurtbař et al., 2022). It was reported that 87.67% of the vertical accelerations occurring when vehicles passed over speed humps were in the “uncomfortable” category. These jolts cause physical discomfort for drivers and passengers, particularly spinal health, and can also lead to premature fatigue in the chassis and suspension systems of high-performance vehicles such as commercial vehicles, ambulances, or heavy-duty trucks (Wilder et al., 1994). Speed humps vary in design across application areas, and proper slope, height, or positioning criteria are often not met (Golbabaei et al., 2025). While speed humps with non-standard heights or angles can effectively control speed, they can also damage vehicles and road infrastructure. Particularly in cold-climate regions such as Erzurum, snowploughs frequently collide with raised humps during winter maintenance, causing significant damage and increasing long-term rehabilitation costs.

In response to these limitations, interest in adaptive and intelligent speed hump systems has increased in recent years (Miracle et al., 2021; Lin and Ho, 2022; Džambas et al., 2023). Dynamic systems capable of recognizing vehicles, processing speed data, and generating a physical barrier only when a violation is detected offer substantial improvements over static designs. For example, RFID-based identification systems have been developed to prioritize special vehicles such as ambulances (Mei and Wang, 2021). Other studies demonstrate that dynamic hump height adjustments can improve cabin comfort and reduce ride harshness for emergency vehicles by up to 20% (Li et al., 2019). Accelerometer-based detection, computer vision, and sensor fusion techniques are also increasingly used to recognize humps, verify positioning, and support autonomous vehicle navigation (Chen et al., 2019; Sirbu et al., 2021).

While these technological advancements represent progress in the design of speed-calming devices, they still rely heavily on localized sensing, which depends on the physical installation of devices and the availability of high-quality empirical speed data. Yet, many urban areas, particularly in developing countries, lack detailed, high-resolution traffic speed datasets needed to systematically evaluate where speed humps should be placed, how frequently violations occur, or how different environmental conditions influence speeding behaviour. This data scarcity limits municipalities' ability to make informed, network-wide decisions for speed management. Accordingly, the main objective of this study is to develop and evaluate a selective urban speed-management approach based on an adaptive speed hump concept. In this framing, Generative AI (GenAI) is used only as a decision-support tool to estimate network-wide speeding risk and prioritize candidate locations under data scarcity, while the laboratory prototype demonstrates how selective actuation operates at a single installation point.

Generative modelling has recently gained attention in transportation because it enables distributional simulation and scenario sampling when empirical data are limited. GenAI techniques can generate synthetic samples from learned conditional distributions, supporting probabilistic inference and scenario exploration in data-poor environments. Recent studies highlight the role of generative models in traffic forecasting, safety analysis, and

network optimization, enabling agencies to analyse infrastructure performance even when empirical data is limited (Aceto et al., 2024; Ray et al., 2024; Basaran and Dressler, 2025; Bibri and Huang, 2025; Kim et al., 2025). By producing synthetic speed distributions conditioned on roadway and environmental characteristics, GenAI can reveal latent risk patterns, support proactive safety analysis, and improve the prioritization of traffic-calming interventions.

Importantly, GenAI is not used as an additional objective or for controlling the prototype. In this study, GenAI is employed solely as a generative (probabilistic) speed modelling tool under data scarcity. The key requirement is to estimate the probability of speeding events, which depends on the full conditional distribution of speeds rather than a single point estimate. The generated speed samples enable (i) segment-level speeding probability estimation, (ii) scenario-based risk surfaces, and (iii) propagation of uncertainty into the Speed-Calming Suitability Index (SSI)-based prioritization used for deployment planning. In this manuscript, ‘generative’ specifically refers to learning a conditional speed distribution  $p(v|X)$  and sampling synthetic speeds to estimate exceedance probabilities such as  $P(v > SL + 5)$ . While high-capacity GenAI architectures (GAN/VAE/diffusion/autoregressive models) can also be used for conditional generation, we adopt a lightweight parametric generator (MLP mean + residual-based dispersion) because it is transparent, computationally efficient, and sufficient for risk inference under sparse observations.

Building on these advancements, this study integrates generative modelling with urban speed management. A synthetic urban road network is constructed, and ground truth speeds are assigned using scenario-based behavioural rules. Using only a sparse subset of edge-level observations, a generative regression model learns the conditional distribution of speeds and produces synthetic samples to estimate speed violation risk across the entire network. These risk values are then combined with pedestrian exposure, traffic volume, equity considerations, and winter maintenance costs to compute a SSI, enabling data-informed selection of optimal locations for adaptive speed control infrastructure.

In parallel, a prototype smart speed hump system equipped with ultrasonic sensors and a servo-driven movable surface is developed to demonstrate how the model's outputs can translate into real-world applications. Although the physical system addresses operational needs at the micro-scale, the generative modelling framework introduces a macro-scale, citywide planning capability, bridging the gap between intelligent device design and network-level transport safety strategies. This study focuses specifically on speed management and traffic calming, bridging risk-informed planning and selective physical actuation. In this paper, we use the term speed hump to denote a vertical traffic-calming device with a smooth, rounded profile intended to reduce operating speeds. The device proposed here is an adaptive (actuated) speed hump, meaning it remains flush with the pavement for compliant speeds and temporarily introduces a raised profile only when a violation is detected.

## 1.1 Research gap

Although existing studies provide important insights into optimizing the geometry of speed humps, improving speed

transition control in autonomous vehicles, and exploring alternative physical design concepts, several practical limitations of current speed-calming practices remain insufficiently addressed. In particular, traditional speed humps still cause substantial discomfort for compliant drivers, impose unnecessary dynamic loads on heavy vehicles, and increase operational challenges during winter maintenance. The literature offers very limited analysis of how winter maintenance equipment, such as snowploughs, damages speed humps, contributes to long-term deterioration of pavement layers, and increases municipal maintenance expenditures.

Moreover, current research generally assumes the availability of detailed traffic speed data when evaluating the need for speed-calming infrastructure. However, many cities, especially in developing regions, face significant data scarcity and lack the fine-grained speed measurements required to understand where and under what conditions speeding behaviour occurs. This creates a methodological gap in determining the optimal location and necessity of speed-control devices at a network-wide scale. Traditional engineering approaches focus on localized measurements, whereas systematic, city-scale assessment remains underdeveloped.

At the same time, while intelligent speed hump concepts have been proposed, existing prototypes predominantly focus on micro-scale actuation and sensing, lacking integration with broader analytical frameworks for evaluating urban speeding risk. There remains a clear need for a sustainable, low-cost system that both (i) measures vehicle speed using affordable, embedded sensing technologies, and (ii) delivers targeted physical warnings only to vehicles violating speed limits. However, this need must now be complemented by network-scale analytical tools capable of identifying high-risk segments before physical deployment.

This reveals a two-fold research gap:

- i. The absence of distributional (probabilistic) modelling workflows that can estimate threshold-exceedance risk in data-scarce cities without dense segment-level speed measurements, and
- ii. The lack of adaptive smart speed hump systems whose deployment is coupled to a risk-informed site-selection process rather than point-wise prototype demonstrations. Together, these gaps highlight the necessity of combining probabilistic generative synthetic speed modelling with intelligent, selectively actuated smart speed hump infrastructure, enabling a new generation of sustainable and context-aware urban speed management systems.

Together, these gaps highlight the necessity of combining probabilistic generative synthetic speed modelling with intelligent, selectively actuated speed-calming infrastructure, enabling a new generation of sustainable and context-aware urban speed management systems.

## 1.2 Contribution of the study

To achieve the study's single overarching objective of selective speed management, we develop and evaluate two

tightly coupled components: (i) a probabilistic generative decision-support module for risk-informed site prioritization under data scarcity, and (ii) a prototype adaptive speed hump that demonstrates selective actuation at the point of installation.

