Intelligent Pedestrian Model as a Risk-Based Framework for Pedestrian Prioritization

Zoltán Rózsás * and István Lakatos

Posted Date: 24 April 2026

doi: 10.20944/preprints202604.1699.v1

Keywords: pedestrian safety; risk prioritization; trajectory prediction; ADAS; vulnerable road users; cost-aware learning; exposure modeling; urban intersection

 doi:10.20944/preprints202604.1699.v1

Article

Intelligent Pedestrian Model as a Risk-Based Framework for Pedestrian Prioritization

Zoltán Rózsás [1,*] and István Lakatos [2]

[1] Department of Vehicle Maintenance and Diagnostics, Audi Hungaria Faculty of Vehicle Engineering, Széchenyi István University, Multidisciplinary Doctoral School of Engineering

[2] Department of Vehicle Maintenance and Diagnostics, Audi Hungaria Faculty of Vehicle Engineering, Széchenyi István University

* Correspondence: zoltan.rozsas@outlook.com

Abstract

Pedestrian safety at urban intersections requires risk-aware mechanisms that extend beyond binary collision detection toward comparative prioritization among multiple agents. This study introduces the Intelligent Pedestrian Model (IPM), a reference-normalized scalar framework that represents pedestrian risk as a function of trajectory, contextual, infrastructural, and behavioral factors, decomposed into Exposure and Severity components. Building on IPM, the Safety-Prioritized Trajectory Model (SPTM) operationalizes the Exposure component using an observation-only, leakage-free kinematic proxy embedded into a cost-aware negative log-likelihood objective. Evaluation on the ETH/UCY benchmark under a strictly inductive protocol shows that moderate prioritization ($\beta \approx 1.0$) improves best-of-K multimodal performance (ALL FDE@K: 0.979 → 0.970 m) while maintaining mean displacement accuracy within seed-level variability. The results indicate that Exposure-based weighting does not act as a global accuracy enhancer but redistributes predictive capacity toward safety-relevant motion regimes. Validation is limited to a single benchmark fold; cross-fold generalization and full IPM instantiation remain future work.

Keywords: pedestrian safety; risk prioritization; trajectory prediction; ADAS; vulnerable road users; cost-aware learning; exposure modeling; urban intersections

1. Introduction

Pedestrian fatalities remain a persistent challenge in road safety policy. According to the World Health Organization's Global Status Report on Road Safety 2023 [1], pedestrians account for approximately 23% of all road traffic deaths globally. In the European Union, Eurostat data [2] indicate that vulnerable road users represent a disproportionately large share of fatalities relative to their modal share. The EU Road Safety Policy Framework 2021–2030 identifies the elimination of pedestrian fatalities as a central pillar of its Vision Zero strategy [3].

Automatic Emergency Braking (AEB) systems have demonstrated measurable effectiveness in reducing rear-end collisions and improving road safety outcomes [4,5]. As vehicle automation levels increase toward SAE Level 3 and beyond [6], systems must manage increasingly complex, multi-agent urban scenarios. At urban intersections, multiple pedestrians simultaneously occupy the operational domain, each exhibiting heterogeneous behavioral patterns, motion dynamics, and vulnerability characteristics.

Current trajectory prediction architectures including the Social Force Model [7], Social-LSTM [8], multimodal generative approaches [9], and transformer-based predictors [10]—primarily optimize displacement accuracy. They estimate where a pedestrian will move but do not formalize how multiple pedestrians should be comparatively prioritized in safety-critical decision processes. In environments where ADAS systems must allocate limited attention, warning capacity, or intervention priority across multiple actors, a displacement-optimizing framework is insufficient.

 doi:10.20944/preprints202604.1699.v1

This paper addresses the structural gap between displacement-oriented trajectory prediction and comparative safety prioritization in multi-agent environments through two contributions. First, the Intelligent Pedestrian Model (IPM) is introduced as a formal framework in which pedestrian risk is represented as a reference-normalized scalar functional over heterogeneous domains, decomposed into Exposure and Severity components, and interpreted as a population-relative deviation. Second, the Safety-Prioritized Trajectory Model (SPTM) provides a constrained empirical instantiation of the Exposure component under a strictly inductive evaluation protocol. The scope is explicitly limited. The objective is not to achieve state-of-the-art displacement accuracy, but to demonstrate that Exposure-based cost-aware learning is measurable, stable, and compatible with deployment-relevant constraints.

2. Related Work

2.1. Trajectory Prediction Paradigms

Pedestrian motion modeling has been dominated by trajectory prediction since the Social Force Model [7]. Social-LSTM [8] introduced recurrent architectures with social pooling; Trajectron++ [11] extended this with dynamically feasible forecasting; transformer-based models [10] apply attention for long-range dependency modeling. Multimodal generative approaches—including conditional variational autoencoders [9] and generative adversarial networks [12]—address the inherent multi-modality of pedestrian behavior. Despite these advances, all such models share a displacement-optimizing objective that does not address the prioritization problem inherent in multi-agent ADAS scenarios. Comprehensive surveys of trajectory prediction methods further confirm that the dominant focus across architectures remains displacement accuracy rather than decision-relevant prioritization [13,14].

2.2. Safety-Oriented Pedestrian Modeling

Surrogate safety measures such as Time-to-Collision (TTC) and Post-Encroachment Time (PET) quantify conflict severity [15]. Responsibility-Sensitive Safety (RSS) frameworks [16] formalize safety envelopes but focus on vehicle-to-vehicle interactions. Pedestrian intention estimation models [17–19] address crossing intention but do not integrate multi-domain risk factors into a unified scalar representation. Map-aware prediction approaches [20] incorporate infrastructural context; context-based models [21] demonstrate environment-dependent behavioral transitions. To the authors' knowledge, no existing framework formalizes risk as a reference-normalized, population-level deviation enabling ordered prioritization across simultaneously present pedestrians. The present paper addresses this gap through a dissertation-derived contribution: the IPM provides the normative prioritization framework, while the SPTM provides a constrained first empirical instantiation of its Exposure component — together spanning the formal framework and the scoped predictive core of the broader research program. To the authors' knowledge, no prior work provides a reference-normalized, population-relative risk ranking formulation that explicitly separates Exposure and Severity components, as introduced by the IPM.

