A Trend-Guided Multi-Fidelity Framework for Efficient and Accurate Occupant Injury Prediction

Tianle Lu, Yi Yang, Jialong Chang, Zijian He, Qing Zhou, Bingbing Nie

I. INTRODUCTION

Accurate occupant injury prediction is critical for optimising restraint systems and informing safety decisions in unavoidable collisions [1]. However, data-driven modelling is fundamentally constrained by the scarcity of high-quality data. High-fidelity (HF) finite element (FE) simulations offer superior biofidelity but are computationally prohibitive, limiting data scale. Conversely, low-fidelity (LF) simulations (e.g. multi-body) are efficient and capture reliable injury trends, yet introduce systematic biases [2]. Existing methods struggle to bridge this gap as pure data-driven models tend to overfit on sparse HF data, and direct fusion often fails due to differences in data structure. To address this, we propose a trend-guided multi-fidelity network (T-MFNet), which extracts injury trends from abundant LF data to guide learning from sparse HF data. This design enables HF-level prediction accuracy with significantly reduced computational cost, facilitating scalable and data-efficient injury prediction.

II. METHODS

The proposed T-MFNet is built on a 'trend + residual' decomposition strategy. Instead of learning the complex mapping from scratch, the model uses LF data (L_l) to establish a baseline injury trend and HF data (L_h) to learn a refined residual correction. The framework comprises four key modules (Fig. 1):

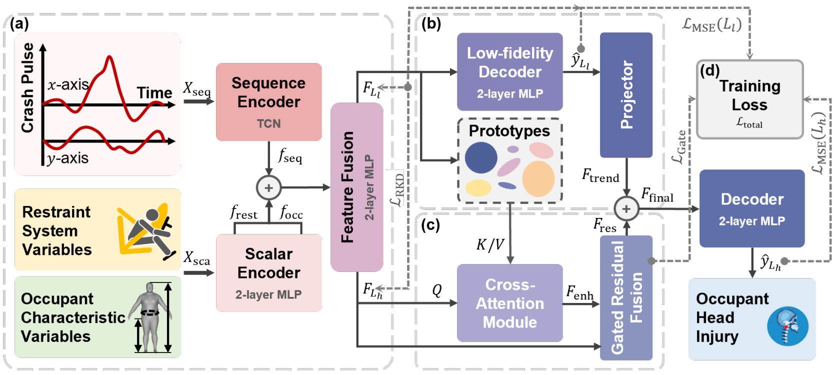

Fig. 1. The proposed T-MFNet for occupant injury prediction.

(a) Shared Multi-Modal Feature Encoding: both L_l and L_h inputs are processed through an identical encoding architecture to map heterogeneous data, including time-series crash pulses X_{seq} and scalar variables X_{sca}, into a unified feature space, yielding initial feature representations F_{L_l} and F_{L_h}. Specifically, a temporal convolutional network (TCN) extracts kinematic features f_{seq} from bi-axial vehicle crash pulses. Simultaneously, a group-wise multi-layer perceptron (MLP) encodes scalar parameters, including restraint system variables f_{rest} (e.g. seat-belt usage, time to fire, and belt load-limiter force) and occupant characteristics f_{occ} (e.g. dummy size, type and sitting posture), preserving their distinct physical semantics before fusion.

(b) LF-Prior Retrieval & Trend Extraction: this module extracts knowledge from the abundant LF data. First, a learnable prototype bank is constructed to capture typical crash patterns and injury monotonicity [1] from the LF dataset, serving as a reference library. Second, an explicit trend injection mechanism utilises the LF decoder to generate a preliminary injury prediction $\hat{y}_{L_l}$. This prediction acts as a coarse injury trend, which is then projected back into the feature space to serve as a reference foundation F_{trend} for the subsequent residual learning.

(c) Trend-Guided Residual Learning: to correct the systematic bias of the LF trend, the HF branch learns a residual correction term. A prototype cross-attention mechanism matches the current HF collision condition (Query, Q) with similar patterns in the LF reference library (Key/Value, K/V). This enables the model to 'look up' relevant prior knowledge for unseen HF cases, compensating for data sparsity. Subsequently, a gated residual fusion module computes the residual feature F_{res} by adaptively weighting the contributions of observed HF features and retrieved priors F_{enh}. This mechanism allows the model to selectively rely on LF priors or HF

B. Nie (e-mail: nbb@tsinghua.edu.cn) is an Associate Professor, T. Lu is a PhD student, Y. Yang and Z. He are MSc students and Q. Zhou is a Professor, all in the School of Vehicle and Mobility at Tsinghua University, China. J. Chang is a MSc student in the School of Technology at Beijing Forestry University, China.

observations to achieve optimal prediction performance. The final feature F_{final} is obtained by superimposing this residual correction onto the baseline trend:

$$F_{\text{final}} = F_{\text{trend}} + F_{\text{res}} \tag{1}$$

Finally, a decoder maps F_{final} to the HF occupant head injury prediction $\hat{y}_{L_h}$ (Head Injury Criterion, HIC).