First, the study introduces a conditional probabilistic generative modelling framework capable of estimating speed violation risk at the network scale on a synthetic testbed designed to mimic heterogeneous urban roadway conditions under data scarcity. By generating synthetic speed distributions from roadway, environmental, and temporal attributes, the proposed model enables network-wide identification of high-risk segments, overcoming a long-standing limitation of traditional engineering approaches that rely solely on localized measurements. The framework further integrates multiple planning-relevant factors, such as pedestrian exposure, traffic volume, equity considerations, and winter maintenance cost, into a comprehensive SSI that supports data-driven and sustainable site selection for traffic-calming infrastructure.

Second, the study develops a low-cost smart speed bump prototype that operationalizes selective and adaptive speed control. The implemented system incorporates real-time speed detection via ultrasonic sensors, a servo-driven actuation mechanism controlled by an Arduino-based controller, and a solar-powered supply that enables energy-independent operation. Unlike traditional fixed speed bumps, this mechanism creates a physical barrier only when a vehicle exceeds the speed limit, thereby maintaining comfort for compliant drivers while reducing unnecessary structural loads on vehicles and pavements. The adaptive design also mitigates the well-documented issue of winter maintenance equipment damaging fixed installations, offering a practical and climate-resilient alternative for municipalities.

Together, these contributions provide a unified approach that links generative, data-informed planning with intelligent, selectively actuated roadside infrastructure. The integration of macro-scale risk mapping and micro-scale hardware implementation offers a scalable pathway for smart transportation systems, particularly in developing urban environments where resource constraints and data scarcity pose persistent challenges.

## 2 Materials and methods

The study comprises two linked modules: (i) a conditional probabilistic modelling module that generates synthetic speed samples to estimate exceedance risk and compute SSI for deployment prioritization on a controlled grid testbed, and (ii) a prototype smart speed bump that demonstrates selective actuation at a single installation point. The modelling module answers where to deploy; the prototype addresses how to implement selective actuation.

At a high level, the modelling module (Sections 2.1–2.3) takes as input the road network, segment attributes, scenario descriptors, and speed limits, and outputs (i) a segment-level speeding risk map  $R_e$  and (ii) an SSI-based top-K shortlist of candidate deployment locations. The detailed computational steps and notation are provided in the step-by-step workflow below.

## 2.1 Synthetic road network generation and ground-truth speed modelling

To evaluate urban speeding behaviour under controlled and reproducible experimental conditions, a synthetic grid-based urban road network was constructed. The network includes roadway attributes such as road type, number of lanes, slope category, winter severity, and adjacent land-use structure. Four operational scenarios were modelled: peak hour, night-time, weekend, and winter conditions, representing typical temporal variations in urban traffic dynamics.

The grid network is used as a controlled experimental testbed rather than as a digital twin of a specific city. Its purpose is to provide (i) a transparent and reproducible environment with known ground-truth speed-generation rules, and (ii) heterogeneous segment attributes to stress test the GenAI pipeline under sparse observations. Therefore, Figure 4 should be interpreted as a proof-of-concept demonstration of the workflow, not as a calibrated citywide deployment model.

Ground-truth speed values were generated as a function of roadway characteristics and scenario-specific behavioural modifiers. A stochastic noise term was included to represent natural speed variability. This synthetic dataset serves as the reference baseline against which the performance of the generative model was evaluated.

## 2.2 Sparse observation sampling and generative speed modelling

The proposed generative component is implemented as a parametric conditional probabilistic generator: an MLP predicts the conditional mean speed and uncertainty is represented via a residual-based dispersion term, from which synthetic speeds are sampled using a Gaussian distribution. This choice is deliberate. The objective is not high-fidelity content generation, but distributional sampling to compute exceedance probabilities ( $P(v > SL + 5)$ ) and to propagate uncertainty into risk and SSI estimates under sparse observations. We acknowledge that this differs from commonly referenced GenAI architectures (GAN/VAE/diffusion/autoregressive models); however, the required functionality here is conditional distribution learning and efficient scenario sampling for risk inference.

Rationale for using a generative model: A deterministic predictor would provide only a mean speed estimate, whereas the proposed workflow requires estimating the probability of exceeding a threshold ( $P(v > SL + 5)$ ). Therefore, we model speeds probabilistically by combining the predicted mean with a residual-based uncertainty estimate and generating synthetic samples. This approach enables scenario libraries and risk inference in data-poor settings without assuming full network-wide measurements.

Because real cities rarely possess comprehensive segment-level speed records, only a subset of road segments in the synthetic network was sampled ( $\approx 30\%$ ) to mimic real-world data scarcity. These observations were used to train a multilayer perceptron (MLP) model that learns the conditional distribution of speeds from roadway and environmental attributes.

Instead of producing a single deterministic estimate, the model outputs a predicted conditional mean speed,  $\hat{\mu}$ , and an uncertainty term expressed as the residual standard deviation,  $\hat{\sigma}$ . Synthetic speeds are then sampled from a Gaussian distribution parameterized by  $(\hat{\mu}, \hat{\sigma}^2)$ , as shown in Equation 1.

$$\tilde{v}_{e,s} \sim N(\hat{\mu}_{e,s}, \hat{\sigma}^2) \quad (1)$$

where  $\hat{\mu}_{e,s}$  is the predicted conditional mean speed and  $\hat{\sigma}$  is the residual standard deviation (thus  $\hat{\sigma}^2$  is the variance).

A total of 100–200 synthetic samples were generated per segment to estimate the distribution of speeding likelihood.

Let  $G = (V, E)$  denote the road network, where  $V$  is the node set and  $E$  is the edge (road segment) set. An edge is indexed by  $e \in E$ , and scenarios are indexed by  $s \in S$ . For segment  $e$  under scenario  $s$ ,  $X_{e,s}$  denotes the input feature vector constructed from segment attributes and scenario descriptors, and  $\hat{\mu}_{e,s}$  denotes the predicted conditional mean speed. The residuals are defined as  $r_i = v_i - \hat{\mu}_i$  over the training observations, and the uncertainty term is defined as the residual standard deviation  $\hat{\sigma} = \text{std}(r)$ ; hence  $\hat{\sigma}^2$  is the corresponding variance. Synthetic (generated) speeds are denoted by  $\hat{v}_{e,s}$ , sampled from *Normal* ( $\hat{\mu}_{e,s}, \hat{\sigma}^2$ ). The posted speed limit for segment is  $SL_e$ , and denotes the tolerance margin (set to 5 km/h in this study). The speeding probability under scenario is  $P_{e,s} = P(\hat{v}_{e,s} > SL_e + \delta)$ , and the aggregated segment risk is  $R_e = \sum_s w_s P_{e,s}$ , where  $w_s$  is the scenario weight. The Speed-Calming Suitability Index is denoted by  $SSI_e$ , computed by combining risk with the normalized planning criteria  $P_e$  (pedestrian exposure),  $T_e$  (traffic volume),  $E_e$  (equity factor), and  $We$  (winter maintenance difficulty), using weights  $\alpha, \beta, \gamma, \delta', \eta$ . Finally,  $N$  denotes the number of synthetic speed samples generated per  $(e, s)$ , and denotes the number of top-ranked candidate segments selected for deployment.

### 2.2.1 Baseline models for comparison

To assess whether the proposed conditional probabilistic generator provides measurable benefits beyond standard prediction approaches, we benchmarked it against widely used alternatives under the same sparse-observation setting. All baseline models use the same feature representation  $X_{e,s}$  (segment attributes concatenated with scenario descriptors), the same training–test split induced by the sampling ratio  $p$ , and the same target variable (scenario-specific speed observations). We include both point-prediction baselines and uncertainty-aware probabilistic baselines: (i) linear regression with L2 regularization (Ridge) as a transparent parametric benchmark; (ii) Random Forest regression as a nonlinear tree-ensemble benchmark; (iii) Gradient Boosting regression as a strong general-purpose predictor as a probabilistic baseline that directly provides predictive uncertainty; (iv) quantile regression to estimate conditional quantiles and approximate exceedance risk; and (v) an MC-dropout neural network as an approximate Bayesian deep-learning baseline for uncertainty estimation. For fair comparison, each model is evaluated using both accuracy metrics (mean-speed error) and decision-oriented metrics relevant to deployment (exceedance-risk quality and site-selection agreement), as detailed in Section 3.1.1.