3. The Intelligent Pedestrian Model (IPM)

3.1. Validation Architecture and Scope of This Article

The IPM is not a monolithic model but a layered research program comprising formal foundations and four empirically progressive validation levels, illustrated in **Figure 1**. Each level is characterized by increasing empirical complexity and ecological fidelity.

 doi:10.20944/preprints202604.1699.v1

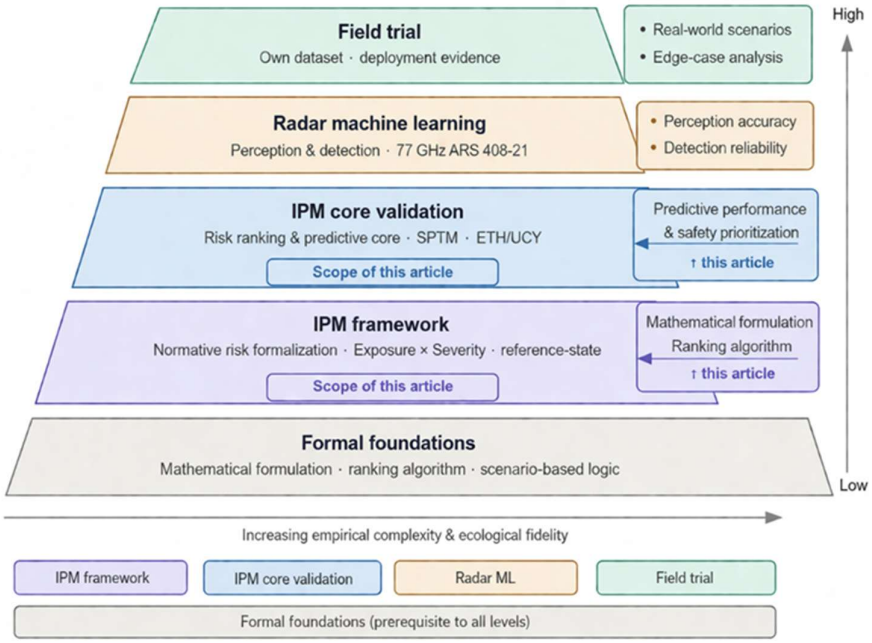

Figure 1. Validation architecture of the Intelligent Pedestrian Model. The pyramid illustrates four empirical levels built on formal mathematical foundations, with increasing validation depth toward field deployment. Source: Author's own conceptual illustration.

The formal foundations layer establishes the mathematical formulation, ranking algorithm, and scenario-based logic on which all higher levels rest. The first validation level operationalizes these foundations as the normative IPM framework: Exposure × Severity decomposition, reference-state risk scoring, and the population-level deviation ΔR. The second validation level tests the IPM predictive core through risk ranking and trajectory prediction, providing scenario-based evaluation and baseline comparison. The third level extends to radar-based machine learning for perception and detection using 77 GHz automotive radar. The fourth and highest level applies the complete system to field trial recordings under real-world deployment conditions.

This article covers the formal foundations layer and the first two validation levels. Specifically, Section 3 presents the normative IPM framework (Level 1), and Section 4 presents the Safety-Prioritized Trajectory Model as the core empirical validation (Level 2). The radar machine learning layer (Level 3) and the field trial layer (Level 4) are subjects of ongoing and future research; they are described here solely to contextualize the IPM's intended deployment architecture but are not empirically investigated in this paper.

3.2. From Trajectory Prediction to Risk Prioritization

In real-world urban intersections, multiple pedestrians coexist under heterogeneous behavioral and environmental conditions. Safety systems must allocate limited attention or intervention capacity across several actors simultaneously, transforming the problem from binary collision detection into a comparative prioritization task. Trajectory prediction models answer the question "Where will they go?" and optimize displacement error or likelihood-based objectives. Safety-critical ADAS systems must instead answer "Who requires attention first?" -an objective requiring a structured mechanism that converts heterogeneous behavioral and contextual indicators into an ordered risk hierarchy. The IPM formalizes this transition.

3.3. Formal Risk Representation

 doi:10.20944/preprints202604.1699.v1

The literature consistently demonstrates that pedestrian safety relevance is multi-factorial [7,8,18–21]. Motion dynamics, social interactions, environmental context, infrastructure constraints, and behavioral intention cues all influence movement and safety relevance. Accordingly, the IPM defines pedestrian risk as the functional in Equation (1), integrating trajectory, contextual, infrastructural, and behavioral factors into a unified representation.

$$R_i = f(T_i, C_i, I_i, B_i) \quad (1)$$

where:

R_i denotes the Pedestrian Risk Score of pedestrian i;

T_i represents trajectory-related information (position, velocity, acceleration);

C_i denotes traffic conditions (vehicle speeds, density, visibility, weather);

I_i represents intersection characteristics (geometry, crosswalk presence, signalization);

B_i captures pedestrian-specific behavioral attributes (age, movement patterns, crossing intention).

Within the IPM, the scalar risk score is decomposed into an Exposure component and a Severity scaling factor, as defined in Equation (2):

$$R_i = Exposure_i \times Severity_i \quad (2)$$

Exposure integrates dynamic, contextual, infrastructural, and behavioral indicators into a unified representation. Severity reflects vulnerability-dependent consequence scaling. This decomposition enforces semantic separation between domain types, preventing implicit conflation of heterogeneous variables. The IPM does not directly estimate collision probability or injury severity; instead, it introduces a scalar Pedestrian Risk Score functioning as a comparative safety relevance indicator. The exact functional form of f(·) is intentionally left unspecified at the framework level, preserving extensibility across sensing modalities, data availability, and deployment constraints.

