SSRG International Journal of Electronics and Communication Engineering
ISSN: 2348-8549/ https://doi.org/10.14445/23488549/IJECE-V13I1P110

Volume 13 Issue 1, 107-126, January 2026

Original Article

Integrating Reinforcement Learning and Contextual Analysis for Enhanced Pedestrian Behavior Prediction

Giribabu Sadineni[1], M. Sri Lakshmi[2], Srinivasa Rao Madala[3], P.V.V.S.D Nagendrudu[4], Buradagunta Swathi Sri[5], Naresh Kumar Bhagavatham[6]

[1]*Department of CSE, PACE Institute of Technology and Sciences, Ongole, Andhra Pradesh, India.*
[2]*Department of Computer Science and Engineering, G.Pullaiah College of Engineering and Technology (Autonomous), Kurnool, Andhra Pradesh, India.*
[3]*Department of Artificial Intelligence and Data Science, Chaitanya Bharathi Institute of Technology(A), Gandipet, Hyderabad, Telangana, India.*
[4]*Department of AI&ML, Aditya University, Surampalem, Andhra Pradesh, India.*
[5]*Department of CSE, Koneru Lakshmai Education Foundation, Vaddeswaram, Andhra Pradesh, India.*
[6]*Associate Professor, Department of CSE, Vignana Bharathi Institute of Technology, Ghatkesar, Hyderabad, India.*

[2]*Corresponding Author : srilakshmicse@gpcet.ac.in*

Received: 06 November 2025 Revised: 08 December 2025 Accepted: 07 January 2026 Published: 14 January 2026

Abstract *- The rapid growth and complexity of urban traffic necessitate advanced solutions to enhance pedestrian safety, particularly at zebra crossings. This research presents a comprehensive framework that integrates Reinforcement Learning (RL) and context-aware prediction to achieve accurate, real-time pedestrian behavior forecasting. The proposed model is dynamic as it uses deep learning solutions, such as Deep Q-Learning (DQN) and Proximal Policy Optimization (PPO), and adjusts to the changing urban environment based on the given context of the time of the day, weather conditions, and traffic lights. Evaluation of the model based on the Pedestrian Intention Estimation (PIE) dataset shows that the developments of the model have strong performance, with an accuracy of 92.5, a precision of 91.1, a recall of 90.3, and an F1 score of 90.7. It performs better than the current models regarding computational efficiency and faster inferences for real-time deployment with reduced memory consumption. Adaptive learning in constant updates and smooth integration with the traffic management systems towards proactive safety approaches are primary contributions. The framework, in addition to improving traffic safety in the cities and aiding in the development of autonomous vehicle systems, also identifies other areas for future research topics, such as controlling the unpredictable actions of pedestrians and managing consumption rate aspects at a larger scale of challenges.*

Keywords *- Pedestrian behavior prediction, Reinforcement learning, Context-aware prediction, Urban traffic safety, Real-time analysis, Autonomous vehicles.*

1. Introduction

Prediction of the actions of pedestrians at a zebra crossing is a required element of the safety of pedestrians in city traffic and the development of opportunities for self-driving vehicles. [1]. This has evolved over the years, and nowadays one can find intricate models using computer vision and machine learning as opposed to crude methods of observation. The research has been anchored on the earlier research, including one in which a bottom-up pose-estimation framework utilizing Convolutional Neural Networks (CNNs) was used to classify the poses and movement of pedestrians. [2]. Further, the transfer learning, along with the use of pre-trained models in the context of Predictive Pedestrian Analytics on Safety Enhancement (PPASE) [3], has also helped in the framework of real-time prediction of pedestrian behaviors. These measures have enhanced our intelligence and capacity in situations of predicting the actions of humans on the road; there are several traps that may be required by prior studies and creative use.

Giving real-time relevance through computer complexity reduction, data speed, and reduction of latency; balancing accuracy in the various environmental factors determined by weather, evolving urban infrastructure, and unpredictable human behaviour; and giving applicability in terms of different designs and infrastructures, and data given by different urban environments are the main selected dilemmas. The presented paper [4] incorporates those difficulties with the help of reinforcement learning and context-sensitive prediction, applying adaptive learning methods to the increased accuracy of real-time prediction, paying attention to the involvement of the contextual data and the possibility of obtaining more accurate models applicable in different urban settings [5].