## 2.3 Speed violation risk and Speed-Calming Suitability Index (SSI)

For each segment, the probability of exceeding the posted speed limit plus a tolerance margin ( $\delta$ : 5 km/h) was computed as shown in Equation 2.

$$R_e = \left(\frac{1}{N}\right) \sum (\hat{v}_e^n > SL_e + \delta) \quad (2)$$

Scenario-specific risks were aggregated through a weighted formulation representing real-world exposure patterns.

Scenario weights  $w_s$  reflect the relative exposure of the road network to each operating condition (peak-hour, night-time, weekend, winter). We enforce  $w_s \geq 0$  and  $\sum_s w_s = 1$ . In the proof-of-concept synthetic testbed, we adopt a transparent equal-weight baseline to avoid subjective tuning and to keep the demonstration policy-neutral.

In real-world applications,  $w_s$  can be calibrated from readily available aggregate evidence without requiring dense segment-level speed measurements. Specifically, we recommend setting weights proportional to the expected exposure share of each scenario, as shown in Equation 3.

$$w_s = \frac{\lambda_s}{\sum_{s' \in S} \lambda_{s'}} \quad (3)$$

where  $\lambda_s$  can be derived from (i) the fraction of time the scenario applies (hours per week/year), (ii) citywide or corridor-level traffic counts by time-of-day/day-type (peak vs. off-peak shares from loop detectors or short-term counts), and (iii) seasonal frequency for weather-driven scenarios (number of winter-storm days). This calibration allows the weighted aggregation to represent how often and how intensely drivers are exposed to each scenario in a given municipality. To test robustness, we perform a scenario-weight sensitivity analysis by perturbing  $w_s$  within plausible ranges while maintaining  $\sum_s w_s = 1$ , and we report changes in the top- $K$  ranked segments. This ensures that SSI-driven deployment recommendations are not an artefact of a single scenario-weight setting.

To provide actionable infrastructure planning guidance, a SSI was computed by integrating:

- predicted speed-violation risk,
- pedestrian exposure intensity,
- traffic volume category,
- equity indicators (income-based area weighting),
- winter maintenance operational difficulty.

The step-by-step workflow below summarizes the proposed conditional probabilistic generative decision-support module for speed management. Given a road network and segment attributes, the method learns a probabilistic speed model from sparse observations, generates synthetic speed samples, and propagates uncertainty to estimate speeding probability at each segment. These risk estimates are then combined with planning criteria to compute SSI, which is finally used to rank candidate segments for targeted deployment of adaptive speed-control interventions.

**Input:** Road network  $G$  (edges, nodes); roadway attributes  $A$  (road type, lanes, slope, winter severity, land use); traffic scenarios  $S$ ;

speed limits  $SL$ ; sampling ratio  $p$ ; synthetic sample count  $N$ ; scenario weights  $w_s$ ; planning criteria  $\{P_e, T_e, E_e, W_e\}$ .

**Step 1:** Generate the synthetic grid network and assign roadway attributes  $A_e$ . For each scenario  $s \in S$ , compute baseline mean speeds and generate observed speeds with Gaussian noise, as defined in Equations 4–6.

$$\mu_{e,s}^{true} = f(A_e, s) \quad (4)$$

$$v_{e,s}^{obs} = \mu_{e,s}^{true} + \varepsilon \quad (5)$$

$$\varepsilon \sim N(0, \sigma^2) \quad (6)$$

**Step 2:** Randomly sample  $p \cdot |E|$  edges to emulate sparse observations. Construct feature vectors and build training pairs, as defined in Equations 7–9.

$$p \cdot |E| \quad (7)$$

$$X_{e,s} = [A_e \parallel s] \quad (8)$$

$$(X_{e,s}, v_{e,s}^{obs}) \quad (9)$$

where  $X_{e,s}$  is the concatenation ( $\parallel$ ) of segment attributes and scenario descriptor.

**Step 3:** Train the neural-network regressor to estimate mean speed. Compute residuals and estimate uncertainty. Uncertainty is estimated as the standard deviation of residuals,  $\hat{\sigma} = std(r)$ . The generative model is then defined as  $\tilde{v}_{e,s} \sim N(\hat{\mu}_{e,s}, \hat{\sigma}^2)$ , shown in Equations 10–13. Here,  $N(\cdot)$  is parameterized by (mean, variance), i.e., the second parameter is  $\hat{\sigma}^2$ .

$$\hat{\mu}_{e,s} = NN(X_{e,s}) \quad (10)$$

$$r = v_{e,s}^{obs} - \hat{\mu}_{e,s} \quad (11)$$

$$\hat{\sigma} = std(r) \quad (12)$$

$$\tilde{v}_{e,s} \sim N(\hat{\mu}_{e,s}, \hat{\sigma}^2) \quad (13)$$

**Step 4:** For each  $(e, s)$  pair, generate  $N$  synthetic speeds by sampling from the generative model, as shown in Equation 14.

$$\tilde{v}_{e,s}^{(i)} \sim N(\hat{\mu}_{e,s}, \hat{\sigma}^2) \quad (14)$$

**Step 5:** Estimate speeding probability and aggregate across scenarios, as defined in Equations 15–16.

$$P_{e,s} = P(\tilde{v}_{e,s} > SL + \delta) \quad (15)$$

$$R_e = \sum_s w_s P_{e,s} \quad (16)$$

**Step 6:** Normalize criteria  $\{R_e, P_e, T_e, E_e, W_e\}$ . Compute the Speed-Calming Suitability Index, as shown in Equation 17.

$$SSI_e = \alpha R_e + \beta P_e + \gamma T_e + \theta E_e + \eta W_e \quad (17)$$

**Step 7:** Rank edges in descending order of  $SSI_e$ . Identify top- $K$  candidate segments and evaluate performance (correlation with ground truth and top- $K$  overlap).

**Output:** Risk map  $R_e$ ; Speed-Calming Suitability Index map  $SSI_e$ ; ranked list of priority segments for adaptive speed-control deployment.

Steps 1–3 establish the probabilistic speed-generation mechanism: Step 1 defines controlled baseline speeds and observed noisy samples; Step 2 emulates data scarcity through sparse sampling and feature construction; and Step 3 trains a predictive model and estimates uncertainty, enabling a generative

speed distribution. Step 4 creates a scenario library by sampling synthetic speeds, while Step 5 converts these samples into a segment-level speeding probability and risk score. Step 6 aggregates risk with additional criteria to compute SSI, and Step 7 operationalizes the output by ranking segments and producing a shortlist of priority locations for implementation.

### 2.3.1 SSI criteria construction and feature encoding

The SSI integrates speeding risk with planning-oriented criteria to support deployable, segment-level recommendations. Each segment  $e$  is represented by a feature vector that combines (i) roadway attributes (road type, lane count, slope class, winter severity, adjacent land-use) and (ii) a scenario descriptor  $s$ . Categorical attributes (road type, land-use class, scenario ID) are encoded using one-hot encoding; ordinal attributes (traffic volume class) are encoded as ordered integers; and continuous attributes (slope proxy, winter severity proxy) are kept as scalars.