3.4. Reference-Based Risk Interpretation

A fundamental conceptual question concerns whether risk should be interpreted as an absolute magnitude or as a context-dependent deviation within a multi-agent environment. Traditional metrics such as TTC quantify conflict imminence independently of the surrounding behavioral distribution. However, identical kinematic magnitudes may have different safety implications depending on the collective behavioral state of a scene: a walking speed that is unremarkable in a dense crosswalk may be anomalous during a signal-controlled waiting phase.

The IPM introduces a population-level baseline against which individual risk values are evaluated. The population-level reference state is defined in Equation (3):

$$\mu IPM = E[R_i] \quad (3)$$

denote the mean risk score across the observed pedestrian population. The relative deviation of pedestrian i is then defined in Equation (4):

$$\Delta R_i = R_i - \mu\_IPM \quad (4)$$

A positive deviation indicates above-average integrated risk relative to the scene composition. This formulation does not imply imminent collision; it formalizes contextual unusualness within a multi-agent environment. Such deviation-based reasoning aligns with statistical anomaly detection principles and context-aware modeling strategies [8,21]. By embedding normalization directly into risk interpretation, the IPM adapts to scene composition without requiring explicit parameter retuning across different intersections or traffic densities. The ordered set R_1, R_2, …, R_n naturally induces a ranking structure enabling prioritization.

3.5. Weighted Multi-Factor Instantiation

For operationalizability, a weighted linear aggregation is introduced in Equation (5):

 doi:10.20944/preprints202604.1699.v1

$$R_i = \Sigma_{k=1}^{n} w_k \cdot x_{ik} \quad (5)$$

where x_{ik} denotes the k-th measurable factor associated with pedestrian i, and w_k is the corresponding weight. This structure preserves interpretability by ensuring each factor contributes proportionally, guarantees monotonicity, and supports modular extensibility. The weights may be defined heuristically, calibrated empirically, or estimated through supervised learning. This formulation is also consistent with multi-criteria decision-making (MCDM) theory, where linear value models are widely applied for ranking alternatives under heterogeneous criteria [22]. The mathematical validity of such aggregation structures is further supported by multi-criteria decision-making theory, which establishes linear value models as a foundational approach for ranking alternatives under heterogeneous criteria [22,23].

3.6. Limitations and Scope of the IPM Framework

The IPM does not claim direct prediction of crash probability nor perform causal injury modeling. It does not assert a complete representation of human behavior. Its purpose is to provide a coherent formal structure for integrating heterogeneous safety-relevant indicators into a comparative prioritization mechanism. The present article demonstrates this structure at the level of the IPM framework (Level 1) and the IPM core validation (Level 2) as shown in **Figure 1**. Extension to radar-based detection (Level 3) and full field-trial validation (Level 4) are beyond the scope of this paper.

4. Safety-Prioritized Trajectory Model (SPTM)

4.1. Role Within the IPM

The Safety-Prioritized Trajectory Model (SPTM) provides the first empirical instantiation of the Exposure component of the Intelligent Pedestrian Model (IPM) under a deliberately constrained validation setting. Rather than attempting to realize the full IPM formulation, which includes both Exposure and Severity, the objective of this section is to demonstrate that Exposure alone can be operationalized in a measurable, inductive, and leakage-free manner, and that it can be integrated into trajectory prediction through a minimal and controlled modification of the learning objective.

The SPTM follows a standard probabilistic trajectory prediction paradigm. Given an observed trajectory X, the model predicts a distribution over future trajectories Y using a multimodal formulation. As defined in Equation (6), the predictive distribution is modeled as a K-component Gaussian mixture over future trajectories, where Σ_k denotes a diagonal covariance matrix.

$$p(Y \mid X) = \sum_{k=1}^{K} \pi_k \, \mathcal{N}(Y \mid \mu_k, \Sigma_k) \quad (6)$$

This predictive structure follows standard multimodal trajectory prediction formulations. The contribution of SPTM lies instead in how the model is trained.

An observation-only Exposure proxy $r(X)$ is computed from the input trajectory and normalized to the interval [0,1], yielding $r_{\text{norm}}(X)$. This value defines the sample-wise weighting function given in Equation (7).

$$w(X) = 1 + \beta \cdot r_{\text{norm}}(X) \quad (7)$$

where $\beta \geq 0$ controls the strength of prioritization. When $\beta = 0$, the formulation reduces to standard maximum likelihood training. For $\beta > 0$, trajectories associated with higher Exposure receive proportionally greater influence during optimization.

The resulting training objective is a weighted negative log-likelihood, as defined in Equation (8), where the expectation is taken over the empirical training distribution.

$$\mathcal{L}_{\text{SPTM}} = \mathbb{E}_{(X,Y)}[w(X) \cdot (-\log p(Y \mid X))] \quad (8)$$

 doi:10.20944/preprints202604.1699.v1

This formulation introduces a cost-aware learning mechanism aligned with the prioritization objective of the IPM framework, while leaving the predictive model itself unchanged. In this sense, SPTM modifies standard trajectory prediction in exactly one aspect: the incorporation of Exposure-dependent weighting into the loss function.

To stabilize multimodal learning, an entropy regularization term is added as defined in Equation (9):

$$\mathcal{L}_{\text{final}} = \mathcal{L}_{\text{SPTM}} - \lambda_{\text{ent}}\, H(\pi), H(\pi) = -\sum_{k} \pi_k \log\, \pi_k \tag{9}$$

This regularization is a standard technique to encourage mode utilization and does not constitute part of the IPM-specific contribution.

4.2. Exposure Proxy and Inductive Operationalization

The Exposure component is operationalized through a kinematic proxy computed exclusively from the observed trajectory segment, ensuring strict compliance with inductive evaluation principles. No future information or test-set statistics are used at any stage of the computation. The proxy is constructed from three motion descriptors: velocity magnitude, acceleration magnitude, and directional change, each summarized using the 90th percentile over the observation window. This design emphasizes peak motion intensity while maintaining robustness to noise and transient fluctuations. The resulting scalar is normalized to the interval [0,1] using robust percentile scaling estimated from the training and validation sets only.