Pedestrian predictability at zebra crossings is a significant issue in traffic safety in urban areas and the development of autonomous cars. The field has evolved throughout the years because people used straightforward methods of observation to sophisticated ones that entail computer vision and machine learning [6]. The recent developments in the prior research areas, including the development of a bottom-up pose-estimation model based on Convolutional Neural Networks (CNNs) to identify the pedestrian poses and motions, have furnished the backdrop of the present research [7]. In addition, the transfer learning and pre-trained models in the Predictive Pedestrian Analytics framework application have an advanced real-time forecast of pedestrian performance. These attempts have improved our knowledge and capabilities in anticipating the directions of pedestrian behavior, but, nonetheless, there are several issues that still have to be examined and proposed through innovation.

The difficult issues are assurance of real-time applicability through computational complexity surmounting, speed of data processing, and reduced latency; accuracy under a wide range of environmental conditions, due to changing weather conditions and other dynamic urban environments, and variability in designs, infrastructures, and data availability in different urban contexts. The solutions of this paper are combining the concepts of reinforcement learning and context-sensitive methods of prediction, utilizing adaptive learning algorithms to produce better real-time predictions, using context-effect information to generate more accurate predictions, and making sure that the models are flexible and generalizable in multiple urban settings. This is also done to provide a strong improvement in pedestrian safety and behavior forecasting on zebra crossings.

Even with the gains made, the modern models still face argumentative constraints, primarily in real-time applicability, accuracy in the different environmental settings, and scalability in the dissimilar city environments. Even though the existing methods are relatively effective in the investigation of approach accuracy under a control situation, they have a general failure in dynamic and heterogeneous urban conditions. More than that, the models are not even scalable, and, therefore, their implementation in real-life situations is prohibited in multiple ways. These problems therefore raise a need to have more flexible and context-sensitive prediction of pedestrian behavior methodologies.

To establish a superior model that can be expected to understand pedestrian behavior at a zebra crossing to retain safety is the specific goal of the research. The suggested framework will be supported on the principles of reinforcement learning and context-sensitive prediction strategies that will add flexibility to the model, together with precision. This research will answer the following research questions:

1. How can reinforcement learning be effectively applied to predict pedestrian behavior in real-time?
2. What role does contextual data play in enhancing the accuracy of pedestrian behavior predictions?
3. How can the proposed framework be scaled to accommodate diverse urban environments?

1.1. Key Contributions of the Paper

1. Reinforcement Learning and Context-Aware Integration for Enhanced Predictive Accuracy: This research introduces a novel framework that synergizes Reinforcement Learning (RL) algorithms with context-aware prediction models to improve the precision and robustness of real-time pedestrian behavior forecasts at zebra crossings. This combination takes advantage of environmental and pedestrian-specific data, which leads to a more detailed insight into the way pedestrians make decisions. [8]
2. Adaptive Learning Mechanisms to Mitigate Computational Challenges: The proposed model incorporates adaptive learning techniques that continuously update parameters with newly acquired data, effectively addressing challenges related to computational complexity and latency. This mechanism ensures that real-time predictions remain precise and responsive, vital for practical deployment in dynamic traffic environments [9, 10].
3. Contextual Data Integration for Comprehensive Behavior Prediction: The integration of time of day, weather conditions, and traffic lights in the RL model of prediction, the research outlines a strong strategy towards prediction that ensures a large improvement in prediction. The predictive system is more realistic and reliable since all the information in the present circumstance will be considered in the predictive system, which would mirror the actual behavior on the road of a pedestrian in various circumstances [11].

1.2. Significance of the Study

The results of the study will significantly affect the safety and appropriateness of urban traffic and the implementation of vehicle systems during autopilot operations. Better forecasts of the behavior of pedestrians that are stronger and more accurate, a better proposed framework will not only improve the performance of the traffic management systems, but it will also lead to the safety of the pedestrians compared to their vehicles.