The planning criteria are defined as follows:  $P_e$  (pedestrian exposure) is proxied from adjacent land-use intensity (residential/commercial segments assigned higher exposure than industrial);  $T_e$  (traffic volume) is encoded as an ordinal class derived from lane count and functional road type (local/collector/arterial);  $E_e$  (equity factor) is represented as an area-based vulnerability proxy (lower-income or higher-vulnerability zones assigned higher values to prioritize interventions); and  $W_e$  (winter maintenance difficulty) is proxied using winter severity and slope category, reflecting higher operational difficulty for snow/ice maintenance. These proxies are used in the synthetic grid experiment to ensure a transparent and reproducible SSI computation under data scarcity.

### 2.3.2 Normalization

Because the SSI aggregates heterogeneous criteria with different units and ranges, each criterion is normalized to [0,1] prior to aggregation. For a criterion value  $x_e$  on segment  $e$ , min-max normalization is applied over all segments in the study network, as shown in Equation 18.

$$x_e(\text{norm}) = (x_e - \min(x)) / (\max(x) - \min(x)) \quad (18)$$

All criteria are oriented such that larger normalized values indicate higher priority (higher speeding risk, higher pedestrian exposure, higher winter difficulty). The normalized terms are then used in the SSI computation.

### 2.3.3 Weight assignment

The SSI combines normalized criteria using weights ( $\alpha, \beta, \gamma, \theta, \eta$ ), constrained to sum to 1. In the absence of an agency-specific policy weight set, this study adopts an equal-weight baseline ( $\alpha = \beta = \gamma = \theta = \eta = 0.2$ ) to avoid subjective tuning and to provide a neutral proof-of-concept. The framework is intentionally designed so that agencies may replace these values with context-specific weights (higher  $\beta$  in pedestrian-priority zones, higher  $\eta$  in severe-winter municipalities).

### 2.3.4 Sensitivity analysis

To assess the robustness of deployment recommendations to weighting choices, we outline two complementary sensitivity analyses that can be applied to the SSI framework: one-at-a-time

(OAT) perturbations and Monte Carlo sampling of feasible weight vectors. In OAT, one weight is varied (by  $\pm 20\%$ ) while the remaining weights are re-normalized to maintain a unit sum, and the resulting changes in the top-K ranked segments can be tracked. In the Monte Carlo analysis, multiple feasible weight vectors are sampled on the simplex ( $\alpha, \beta, \gamma, \theta, \eta \geq 0$ ;  $\text{sum} = 1$ ), and the stability of the priority ranking is quantified by the frequency with which segments appear in the top-K. Together, these analyses provide a structured way to assess whether SSI-based site selection is dominated by a single criterion or remains stable across plausible policy preferences. The SSI criteria, their operational definitions, encoding choices, and normalization procedures are summarized in Table 1.

All criteria are oriented so that larger values indicate higher deployment priority. In this study,  $\delta$  is set to 5 km/h, and scenario weights  $w_s$  satisfy  $\sum_s w_s = 1$ .

$$SSI_e = \alpha R_e^{\text{norm}} + \beta D_e^{\text{norm}} + \gamma T_e^{\text{norm}} + \theta E_e^{\text{norm}} + \eta W_e^{\text{norm}}, \text{ with } \alpha + \beta + \gamma + \theta + \eta = 1$$

## 2.4 Micro-scale adaptive speed hump prototype

While the generative model supports macro-scale decision-making, a physical adaptive speed hump prototype was developed to demonstrate how risk-informed site selection can be operationalized at the segment level. The prototype functions as an electromechanical system that detects vehicle speed in real time and creates a physical barrier only for drivers exceeding the speed limit.

The system comprises three main components:

- i. Detection unit,
- ii. Decision algorithm,
- iii. Mechanical actuation module.

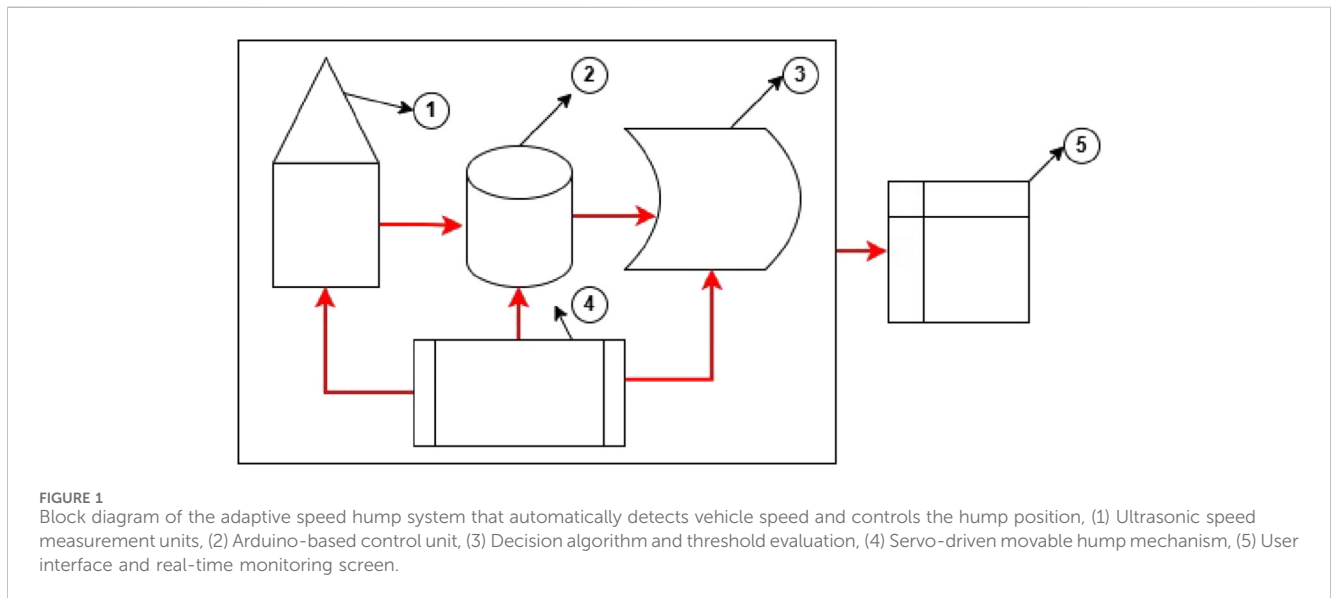
An Arduino Uno microcontroller serves as the control unit. Vehicle speed is detected using two HC-SR04 ultrasonic sensors mounted on the platform surface. The distance between sensors remains fixed, and vehicle speed is computed based on the measured passage time. When the calculated speed exceeds the user-defined threshold, a servo-motor-driven mechanism raises the hump module; otherwise, the platform remains flush with the surface. The hump module was fabricated using a 3D printer and scaled to 14 cm for the single-lane laboratory prototype.

All system logic was programmed in the Arduino IDE. The speed threshold was implemented as a user-adjustable parameter, enabling real-time calibration. A Processing-based graphical interface was developed to display the instantaneous speed, hump position, and threshold value. Users can increase or decrease the threshold and monitor system behaviour visually during operation.

To ensure autonomous operation, a solar-panel-assisted power unit was integrated into the prototype. The low energy consumption of the microcontroller-sensor-servo assembly significantly reduces operational cost and maintenance needs, offering a scalable solution for field deployment. Figure 1 illustrates the general workflow of the adaptive speed hump system. Ultrasonic sensors detect vehicle

TABLE 1 Definition and operationalization of SSI criteria.