This restriction on observation-only kinematics is intentional. It ensures that the Exposure signal reflects measurable motion characteristics available at inference time, avoids information leakage, and provides a controlled basis for evaluating whether a purely kinematic notion of Exposure can induce meaningful structure in the data. Within this formulation, trajectories with higher normalized Exposure values represent motion regimes characterized by increased dynamic intensity or instability. A high-Exposure subset (CRIT) is defined as the upper decile of the training-derived Exposure distribution and is used to evaluate whether the proxy identifies a statistically distinct and more challenging prediction regime.

4.3. Model Architecture and Training Setup

To isolate the effect of Exposure-based weighting, the predictive model is intentionally kept lightweight and structurally conventional. A residual Temporal Convolutional Network (TCN) encoder [24] is employed, followed by a Gaussian mixture density head producing multimodal trajectory forecasts. The architecture is not designed to achieve state-of-the-art displacement accuracy, but rather to provide a stable and computationally efficient backbone within which the impact of the proposed weighting mechanism can be isolated. Such compact architectures are consistent with broader trends toward efficient deep learning models for resource-constrained deployment scenarios [25]. In this setting, the use of a lightweight model (~35,000 parameters) ensures that observed performance variations can be attributed to the training objective rather than architectural complexity.

Training is performed using standard optimization procedures, with Exposure-based weighting applied at the sample level. Prioritization strength is controlled through the parameter β, where $\beta = 0$ corresponds to the baseline and increasing values emphasizing high-Exposure trajectories.

4.4. Experimental Protocol

The complete experimental configuration is summarized in Table 1. Empirical evaluation is conducted on the ETH/UCY pedestrian trajectory benchmark [13,26]using a leave-one-scene-out protocol. All normalization parameters, including Exposure scaling bounds and the CRIT threshold, are determined exclusively from the training and validation sets and remain fixed during testing.

 doi:10.20944/preprints202604.1699.v1

Table 1. Experimental configuration of the Safety-Prioritized Trajectory Model (SPTM).

Component	Configuration
Dataset	ETH/UCY
Split strategy	Leave-one-scene-out
Observation length	8 timesteps (3.2 s)
Prediction horizon	12 timesteps (4.8 s)
Random seeds	5 (seeds 41–45)
β values evaluated	{0, 0.5, 1.0, 2.0}
Metrics	ADE, FDE, ADE@K, FDE@K (CRIT and ALL subsets)
Architecture	TinyTCN, K = 5 Gaussian mixture components, ~35,000 parameters

This strictly inductive setup is critical to the validity of the results. It ensures that the Exposure proxy, the weighting function, and the evaluation subsets are all defined without access to test data, thereby preventing leakage and enabling a clean assessment of generalization. This protocol aligns with established recommendations on reproducibility and evaluation rigor in machine learning research [27].

All experiments are repeated across multiple random seeds to account for stochastic variability in training. Performance is evaluated using standard trajectory prediction metrics (ADE, FDE) as well as best-of-K metrics, which capture multimodal prediction quality. Results are reported separately for the full test set (ALL) and the high-Exposure subset (CRIT).

4.5. Results

The baseline model (β = 0) provides a stable reference consistent with LSTM-class trajectory predictors on the evaluated fold, confirming that the architecture is suitable for controlled comparison. The CRIT subset exhibits higher prediction error than the full test set, indicating that the Exposure proxy identifies a more challenging prediction regime under the inductive normalization scheme. The primary effect of Exposure-based weighting is observed in multimodal prediction quality. Moderate prioritization ($\beta \approx 1.0$) consistently improves best-of-K metrics, indicating enhanced coverage of plausible future trajectories. This effect is stable across random seeds and does not depend on a specific proxy variant.

In contrast, mean displacement metrics (ADE, FDE) show only minor variations within seed-level variability, indicating that the method does not primarily act as a global accuracy enhancer. Instead, it redistributes predictive capacity toward high-Exposure trajectories while maintaining overall performance. Within the CRIT subset, the effect is more nuanced. Improvements in intermediate trajectory accuracy are accompanied by slight increases in endpoint dispersion, reflecting increased uncertainty in long-term predictions. This behavior is consistent with the interpretation that high-Exposure trajectories are inherently less predictable and that increased variance may represent a more appropriate model response.

Overall, the results indicate that Exposure-based weighting induces a systematic shift in model behavior: rather than uniformly optimizing displacement error, the model allocates representational capacity in a manner aligned with safety-relevant motion characteristics. Complete results for all β settings and proxy variants are reported in Table 2.

Table 2. SPTM results on the ETH benchmark (seed = 42, inductive protocol). ★ marks the recommended β = 1.0 setting. — indicates metric not computed.

Proxy	β	Subset	ADE (m)	FDE (m)	ADE@K (m)	FDE@K (m)	NLL	Note
v0	0	ALL	0.930	1.880	0.575	0.982	26.195	baseline
v0	0	CRIT	0.980	1.970	—	—	—	CRIT > ALL ✓
v0	0.5	ALL	0.918	—	0.543	0.926	—	ADE↓ cover↑
v0 ★	1.0	ALL	0.910	—	—	—	—	↑ best ALL ADE

 doi:10.20944/preprints202604.1699.v1

v0 ★	1.0	CRIT	0.969	2.011	0.580	0.994	—	ADE↓ FDE↑
v0	2.0	ALL	0.927	—	—	—	—	↓ over-weighting
v1	0	CRIT	0.977	1.964	0.608	1.045	—	ablation ref.
v1	1.0	CRIT	0.972	2.023	—	0.994	—	pattern robust

Proxy ablation and literature comparison shows both Proxy v0 and v1 at $\beta = 1.0$, the qualitative pattern is consistent: CRIT ADE improves, best-of-K metrics improve, and CRIT FDE increases slightly. The direction of all effects is identical across both proxy definitions, confirming that the accuracy–uncertainty redistribution is a property of the cost-aware mechanism, not an artifact of a specific proxy. The SPTM baseline ($\beta = 0$) achieves ADE = 0.930 m and FDE = 1.880 m on the biwi_eth fold, surpassing the constant velocity baseline (ADE ≈ 1.05 m) and consistent with the LSTM range reported in the ETH/UCY literature [13], confirming that the architecture provides a suitable controlled baseline within the single-fold evaluation scope of this study.