(d) Composite Loss Function: the model is trained with a multi-objective loss function $\mathcal{L}_{\text{total}}$. In addition to the standard regression loss ($\mathcal{L}_{\text{MSE}}$) for prediction accuracy, we introduce a relational knowledge distillation (RKD) loss ($\mathcal{L}_{\text{RKD}}$) to guide cross-domain alignment [3]. Specifically, we hypothesise that if two scenarios are structurally similar in the LF domain, they should remain proximate in the HF domain. Furthermore, a gating regularisation loss ($\mathcal{L}_{\text{Gate}}$) is included to encourage the model to adaptively balance the weights between prior knowledge and current observations, preventing overfitting to either source.

III. INITIAL FINDINGS

Using two simulation datasets (4,849 cases from MADYMO and 3,598 cases from FE [4]), we compared the T-MFNet with a single-fidelity baseline (TCN-MLP) and a standard transfer learning model (TL). HIC predictions were also mapped to AIS levels for classification evaluation. With all L_l cases and 200 L_h cases, T-MFNet outperforms both baselines (Table I), reducing RMSE by 11.8% compared to TCN-MLP and MAE by 9.6% compared to TL. Crucially, it achieves the highest G-mean (0.623), with predictions more concentrated along the diagonal of the confusion matrix (Fig. 2), indicating reduced underestimation bias for high-risk injuries. Figure 3 further highlights superior sample efficiency, where T-MFNet trained on just 200 HF samples (MAE ≈ 104) matches the performance of the TCN-MLP baseline trained on 500 samples (MAE ≈ 99), suggesting a ~60% saving in the high-cost L_h simulation budget. This indicates that abundant LF data capture macroscopic injury-related loading-response trends, allowing sparse HF data to refine fidelity-specific nonlinear responses. Consequently, the proposed model effectively bridges the fidelity gap, enabling accurate prediction even with minimal HF data.

TABLE I
PREDICTION PERFORMANCE OF THE THREE MODELS

Model	HIC: RMSE (↓) / MAE (↓) / R^2	AIS: Accuracy / G-mean / F1-score
TCN-MLP	205.838	0.677
	121.490	0.600
	0.779	0.625
TL	190.119	0.677
	115.720	0.597
	0.812	0.619
T-MFNet	**181.597**	**0.699**
	104.619	**0.623**
	0.828	**0.642**

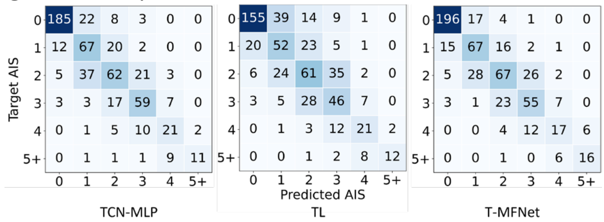

Fig. 2. Head injury AIS level confusion matrices.

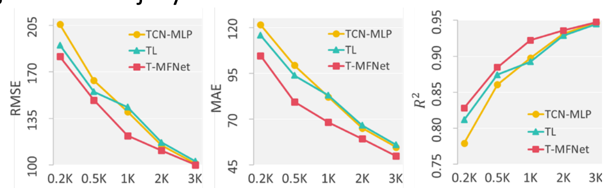

Fig. 3. Evaluation with different training set sizes.

IV. DISCUSSION

Efficient and accurate injury prediction is critical for restraint optimisation and safety decision-making, yet current models require large HF datasets that are unavailable early in development. T-MFNet addresses this by decomposing prediction into an LF-derived physics-based trend and an HF residual correction, improving sample efficiency and robustness. Despite these advances, the framework implicitly assumes that LF data preserve broadly consistent injury trends for HF prediction. Moreover, the current evaluation is limited to rear-end collisions and requires validation in other crash scenarios. Overall, the proposed framework offers a cost-efficient approach for occupant injury prediction and supports virtual safety assessment in HF data-scarce settings.

V. ACKNOWLEDGEMENTS

This work was supported by the National Natural Science Foundation of China (52572413).

VI. REFERENCES

[1] Lu, T., *et al.*, *IRCOBI*, 2025.
[2] Fahlstedt, M., *et al.*, *Traffic Inj Prev*, 2016.
[3] Park, W., *et al.*, *CVPR*, 2019.
[4] Wang, Q., *et al.*, *J Biomech Eng*, 2024.