Criterion	Definition (what it represents)	Operationalization/Proxy in this study	Encoding (raw)	Normalization (to [0,1])	Direction (higher means)
$R_e$ (speed violation risk)	Probability that typical speeds exceed the limit on segment (e), aggregated across scenarios	Scenario-weighted exceedance probability: $R_e = \sum_s w_s P_{e,s}$	Continuous	Min-max over all segments: $R_e^{norm} = \frac{R_e - \min(R)}{\max(R) - \min(R)}$	Higher speeding risk (higher priority)
$P_e$ (pedestrian exposure)	Potential pedestrian exposure along segment (e)	Land-use intensity proxy (residential/commercial assigned higher exposure than industrial)	Ordinal or continuous score	Min-max: $P_e^{norm} = \frac{P_e - \min(P)}{\max(P) - \min(P)}$	More pedestrian exposure (higher priority)
$T_e$ (traffic volume)	Relative motor-vehicle volume/flow level on segment (e)	Functional class and lane count proxy (arterial + more lanes → higher volume class)	Ordinal class (e.g., 1-5)	Min-max: $T_e^{norm} = \frac{T_e - \min(T)}{\max(T) - \min(T)}$	Higher traffic volume (higher priority)
$E_e$ (equity factor)	Priority for interventions in vulnerable/underserved areas	Area-based vulnerability proxy (higher score for higher vulnerability zones)	Ordinal or continuous score	Min-max: $E_e^{norm} = \frac{E_e - \min(E)}{\max(E) - \min(E)}$	Higher equity priority (higher priority)
$W_e$ (winter maintenance difficulty)	Relative operational difficulty for winter maintenance on segment (e)	Winter severity + slope proxy (higher severity and steeper slope → higher difficulty)	Ordinal or continuous score	Min-max: $W_e^{norm} = \frac{W_e - \min(W)}{\max(W) - \min(W)}$	Higher winter operational difficulty (higher priority)



presence and generate speed information. This data is relayed to the Arduino-based decision module, which evaluates speed relative to the threshold. If exceeded, the servo motor activates the movable hump. For compliant drivers, the hump remains level with the roadway.

To reduce false triggers and missed violations under noisy speed estimates, the actuation logic can be implemented with (i) a decision margin and (ii) hysteresis/persistence. Specifically, the hump is raised only if the estimated speed exceeds an upper threshold  $T_{high}$  and remains raised until the speed falls below a lower threshold  $T_{low}$ , where  $T_{high} = SL + \delta + m$  and  $T_{low} = SL + \delta - h$ . Here,  $m$  is a safety margin (proportional to an estimated speed standard deviation) and  $h$  is a hysteresis band. Alternatively, the system can require the violation condition to be met for two consecutive measurements before actuation.

To provide contextual evidence on the prevalence of speeding in a typical urban setting, a short field observation was conducted in Erzurum on a road segment where the posted speed limit is 50 km/h.

Vehicle speeds were estimated using a travel-time method over a measured 50 m distance: observers recorded the time at the entry and exit points, and speed was computed as distance divided by elapsed time. The objective of this observation was to quantify how frequently the posted urban limit is violated under natural flow conditions, rather than to evaluate the prototype.

Using this approach, the average speed under natural traffic flow was 46.08 km/h, while a substantial share of vehicles travelled near or above the 50 km/h limit. To check the plausibility of the field measurement procedure, repeated passes with a test vehicle at an intended 50 km/h were conducted, yielding a measured value of 54.08 km/h, which indicates a measurement deviation on the order of  $\pm 8$  km/h for individual runs. Figure 2 summarizes the instantaneous speed estimates obtained for 50 vehicles.

The modular laboratory-scale model includes an Arduino control unit, ultrasonic sensors, a servo motor assembly, and a solar-panel-assisted power unit. Key steps of the assembly process are shown in Figure 3.

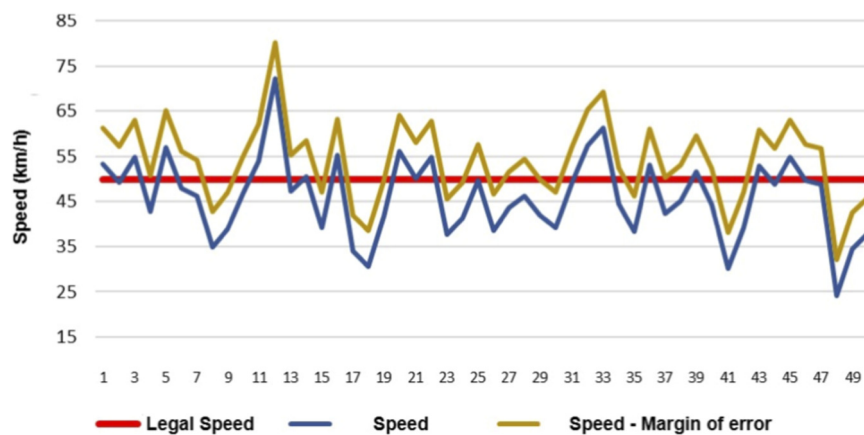


FIGURE 2

Field-based speed compliance observation on a 50 m urban road segment. Speeds were estimated from travel time between entry and exit points to quantify the prevalence of speeding under natural traffic flow (not prototype output).




FIGURE 3

Images from the assembly process of the scale model of the developed adaptive speed hump system. (a) Installation of sensors and electronic components and voltage measurement, (b) Integration of the solar panel and power unit, (c) Placement of the movable speed hump module inside the body.

### 3 Results and discussion

The proposed framework was evaluated at two complementary levels:

- i. A generative, network-wide risk mapping and site-selection experiment on a synthetic urban grid, and
- ii. A micro-scale adaptive speed hump prototype demonstrating selective speed control in practice.

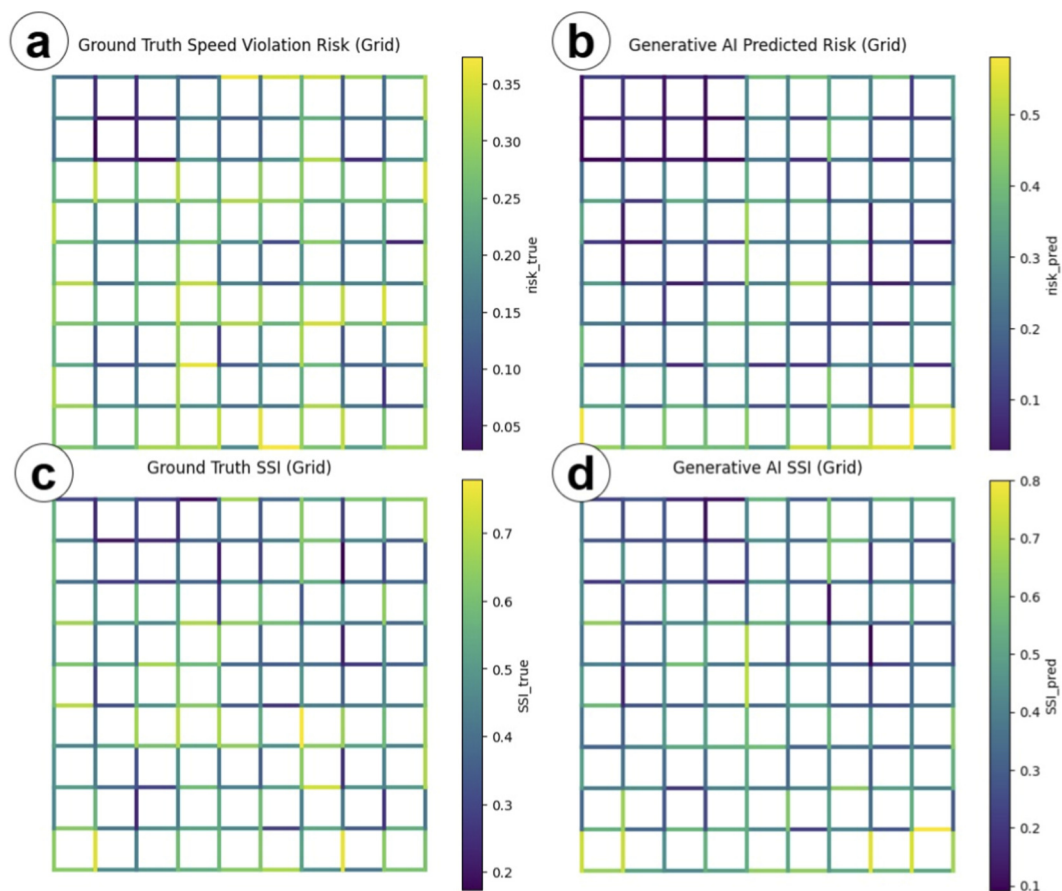
#### 3.1 Performance of the generative risk mapping framework

Figure 4 illustrates the spatial distribution of ground-truth and generative-model-predicted speed violation risk on a synthetic grid testbed designed for controlled experimentation (not a specific city representation). The ground-truth map (Figure 4a) reflects the

behavioural rules embedded in the synthetic data generation process, with elevated risk concentrated along specific arterial segments and intersections. The probabilistic generative prediction (Figure 4b) reproduces these patterns with reasonable fidelity while also smoothing some local fluctuations. Visually, high-risk corridors in the lower and central parts of the network are captured well, whereas a few peripheral segments are slightly overestimated or underestimated.