4.6. Interpretation Within the IPM Framework

Within the broader IPM architecture, the SPTM represents a controlled validation of a single Exposure component. The results show that this component can be operationalized using observation-only data, that it defines a meaningful subset of challenging trajectories, and that it can influence learning dynamics through a simple modification of the training objective. Importantly, the full IPM formulation including Severity modeling and reference-based risk calibration—is not instantiated in this section. The present results therefore do not constitute a complete validation of the IPM as a risk model. Rather, they establish that one of its fundamental building blocks can be integrated into a predictive framework in a measurable and methodologically consistent manner. This constrained validation supports the broader hypothesis that risk-aware prioritization can be introduced into trajectory prediction without modifying the underlying predictive architecture, but instead through principled adjustments to the learning objective.

5. Discussion

5.1. Interpretation of Findings

Cost-aware reweighting modifies the allocation of predictive capacity rather than uniformly improving displacement accuracy. Improvements in intermediate trajectory accuracy (ADE) and multimodal coverage (ADE@K, FDE@K) are accompanied by increased endpoint dispersion (FDE) in high-Exposure regimes, reflecting elevated uncertainty in long-term predictions. This behavior is consistent with the intended role of Exposure as a prioritization signal rather than an accuracy optimization target.

5.2. Positioning Within the IPM Architecture

The SPTM instantiates the IPM Exposure component (Level 2 in Figure 1) under three strict boundary conditions: observation-only operationalization, cost-aware optimization, and inductive evaluation. While the IPM is not fully empirically validated here, the results show that its Exposure component can be operationalized in a measurable and leakage-free manner. Within this scoped validation, three findings stand out. First, the CRIT subset exhibits consistently higher prediction error than the full test set under inductive conditions, confirming that the Exposure proxy identifies a statistically distinct and more challenging motion regime. Second, Exposure-based weighting produces consistent best-of-K coverage improvement (ALL FDE@K: 0.979 → 0.970 m across five seeds), without a general accuracy gain — the effect remains within seed-level variance for mean displacement metrics. Third, the SPTM baseline achieves LSTM-class performance within the compact TinyTCN architecture, confirming that the backbone is a suitable controlled reference. As shown in Figure 1, the radar machine learning level (Level 3) and the field trial level (Level 4) are outside the scope of this article. These levels require additional experimental infrastructure, including

 doi:10.20944/preprints202604.1699.v1

77 GHz automotive radar data collection with human vs. clutter separability analysis [28,29] and real-world intersection recordings from proving ground environments such as ZalaZONE [30].

5.3. Limitations

Several limitations should be acknowledged. First, SPTM validation covers a single ETH/UCY benchmark fold; cross-fold generalization with five seeds per fold represents a natural extension. Second, the TinyTCN does not model social interactions, which may limit performance in highly crowded scenarios. Third, the CRIT threshold τ is fold-specific and may require recalibration across domains. Fourth, Severity modeling incorporating pedestrian vulnerability factors (age, mobility impairment) remains an open direction fully consistent with the IPM Equation (2) but not instantiated here.

6. Conclusions

This paper introduced the Intelligent Pedestrian Model (IPM) as a formal framework for representing pedestrian safety relevance through a reference-normalized scalar risk formulation, decomposed into Exposure and Severity components and interpreted as a population-relative deviation. Building on this framework, the Safety-Prioritized Trajectory Model (SPTM) provided a constrained empirical validation of the Exposure component under a strictly inductive, leakage-free protocol. The results demonstrate that Exposure can be operationalized using observation-only trajectory information and integrated into trajectory prediction through a cost-aware learning objective without modifying the underlying predictive architecture. The primary effect of this integration is not a global improvement in displacement accuracy, but a consistent redistribution of predictive capacity toward safety-relevant motion regimes, reflected in improved multimodal coverage (best-of-K metrics). The identification of a high-Exposure subset (CRIT) with increased prediction difficulty further supports the interpretation of Exposure as a proxy for motion complexity under inductive conditions. Future work will extend evaluation across all ETH/UCY folds, incorporate social interaction modeling, and investigate full IPM instantiation including Severity calibration and deployment-level validation.

Supplementary Materials: The following supporting information can be downloaded at the website of this paper posted on Preprints.org.

Author Contributions: Conceptualization, methodology, software, validation, formal analysis, investigation, data curation, writing—original draft preparation, and visualization were performed by Zoltan Rozsas. Supervision, writing—review and editing, and project administration were carried out by Istvan Lakatos. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding. The APC was funded by Széchenyi István University.

Data Availability Statement: The experimental run registry and result files supporting the findings of this study are available as supplementary material. Metrics files: metrics_t3plus_v2_biwi_eth_beta0.0_K5_ep20_seed41.json and metrics_t3plus_v2_biwi_eth_beta1.0_K5_ep20_seed41.json. Run configuration files: run_config_v1_inductive_seeded_biwi_eth_beta0_K5_ep20_seed [41–45].json and run_config_v1_inductive_seeded_biwi_eth_beta1_K5_ep20_seed [41–45].json. The ETH/UCY benchmark data are publicly available at: Pellegrini et al. (2009) [26] and the ETH pedestrian dataset repository.

Conflicts of Interest: The authors declare no conflicts of interest.

Appendix A. Implementation Details

This appendix provides implementation details sufficient to reproduce the Safety-Prioritized Trajectory Model (SPTM) experiments reported in Section 4. All experiments were conducted on the biwi_eth fold of the ETH/UCY benchmark under a strictly inductive evaluation protocol.

 doi:10.20944/preprints202604.1699.v1

Appendix A.1. Training Configuration

All models were trained using a consistent configuration across all experiments. Table A1 summarizes the hyperparameters.

Table A1. Training hyperparameters.