This has been among the major compositions behind smart solutions of urban traffic and the successful implementation of autonomous vehicles within advanced cities. It is represented in the research paper that addresses the gaps that still exist in predicting pedestrian behavior with the intention of selling them to establish new standards of predicting them to offer efficient and safer traffic through the fixed route that is the town.

The paper is split into the following sections: Section 2 involves the literature review to identify the scientific fields that theorize about predicting the behaviour of pedestrians, reinforcement learning, and the context-aware system to fill the significant gaps and gaps in the knowledge in a manner that makes sense and sensational conclusions. Section 3 identifies the overall strategy that the proposed architecture pursues and explains the methods by which methods of reinforcement learning are combined with context-sensitive prediction processes.

This analysis is reflected in Title 4, which depicts the research findings and performance parameters, as well as a comparison with available models. Its importance, the possible difficulties, and constraints on future research can be identified in the discussion of the results in Section 5. Discussion of contributions to be summarized and research guidelines to be given in the future are concluded in Section 6.

2. Literature Review

The prediction of pedestrian behavior represented one of the most popular topics, as it is rather important to predict this kind of behaviour to make the traffic in the cities safer, as well as to create autonomous vehicle systems. The latest developments used deep learning and computer vision to estimate pedestrian motion better.

Various approaches have been used, including Convolutional Neural Networks (CNNs) [12], Recurrent Neural Networks (RNNs) [13], and Long Short-Term Memory (LSTM) networks [14]. These models hold great potential for the accurate trajectory prediction of pedestrians through their training on huge datasets of pedestrian trajectories.

2.1. Reinforcement Learning and Context-Aware Systems

While previous studies have used RL [15] to enhance pedestrian behavior prediction, recent algorithms apply additional context-aware technology. DQN and PPO, as RL algorithms, have been utilized to learn the best policies for predicting behaviors through environmental interaction [16]. These are so because such algorithms allow the models to adapt to fit into emerging information and enhance prediction accuracy with time.

2.2. Identifying Gaps in Current Research

Although the latest advances in the field of pedestrian behavior prediction have been made, there are still some gaps. The first might be real-time applicability, and in this case, models must perform with high speed but with significant delay. Existing models usually have issues with computational complexity and data processing speed, and thus are not applicable when realistic time constraints need to be met. In addition, the issue that brings out high accuracy of forecasting in different environmental conditions holds. The models should stand strong to withstand changes in weather, lights, and the urban dynamics that can essentially change the behavior of pedestrians.

3. Methodology

Traffic congestion malpractice in urban settings demands superior strategies for enhancing pedestrian safety in zebra crossings. This paper introduces a unified approach applying methods of Reinforcement Learning (RL) and the ability to predict the behavior of pedestrians using the context area accurately and in real-time. The reinforced learning framework, context-aware prediction, and real-time prediction and analysis make up the three major parts of the approach. [17].

3.1. Reinforcement Learning Framework

An RL environment is configured in which the state is the existing observations of pedestrians, possible predictive decisions are realized by actions, and rewards are received based on prediction accuracy. Advanced RL models such as DQN and PPO are utilized to learn policies. The model allows historical data to provide policies that can be used in various scenarios successfully.

Context-Aware Prediction: This feature uses information about the time of day, weather, and road lights during the prediction model. The model enhances its predictive accuracy by deriving useful features. To allow models to be highly responsive and become flexible regarding dynamically changing statements of real-time information, it relies on special context modules to render predictions made by the model dynamically.

Real-Time Prediction and Analysis: It uses a real-time prediction development, which involves a task that involves using contextual information and provides an accurate and real-time prediction of the action of an individual by applying learned controls of an RL system.

Dynamic mechanisms update the prediction using new conditions and data that change. Integration of the predictive analytics and traffic management systems will help to implement and realize proactive safety control, including driver warnings and adjusting traffic lights based on the location as a response to control any possible conflict and maximize the safety of passengers [18].