This qualitative agreement is confirmed quantitatively. The Pearson correlation coefficient between the segment-level ground-truth risk values and the probabilistic generative predictions is 0,659,971, indicating a moderate positive association between the two risk surfaces. Given that the model is trained on only a subset of segments (approximately 30% of the network), this level of correspondence suggests that the generative regression model successfully learns meaningful relationships between roadway attributes, scenario conditions, and speeding behaviour.

Figures 4c,d present the corresponding SSI maps based on ground-truth and predicted risk, respectively. Because SSI combines risk with additional planning dimensions, pedestrian



**FIGURE 4** Proof-of-concept results on a synthetic grid testbed (not a specific city): ground-truth vs. Probabilistic generative-predicted speeding risk and SSI. **(a)** Ground-truth risk; **(b)** AI-predicted risk; **(c)** ground-truth SSI; **(d)** AI-predicted SSI. The testbed provides controlled ground truth to evaluate whether the generative workflow can recover risk patterns and support site prioritization under sparse observations.

exposure, traffic volume, equity indicators, and winter maintenance considerations, the patterns are more spatially diffuse than in the risk maps. Nonetheless, the SSI surfaces derived from ground-truth and probabilistic generative risk show comparable high-priority clusters, especially around central grid corridors and selected peripheral links. This demonstrates that the generative framework is not only able to approximate raw speeding risk, but also to propagate that information into a multi-criteria prioritization index that is relevant for infrastructure planning.

To assess the implications for site selection, segments were ranked according to SSI and the top 20 candidates were compared. The overlap between the top-20 sites based on ground-truth SSI and GenAI-derived SSI is 0.50, meaning that 40% of the highest-priority locations would be identically selected under both approaches. Considering the data scarcity scenario and the stochastic nature of the synthetic sampling, this level of agreement indicates that the proposed generative workflow can provide practically useful guidance for locating adaptive speed-control devices, even when full network-wide speed data are not available. In a real-world context, such a tool would allow agencies to screen large urban areas, identify a manageable shortlist of candidate segments, and then focus field measurements and investment decisions on those locations.

Overall, these findings suggest that the probabilistic generative approach is capable of.

- i. Reconstructing spatial patterns of speeding risk with moderate accuracy under sparse observations, and
- ii. Supporting robust, multi-criteria site-selection decisions for traffic-calming infrastructure. The grid experiment thus serves as a controlled proof-of-concept for how generative models can augment traditional traffic safety analysis in data-constrained urban environments.

As a methodological benchmark, the synthetic testbed allows systematic comparison against known ground truth; real-world deployment would require calibration and validation using measured speeds, which is explicitly acknowledged in the limitations section.

### 3.1.1 Comparison against standard baselines

Table 2 shows that, on this synthetic testbed under sparse observations ( $p \approx 0.30$ ), Ridge regression combined with a residual-based uncertainty mapping achieves the strongest overall performance across exceedance Brier score and decision-oriented metrics (risk-map correlation and top-20 overlap). The proposed

TABLE 2 Baseline comparison on the synthetic grid testbed under sparse observations ( $\rho = 0.30$ ).

Model	Mean speed MAE	Mean speed RMSE	Exceedance brier	Risk-map corr. (r)	Top-20 overlap
Linear regression (ridge)	1.91341	2.358786	0.009676	0.820621	0.65
Random forest	2.67131	3.286092	0.026895	0.574928	0.5
Gradient boosting	2.63415	3.169802	0.019097	0.599695	0.5
Quantile regression (GBR)	3.136839	3.883762	0.026825	0.625241	0.55
MC-dropout MLP	3.491433	4.239826	0.021894	0.628207	0.55
Proposed probabilistic generator (MLP + $\sigma$ )	3.358419	4.070785	0.02548	0.659971	0.5

Mean-speed metrics are computed over unobserved segments; exceedance metrics are computed for the event  $v > S L + \delta$ .

probabilistic generator performs competitively but does not outperform the best baseline in this setting. We therefore interpret the value of the proposed approach primarily as a transparent distributional sampling framework that enables consistent computation of exceedance probabilities and downstream SSI-based prioritization under data scarcity, while acknowledging that simpler baselines may perform better depending on the data-generating process and feature set.

## 3.2 Implications for speed hump site selection

From a planning perspective, the SSI maps demonstrate how generative risk estimation can be translated into actionable decisions. Rather than installing speed humps opportunistically or based solely on local complaints, municipalities could use the proposed framework to.

- generate synthetic speed distributions for different time-of-day and weather scenarios,
- compute segment-level speeding risk under each scenario,
- integrate exposure, equity, and maintenance criteria, and
- rank segments according to their overall suitability for adaptive speed control.

Such a process would enable more transparent, evidence-based justification for the placement of speed-calming devices. In addition, by repeating the analysis under alternative weighting schemes, decision-makers could explore trade-offs, for example, prioritizing equity and pedestrian safety versus minimizing winter maintenance costs, within a consistent analytical framework. In practice, the GenAI module is used to screen the network and shortlist candidate segments for intervention (where-to-deploy). The adaptive speed hump prototype then represents the unit intervention that would be installed at one of these prioritized segments to operationalize selective speed control (how-to-operationalize). Accordingly, the grid experiment does not extrapolate prototype mechanics to a whole city; it demonstrates the decision-support layer that informs deployment targeting. The risk/SSI component is intended as a screening tool to shortlist candidate locations (where-to-deploy). The prototype represents the unit intervention that can be installed at one shortlisted

segment to operationalize selective speed control (how-to-operationalize).

## 3.3 Evaluation of the adaptive speed hump prototype

### 3.3.1 Actually implemented: laboratory-scale prototype and demonstration tests

While the probabilistic generative modelling framework operates at the network scale, the laboratory-scale adaptive speed hump prototype implemented in this study targets the operational feasibility of selective actuation at a single installation point. The prototype is designed to mitigate limitations of fixed vertical traffic-calming devices, which can reduce comfort for compliant drivers and create maintenance burdens in cold-climate cities (snowplough impacts leading to repeated reconstruction and pavement degradation). In this context, the implemented prototype actuates only when a speed violation is detected, allowing compliant vehicles to pass over a flush surface.

During demonstration tests, the implemented sensing and actuation chain operated consistently across different threshold settings: the hump remained level for below-threshold crossings and raised for detected violations. The final laboratory prototype and its components are shown in Figure 5. Because the prototype was not instrumented for systematic logging during the demonstration phase, quantitative performance is reported as an analytical performance envelope derived from the implemented scale geometry (sensor spacing 0.17 m; offset from Sensor-2 to the hump 0.125 m) and nominal controller/actuator timing. The key feasibility condition is that the hump reaches the raised position before a vehicle travels from Sensor-2 to the hump. Table 3 summarizes the resulting timing constraints, expected speed uncertainty due to event-timing jitter, and the implied maximum effective speed range for full deployment in the laboratory-scale setup.