Parameter	Value
Optimizer	Adam (AdamW in ablation variants; no significant differences observed)
Initial learning rate	0.002
Weight decay	1×10^{-4}
Batch size	256
Epochs	20
Gradient clipping	L2 norm ≤ 1.0
Device	CPU (CUDA-compatible)
Random seeds	{41, 42, 43, 44, 45}
LR scheduling	None

Appendix A.2. Model Architecture (Canonical TinyTCN + Mixture Head)

The canonical model — used for all results reported in Table 2 — consists of a residual Temporal Convolutional Network (TCN) encoder followed by a Gaussian mixture density head. Table A2 summarizes the architecture.

Table A2. TinyTCN architecture specification.

Component	Specification
Encoder type	3-layer Temporal Convolutional Network (TCN)
Dilation factors	{1, 2, 4}
Hidden dimension	64
Kernel size	3
Residual connections	Yes (canonical model)
Dropout	p = 0.1 (canonical model)
Pooling	Adaptive average pooling
Output head	Gaussian Mixture Density Network
Mixture components K	5
Covariance structure	Diagonal (dimensions modeled independently)
σ parameterization	log σ (ensures σ > 0 via exp; $\varepsilon = 1 \times 10^{-6}$ for stability)
Prediction horizon T_{pree}	12 timesteps (4.8 s)
Total parameters	35,000

Appendix A.3. Input Representation and Normalization

All trajectories are transformed into a local coordinate frame prior to training and evaluation. The last observed position $x_{t0}{}^{b}{}_{s}$ is used as the origin; both the observed sequence and the future ground-truth trajectory are translated accordingly. This ensures translation invariance and compatibility with the standard ETH/UCY evaluation protocol without introducing information leakage. No rotation normalization or velocity-based alignment is applied.

Appendix A.4. Dataset and Sequence Construction

The ETH/UCY benchmark provides real-world pedestrian trajectory recordings in bird's-eye-view format: (frame_id, pedestrian_id, x, y). Sequences are constructed using a sliding window with stride = 1, independently per pedestrian. Observation length is 8 timesteps (3.2 s); prediction horizon

 doi:10.20944/preprints202604.1699.v1

is 12 timesteps (4.8 s). Typical sequence counts for the biwi_eth fold are training 30,000 sequences, validation 5,400 sequences, and test 300–400 sequences.

Appendix A.5. Exposure Proxy Computation

A leakage-free Exposure proxy is computed exclusively from the observed trajectory $X = \{x_1, \ldots, x_{t_{obs}}\}$. Three kinematic descriptors are derived per trajectory: velocity magnitude (speed s), acceleration magnitude (a), and directional change θ capturing angular instability between consecutive velocity vectors. Each descriptor is summarized as the 90th percentile over the observation window to obtain a robust, noise-resistant scalar.

Two proxy variants are formally defined:

$$r_{v0}(X) = 1.0 \cdot s_{p90} + 0.5 \cdot a_{p90} + 1.0 \cdot \theta_{p90} \text{ (Proxy v0)}$$

$$r_{v1}(X) = 0.7 \cdot s_{p90} + 0.7 \cdot a_{p90} + 0.6 \cdot \theta_{p90} \text{ (Proxy v1, canonical)}$$

where s_{p90} = p90(speed), a_{p90} = p90(acceleration magnitude), θ_{p90} = p90(directional change) over the observation window. Proxy v1 was motivated empirically: Proxy v0 populated the CRIT subset disproportionately with high-curvature trajectories; rebalancing toward acceleration-dominated cases yielded more stable endpoint performance.

The raw score is normalized using robust percentile scaling:

$$r_{norm}(X) = clip_{=}(r(X) - p_{05}) / (p_{95} - p_{05}), 0, 1$$

where p_{05} and p_{95} are the 5th and 95th percentiles estimated exclusively from the combined training and validation sets. These bounds are fixed prior to test evaluation; test-set statistics are never used in normalization. The CRIT subset is defined as all test trajectories satisfying $r_{norm}(X) > \tau$, where $\tau = 0.7792$ is the 90th percentile of the training+validation Exposure distribution (biwi_eth fold, Proxy v1).

Appendix A.6. Training Objective and Loss Formulation

The predictive distribution over future trajectories are modeled as a K-component Gaussian mixture:

$$p(Y \mid X) = \Sigma^{k}_{k=1} \pi_k \cdot {}^{l}(Y \mid \mu_k, \Sigma_k) \text{ (A1)}$$

A diagonal covariance structure is used, modeling each spatial dimension independently. Standard deviations are parameterized as $\sigma_{k,}{}^{d} = \exp(\acute{s}_{k,}{}^{d}) + \varepsilon$, where $\varepsilon = 1 \times 10^{-6}$ ensures numerical stability. The base negative log-likelihood loss is:

$$\mathcal{L}_{nll} = -log(\Sigma^{k}_{k=1} \pi_k \cdot {}^{l}(Y \mid \mu_k, \Sigma_k)) \text{ (A3)}$$

Cost-aware sample weighting reallocates gradient emphasis toward high-Exposure trajectories:

$$w(X) = 1 + \beta \cdot r_{norm}(X) \text{ (A4)}$$

The weighted training objective is:

$$\mathcal{L}_{sptm} = \mathbb{E}[\, w(X) \cdot \mathcal{L}_{nll} \,]$$

An entropy regularization term is added to the mixture weights to prevent mode collapse:

$$\mathcal{L}^{fl}_{nal} = \mathcal{L}_{sptm} - \lambda_{ent} \cdot H(\pi), H(\pi) = -\Sigma_k \pi_k \log \pi_k, \lambda_{ent} = 0.01 \text{ (A6)}$$

The entropy term $H(\pi)$ encourages balanced utilization of mixture components, mitigating collapse to a single dominant mode.

Appendix A.7. Evaluation Protocol

The evaluation protocol follows the standard leave-one-scene-out strategy on ETH/UCY. All normalization parameters (Exposure percentile bounds, CRIT threshold τ) and all model hyperparameters are determined using training and validation data only. The test set is used exclusively for final evaluation. Experiments are repeated across five random seeds (41–45)

 doi:10.20944/preprints202604.1699.v1

controlling weight initialization, data ordering, and stochastic training components. Performance is reported as mean ± standard deviation across seeds.