To address threshold-actuation reliability under speed-estimate fluctuations, we adopt an uncertainty-aware triggering interpretation: the analytical timing envelope motivates conservative actuation rules that reduce false triggers and missed violations. These measures are particularly relevant for low-cost sensors and small-scale prototypes where measurement noise can be non-negligible.



FIGURE 5 General view and system components of the final version of the adaptive speed hump mechanism. Figure contains images of the author(s) only.

TABLE 3 Analytical performance envelope of the scale adaptive speed hump prototype.

Metric	Definition/basis	Value (this study)
Sensor spacing, (d)	Distance between ultrasonic sensors	0.17 m
Actuator offset, (L)	Distance from Sensor-2 (exit) to hump	0.125 m
Available time window	$t_{avail} = L/v$	Velocity-dependent
Representative speed levels (toy-scale)	Used for envelope illustration	0.4, 0.6, 0.8 m/s
Time between sensors	$\Delta t = d/v$	0.425 s (0.4 m/s); 0.283 s (0.6 m/s); 0.213 s (0.8 m/s)
Expected speed uncertainty (timing jitter-based)	Assuming effective event timing uncertainty $\Delta t_{err} \approx 7$ ms (ultrasonic triggering + detection jitter); $\Delta v \approx \frac{d}{\Delta t} \Delta t_{err}$	$\pm 0.0066$ m/s (0.4); $\pm 0.0148$ m/s (0.6); $\pm 0.0264$ m/s (0.8)
Control logic latency	Arduino decision and command issuance (order-of-magnitude)	$\approx 1-5$ ms
Full actuation time (assumption)	Typical micro-servo deployment time used for analytical envelope	0.15 s (conservative)
Detection-to-full-deployment latency (analytical)	$t_{total} \approx t_{logic} + t_{servo}$	$\approx 0.155$ s
Maximum effective speed for full deployment	$v_{max} = L/t_{total}$	$\approx 0.125/0.155 = 0.81$ m/s
Example violation-classification performance (simulation-based)	Assuming Gaussian speed noise $\sigma = 0.02$ m/s around a threshold $v_{thr}$ , 1000 simulated trials with overlapping compliant/violation speeds	Accuracy $\approx 93.2\%$ , precision $\approx 92.3\%$ , recall $\approx 93.1\%$

### 3.3.2 Proposed: full-scale deployment considerations

The following considerations relate to full-scale field deployment and are proposed design requirements rather than implemented features of the laboratory prototype. First, selective actuation in real traffic must provide sufficient time and distance for a driver to perceive the cue and decelerate safely. This motivates one or more of the following deployment strategies: (i) upstream detection with an advance warning distance, (ii) staged actuation

(visual/auditory warning followed by physical deployment), or (iii) conservative triggering policies designed to avoid abrupt or late activation. Full-scale deployment therefore requires calibration of warning distance based on approach speed distributions, driver reaction time, braking capability, and roadway operating conditions.

Second, from a feasibility perspective, a full-scale implementation would require engineering of the mechanical module, enclosure/weatherproofing, and foundation integration for repeated heavy-vehicle loading and long-term exposure to

rain, snow, icing, and salt. Under such conditions, the cost structure is expected to be dominated by the mechanical module, actuator assembly, and environmental sealing, with sensing and control electronics representing a smaller fraction. Based on a component-level estimate, a full-scale smart speed hump could be produced at approximately 3,000 USD per lane, including sensing and electronics, mechanical module, energy system, control interface, and installation. Maintenance requirements would likely be driven primarily by mechanical wear, sealing integrity, and winter operability. When contrasted with recurring reconstruction expenses associated with conventional fixed humps in harsh winter environments, this investment can be competitive and potentially cost-saving over the infrastructure lifecycle.

In summary, the implemented laboratory prototype demonstrates the feasibility of selective actuation at a single point, while the proposed full-scale considerations clarify the practical requirements for real-world deployment. The smart adaptive speed hump concept is attractive because it can.

- avoid unnecessary shocks for drivers complying with the speed limit,
- provide a physical deterrent for violations through selective actuation,
- reduce recurrent reconstruction by remaining flush during winter maintenance operations, and
- support energy-independent operation via a solar-powered configuration.

Coupled with the probabilistic generative risk mapping and site-selection framework, the prototype illustrates how intelligent roadside infrastructure could be deployed strategically at prioritized locations to support safer and more sustainable urban mobility.

## 4 Limits of the study

This study integrates a Generative AI-based risk modelling framework with a prototype-level adaptive speed hump system; however, several limitations should be acknowledged when interpreting the findings. First, the generative modelling component was validated using a synthetic grid network rather than an empirically observed urban road network. Although this approach enables controlled experimentation and isolates methodological behaviour, it does not fully capture the complexity of real-world traffic dynamics, driver heterogeneity, or context-specific behavioural patterns. The model was trained under sparse observations by design, and its accuracy may vary under different sampling ratios, feature uncertainties, or alternative network structures. Furthermore, scenario definitions (peak hour, night, weekend, winter) were simplified representations of typical operating conditions, and additional temporal, weather, or land-use factors could influence real-world speeding behaviour. In addition, the generative module is implemented as a parametric conditional probabilistic model (mean + dispersion) rather than a high-capacity deep generative architecture (VAE/GAN/diffusion/autoregressive models). Future work will benchmark these alternative generative model classes once real-world calibration datasets are available, to

assess whether they yield measurable gains in risk-surface fidelity and site-selection stability.

With respect to the prototype system, the smart speed hump was evaluated at a scaled laboratory level and was not subjected to full-scale mechanical load testing. As such, the structural resilience of the system under the repeated axle loads of passenger cars, heavy-duty vehicles, and emergency vehicles remains unknown. The experimental validation was conducted using a single type of test vehicle, which limits the generalizability of the results. Vehicle-dependent factors such as wheelbase length, suspension stiffness, and tire characteristics may influence speed detection accuracy and hump-vehicle interaction forces.

The prototype was also assessed in a controlled indoor environment. Consequently, the long-term impacts of environmental stressors, such as rain, snow, ice formation, high temperature fluctuations, dust, or salt accumulation, were not investigated. It also remains unclear how the system would perform under maintenance operations, especially in cold-climate regions where snowplough contact poses a recurrent risk to roadside infrastructure. Additionally, this study evaluated only a single-lane configuration; multi-lane or modular installations, which may require synchronized actuation or distributed sensing, were not examined.

These limitations highlight the need for future work involving (i) deployment on real-world road segments with empirically measured speeds, (ii) more extensive mechanical and environmental durability testing, and (iii) evaluation of the generative model on larger, heterogeneous networks under broader operational scenarios. Such advancements will be essential before transitioning the system from a proof-of-concept to a field-ready smart mobility solution. In full-scale deployment, threshold-based actuation should be validated with logged ground-truth speeds and implemented with hysteresis and safety margins to ensure robust triggering under sensor noise and environmental variability.

## 5 Conclusion

This study presents an integrated framework that combines a Generative AI-based risk modelling approach with a prototype-level adaptive speed hump system designed to improve urban traffic safety while minimizing unnecessary driver discomfort. On the computational side, a synthetic grid network was used to evaluate how generative models can approximate speeding behaviour under sparse observations and produce spatially refined speeding risk maps. The framework further introduced a SSI, which aggregates speeding risk with multi-criteria planning factors to support equitable and context-aware site selection for traffic calming infrastructure. The results demonstrate that the generative model captures major spatial trends in both speeding risk and SSI, offering a promising pathway for data-scarce municipalities to prioritize investment locations.