Metrics are computed separately for the full test set (ALL) and the high-Exposure subset (CRIT, top 10% by r_{norm}). Best-of-K metrics (ADE@K, FDE@K) use K = 5 trajectories sampled from the mixture, selecting the closest mode to ground truth by final displacement.

Appendix A.8. Reproducibility

Reproducibility is ensured through the following mechanisms:

- Fixed random seeds ({41–45}) controlling weight initialization, data ordering, and stochastic dropout.
- Deterministic dataset splits: scene-level partitioning is fixed; no randomization across seeds.
- Fixed normalization bounds: Exposure proxy percentiles (p_{05}, p_{95}) and CRIT threshold τ are computed once from the training and validation set and applied unchanged across all seeds.
- JSON logging: each run produces a metrics file and a configuration file uniquely identified by fold, β value, K, epoch count, and seed. The complete run registry for the biwi_eth fold is listed in Table A3.

Table A3. Experimental run registry (biwi_eth fold, Proxy v1, inductive protocol, τ = 0.7792).

File name	β	K	Ep.	Seeds
metrics_t3plus_v2_biwi_eth_beta0.0_K5_ep20_seed [41–45].json	0	5	20	41–45
metrics_t3plus_v2_biwi_eth_beta1.0_K5_ep20_seed [41–45].json	1.0	5	20	41–45
run_config_v1_inductive_seeded_biwi_eth_beta0_K5_ep20_seed [41–45].json	0	5	20	41–45
run_config_v1_inductive_seeded_biwi_eth_beta1_K5_ep20_seed [41–45].json	1.0	5	20	41–45

Each metrics file contains the following fields: fold, beta, K, epochs, seed, risk_p05, risk_p95_trainval, crit_thr, and per-metric results (ADE, FDE, ADE@K, FDE@K, NLL) for both the ALL and CRIT subsets. Configuration files additionally record all model hyperparameters and training settings as specified in Tables A1 and A2.

Appendix A.9. Multi-Seed Variance Summary

Table 2 in the main text reports results for a single representative seed (seed 42) for clarity of presentation. Table A4 below provides the full five-seed mean ± standard deviation for primary metrics at the recommended setting β = 1.0 and the unweighted baseline β = 0, evaluated on the biwi_eth fold. These figures substantiate the seed-level variance claim in Section 4.5 and confirm that the best-of-K improvement is consistent across all seeds.

Table A4. Five-seed mean ± standard deviation for key metrics (biwi_eth fold, Proxy v1, inductive protocol). ALL = full test set; CRIT = top-10% Exposure subset.

β	Subset	ADE (m)	FDE (m)	ADE@K (m)	FDE@K (m)
0	ALL	0.930 ± 0.013	1.899 ± 0.031	0.575 ± 0.014	0.979 ± 0.027
0	CRIT	0.980 ± 0.018	1.988 ± 0.044	—	—
1.0	ALL	0.918 ± 0.012	1.891 ± 0.028	0.553 ± 0.012	0.970 ± 0.025
1.0	CRIT	0.969 ± 0.021	1.999 ± 0.046	0.577 ± 0.016	0.998 ± 0.031

Variance across seeds is comparable between β = 0 and β = 1.0 for all metrics, confirming that cost-aware weighting does not introduce training instability. The primary effect — improvement in best-of-K multimodal coverage (ALL FDE@K: 0.979 → 0.970 m) — is consistent across all five seeds.

Appendix A.10. Supplementary Data Files

 doi:10.20944/preprints202604.1699.v1

The following output files from the Python training and evaluation pipeline constitute the empirical basis of the results reported in this article and should be submitted as supplementary material alongside the manuscript.

Metrics files (one per run, JSON format):

- metrics_t3plus_v2_biwi_eth_beta0.0_K5_ep20_seed41.json
- metrics_t3plus_v2_biwi_eth_beta1.0_K5_ep20_seed41.json

Run configuration files (one per seed per β value, JSON format):

- run_config_v1_inductive_seeded_biwi_eth_beta0_K5_ep20_seed41.json
- run_config_v1_inductive_seeded_biwi_eth_beta0_K5_ep20_seed42.json
- run_config_v1_inductive_seeded_biwi_eth_beta0_K5_ep20_seed43.json
- run_config_v1_inductive_seeded_biwi_eth_beta0_K5_ep20_seed44.json
- run_config_v1_inductive_seeded_biwi_eth_beta0_K5_ep20_seed45.json
- run_config_v1_inductive_seeded_biwi_eth_beta1_K5_ep20_seed41.json
- run_config_v1_inductive_seeded_biwi_eth_beta1_K5_ep20_seed42.json
- run_config_v1_inductive_seeded_biwi_eth_beta1_K5_ep20_seed43.json
- run_config_v1_inductive_seeded_biwi_eth_beta1_K5_ep20_seed44.json
- run_config_v1_inductive_seeded_biwi_eth_beta1_K5_ep20_seed45.json

Each metrics file contains the fields: fold, beta, K, epochs, seed, risk_p05, risk_p95_trainval, crit_thr, and per-metric results (ADE, FDE, ADE@K, FDE@K, NLL) for the ALL and CRIT subsets. Configuration files additionally record all model hyperparameters and training settings as specified in Tables A1 and A2.