On the hardware side, a dynamic speed hump mechanism was designed and implemented using ultrasonic sensors, a microcontroller-based control unit, and a servo-driven actuation system. Experimental validation confirmed that the prototype can detect vehicle speed in real time and raise or lower the hump surface

according to a predefined speed threshold. In doing so, the system preserves comfort for compliant drivers while providing a physical deterrent for violators. Cost analysis suggests that the proposed design can be produced using low-cost components, and the solar-powered configuration supports energy independence and operational sustainability, particularly beneficial in regions with harsh winter climates where traditional speed humps are frequently damaged.

The combined results highlight that both components, the GenAI-enabled risk prioritization tool and the adaptive mechanical speed control system, offer complementary contributions to next-generation smart mobility applications. The modelling framework helps identify where speed-control interventions are most needed, while the prototype system provides a selective, lower-impact alternative to traditional fixed speed humps.

Future work will focus on several enhancements. From a modelling perspective, applying the generative framework to real-world speed datasets, expanding scenario definitions, and integrating richer environmental and behavioural features will further improve predictive realism. From an engineering perspective, future prototypes will incorporate embedded roadway sensors or radar systems to improve speed detection accuracy, as well as modular, lane-specific actuation mechanisms suitable for multi-lane corridors. Field-scale deployment and long-term durability testing under diverse climatic conditions will be essential steps toward integrating the system into broader smart city infrastructures. Together, these directions will help advance both the methodological and operational aspects of intelligent, data-driven traffic calming solutions.

## Data availability statement

The raw data supporting the conclusions of this article will be made available by the authors, without undue reservation.

## Author contributions

ÖK: Methodology, Supervision, Writing – review and editing, Investigation, Conceptualization, Software, Data curation, Validation, Writing – original draft, Resources, Formal Analysis, Visualization, Project administration, Funding acquisition.

## References

- Aceto, G., Giampaolo, F., Guida, C., Izzo, S., Pescapè, A., Piccialli, F., et al. (2024). Synthetic and privacy-preserving traffic trace generation using generative AI models for training network intrusion Detection Systems. *J. Netw. Comput. Appl.* 229, 103926. doi:10.1016/j.jnca.2024.103926
- Afukaar, F. K. (2003). Speed control in developing countries: issues, challenges and opportunities in reducing road traffic injuries. *Inj. Control Saf. Promot.* 10, 77–81. doi:10.1076/icsp.10.1.77.14113
- Alshabibi, N. M. (2025). An impact assessment of speed humps' geometric characteristics and spacing on vehicle speed: an overview. *Infrastructures* 10, 190. doi:10.3390/infrastructures10070190
- Basaran, O. T., and Dressler, F. (2025). XAIomaly: explainable and interpretable Deep Contractive Autoencoder for O-RAN traffic anomaly detection. *Comput. Netw.* 261, 111145. doi:10.1016/j.comnet.2025.111145
- Bibri, S. E., and Huang, J. (2025). Generative AI of things for sustainable smart cities: synergizing cognitive augmentation, resource efficiency, network traffic, cybersecurity, and anomaly detection for environmental performance. *Sustain. Cities Soc.* 133, 106826. doi:10.1016/j.scs.2025.106826
- Chen, Q., Ding, D., Wang, X., Liu, A. X., and Zhao, J. (2019). An efficient urban localization method based on speed humps. *Sustain. Comput. Inf. Syst.* 24, 100341. doi:10.1016/j.suscom.2019.07.004
- Džambas, T., Ivančev, A. Č., Dragčević, V., and Vujević, I. (2023). Safety and environmental benefits of intelligent speed bumps. *Transp. Res. Procedia* 73, 159–166. doi:10.1016/j.trpro.2023.11.904
- Golbabaee, F., Mirhashemi, A., Rabbani, W., Ahern, Z., Denman, S., and Paz, A. (2025). Effectiveness of raised safety platforms: a systematic review of literature. *Transp. Res. Interdiscip. Perspect.* 34, 101671. doi:10.1016/j.trip.2025.101671

## Funding

The author(s) declared that financial support was not received for this work and/or its publication.

## Acknowledgements

The author thanks Erzurum Technical University for providing laboratory facilities and acknowledges the helpful discussions that supported the development of the study.

## Conflict of interest

The author(s) declared that this work was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

## Generative AI statement

The author(s) declared that generative AI was used in the creation of this manuscript. The author confirms and takes full responsibility for the use of generative AI tools in the preparation of this manuscript. Generative AI was used only to assist with language refinement and structural organization of certain sections. All scientific content, methodological design, data generation, analysis procedures, and conclusions were conceived, developed, and verified entirely by the author.

Any alternative text (alt text) provided alongside figures in this article has been generated by Frontiers with the support of artificial intelligence and reasonable efforts have been made to ensure accuracy, including review by the authors wherever possible. If you identify any issues, please contact us.

## Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

- Kim, E. J., Kim, D.-K., Lee, H.-S., and Joo, Y. J. (2025). Generative artificial intelligence for class imbalance in crash severity estimation with mixed data types. *Transp. Res. Rec. J. Transp. Res. Board*, 03611981251387114. doi:10.1177/03611981251387114
- Kosakowska, K. (2022). Evaluation of the impact of speed bumps on the safety of residents - selected aspects. *Transp. Res. Procedia* 60, 418–423. doi:10.1016/j.trpro.2021.12.054
- Li, J., Guo, W., Wang, L., and Chen, S. (2019). Multi-objective optimization of ambulance ride comfort under speed bump. *IEEJ Trans. Electr. Electron. Eng.* 14, 1372–1380. doi:10.1002/tee.22939
- Lin, H.-Y., and Ho, C.-Y. (2022). Adaptive speed bump with vehicle identification for intelligent traffic flow control. *IEEE Access* 10, 68009–68016. doi:10.1109/ACCESS.2022.3186010
- Mei, J., and Wang, J. (2021). “An intelligent dynamic speed hump,” in *2021 IEEE 4th international conference on automation, electronics and electrical engineering (AUTEEE)* (IEEE), 274–277. doi:10.1109/AUTEEE52864.2021.9668805
- Miracle, P. U., Bala, J. A., Aibinu, A. M., and Folorunso, T. A. (2021). “Intelligent speed bump identification and speed control system for autonomous vehicles: a conceptual design,” in *2021 1st international conference on multidisciplinary engineering and applied science (ICMEAS)* (IEEE), 1–6. doi:10.1109/ICMEAS52683.2021.9692403
- Ray, U., Arteaga, C., Ahn, Y., and Park, J. (2024). Enhanced identification of equipment failures from descriptive accident reports using language generative model. *Eng. Constr. Archit. Manag.* doi:10.1108/ECAM-09-2024-1259
- Šarić, Ž., Kučinić, T., Kunštek, A., and Ondruš, J. (2025). The impact of speed bumps on traffic flow speed in urban road networks. 15, 1–18. doi:10.3390/app152212221
- Sirbu, C. L., Tomoiu, C., Fancsali-Boldizsar, S., and Orhei, C. (2021). “Real-time line matching based speed bump detection algorithm,” in *2021 IEEE 27th international symposium for design and technology in electronic packaging (SIITME)* (IEEE), 246–249. doi:10.1109/SIITME53254.2021.9663602
- Wilder, D., Magnusson, M. L., Fenwick, J., and Pope, M. (1994). The effect of posture and seat suspension design on discomfort and back muscle fatigue during simulated truck driving. *Appl. Ergon.* 25, 66–76. doi:10.1016/0003-6870(94)90067-1
- Yurtbaş, E., Kuşkan, E., and Çodur, M. Y. (2022). Analysis of vibration exposure of vehicles during passing through speed bumps. *Pamukkale Univ. J. Eng. Sci.* 28, 777–785. doi:10.5505/pajes.2021.82504