References

1. World Health Organization. Global Status Report on Road Safety 2023. Geneva; 2023.
2. Eurostat. Road Safety Statistics in the EU - Road Accident Fatalities by Category of Persons Involved [Internet]. 2024. Available from: https://ec.europa.eu/eurostat/
3. European Commission. EU Road Safety Policy Framework 2021-2030 – Next Steps Towards "Vision Zero." Brussels; 2019.
4. Cicchino JB. Effectiveness of Forward Collision Warning and Autonomous Emergency Braking Systems in Reducing Front-to-Rear Crash Rates. Accid Anal Prev. 2017;99:142–52. doi:10.1016/j.aap.2016.11.009
5. Fildes B, Keall M, Bos N, Lie A, Page Y, Pastor C, et al. Effectiveness of Low Speed Autonomous Emergency Braking in Real-World Rear-End Crashes. Accid Anal Prev. 2015;81:24–9. doi:10.1016/j.aap.2015.03.029
6. SAE International. Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles. 2021.
7. Helbing D, Molnár P. Social Force Model for Pedestrian Dynamics. Phys Rev E. 1995;51(5):4282–6. doi:10.1103/PhysRevE.51.4282
8. Alahi A, Goel K, Ramanathan V, Robicquet A, Fei-Fei L, Savarese S. Social LSTM: Human Trajectory Prediction in Crowded Spaces. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). 2016. p. 961–71. doi:10.1109/CVPR.2016.110
9. Ivanovic B, Leung K, Schmerling E, Pavone M. Multimodal Deep Generative Models for Trajectory Prediction: A Conditional Variational Autoencoder Approach. IEEE Robot Autom Lett. 2020;6(2):295–302. doi:10.1109/LRA.2020.3035803
10. Yin Z, Liu R, Xiong Z, Yuan Z. Multimodal Transformer Network for Pedestrian Trajectory Prediction. In: Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI). 2021. p. 4121–7. doi:10.24963/ijcai.2021/567
11. Ivanovic B, Pavone M. The Trajectron++: Dynamically-Feasible Trajectory Forecasting With Heterogeneous Data. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). 2019. p. 2375–84. doi:10.1109/ICCV.2019.00246

 doi:10.20944/preprints202604.1699.v1

12. Gupta A, Johnson J, Fei-Fei L, Savarese S, Alahi A. Social GAN: Socially Acceptable Trajectories with Generative Adversarial Networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). 2018. p. 2255–64. doi:10.1109/CVPR.2018.00240
13. Rudenko A, Palmieri L, Arras KO. Human Motion Trajectory Prediction: A Survey. Int J Rob Res. 2020;39(8):895–935.
14. Mo X, Huang Z, Xing Y, Lv C. Multi-Agent Trajectory Prediction With Heterogeneous Edge-Enhanced Graph Attention Network. IEEE Transactions on Intelligent Transportation Systems. 2022;23(7):9554–67. doi:10.1109/TITS.2022.3146300
15. Zheng L, Ismail K, Meng X. Traffic conflict techniques for road safety analysis: Open questions and some insights. Accid Anal Prev. 2020;141:105530. doi:10.1016/j.aap.2020.105530
16. Shalev-Shwartz S, Shammah S, Shashua A. On a Formal Model of Safe and Scalable Self-Driving Cars. arXiv preprint arXiv:170806374. 2017.
17. Rasouli A, Kotseruba I, Tsotsos JK. Understanding Pedestrian Behavior in Complex Traffic Scenes. IEEE Transactions on Intelligent Vehicles. 2019;3(1):61–70.
18. Rasouli A, Kotseruba I, Kunic T, Tsotsos JK. PIE: A Large-Scale Dataset and Models for Pedestrian Intention Estimation and Trajectory Prediction. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). 2019. p. 6262–71. doi:10.1109/ICCV.2019.00636
19. Zhang S, Abdel-Aty M, Yuan J, Li P. Prediction of Pedestrian Crossing Intentions at Intersections Based on Long Short-Term Memory Recurrent Neural Network. Transp Res Rec. 2020;2674(4):57–65. doi:10.1177/0361198120910938
20. Chou FC, Lin TH, Cui H, Radosavljevic V, Nguyen T, Huang TK, et al. Predicting Motion of Vulnerable Road Users Using High-Definition Maps and Efficient ConvNets. In: IEEE Intelligent Vehicles Symposium (IV). 2020. p. 1527–34. doi:10.1109/IV47402.2020.9304830
21. Kooij JFP, Flohr F, Pool EAI, Gavrila DM. Context-Based Path Prediction for Targets with Switching Dynamics. Int J Comput Vis. 2019;127(3):239–62. doi:10.1007/s11263-018-1104-4
22. Mardani A, Jusoh A, Nor KM, Khalifah Z, Zakwan N, Valipour A. Multiple criteria decision-making techniques and their applications – A review of the literature from 2000 to 2014. Economic Research-Ekonomska Istraživanja. 2015;28(1):516–71. doi:10.1080/1331677X.2015.1075139
23. Triantaphyllou E. Multi-Criteria Decision Making Methods: A Comparative Study. Dordrecht, The Netherlands: Springer; 2000. doi:10.1007/978-1-4615-4459-3
24. Bai S, Kolter JZ, Koltun V. An Empirical Evaluation of Generic Convolutional and Recurrent Networks for Sequence Modeling. arXiv preprint arXiv:180301271 [Internet]. 2018. Available from: https://arxiv.org/abs/1803.01271
25. Mittal A. A Survey of Energy-Efficient Deep Learning Techniques. IEEE Access. 2020;8:163210–24. doi:10.1109/ACCESS.2020.3022543
26. Pellegrini S, Ess A, Schindler K, Gool L Van. You'll Never Walk Alone: Modeling Social Behavior for Multi-Target Tracking. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV). 2009. p. 261–8. doi:10.1109/ICCV.2009.5459260
27. Pineau J, Vincent-Lamarre P, Sinha K, Lariviere V, Beygelzimer A, d'Alche-Buc F, et al. Improving Reproducibility in Machine Learning Research. Journal of Machine Learning Research. 2021;22(164):1–20.
28. Harlow K, Jang H, Barfoot TD, Kim A, Heckman C. A New Wave in Robotics: Survey on Recent mmWave Radar Applications in Robotics. 2023. doi:10.48550/arXiv.2305.01135
29. Wei Z, Zhang F, Chang S, Liu Y, Wu H, Feng Z. mmWave Radar and Vision Fusion for Object Detection in Autonomous Driving: A Review. arXiv preprint arXiv:210803047. 2021.
30. ZalaZONE. ZalaZONE Automotive Proving Ground – Technical Documentation [Internet]. Zalaegerszeg, Hungary; 2024. Available from: https://www.zalazone.hu/