Both methods create an elastic predictive framework for the variants of pedestrian behavioral forecasts, which result in safer and smoother traffic in urban areas. It increases the sensor capabilities of self-driving vehicles and contributes to the appearance of the smart traffic control system, where pedestrians are prioritized.

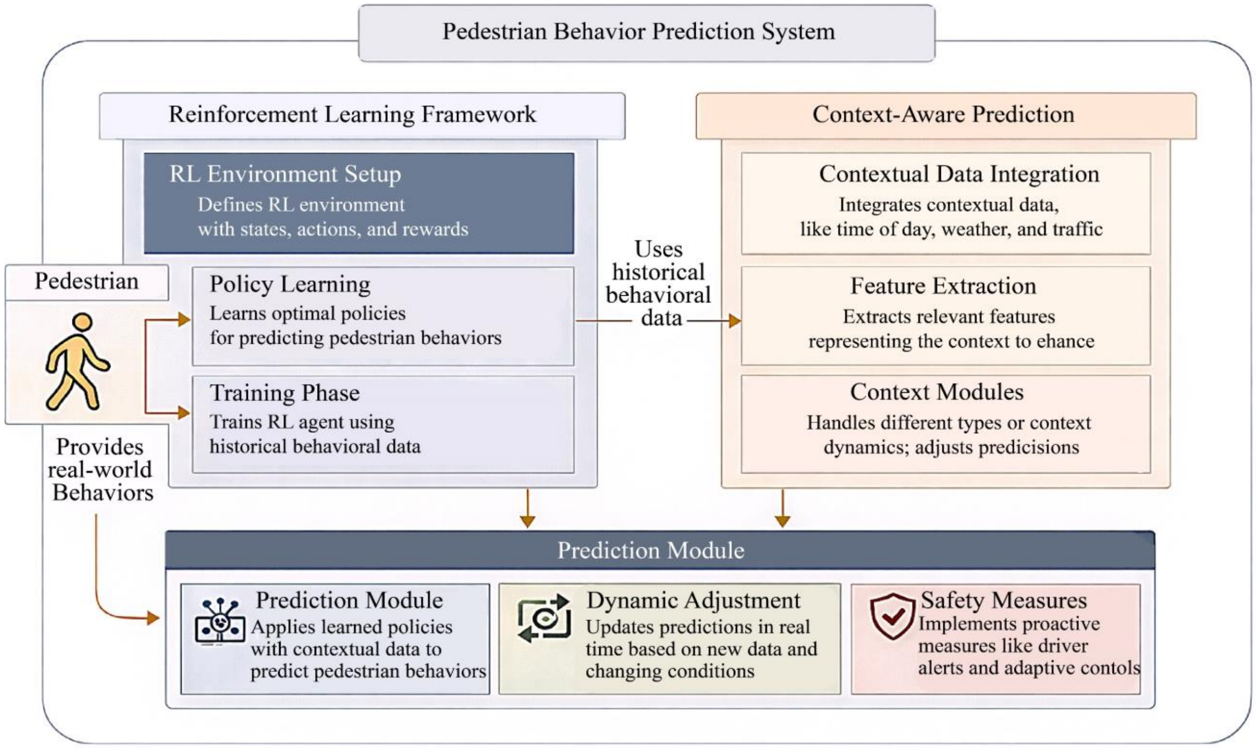

Fig. 1 Pedestrian behavior prediction system architecture

The research methodology used in this study would only combine Reinforcement Learning (RL) and context-sensitive prediction methods in enhancing the prediction of pedestrian behavior at zebra crossings. It is in this section that the RL framework, including environment setup, policy learning, and the stage of training, is explained. This is aimed at designing a real-time predictive model of pedestrian behaviour under adaptive prediction, which involves the Pedestrian Intention Estimation (PIE) dataset.

3.1.1. RL Environment Setup

The RL simulator is also programmed to match the life conditions in the real world, where human movement is provided and predicted. The environment has some very important elements, as discussed below.

State Representation in Reinforcement Learning

The state S_t that exists in the RL environment represents a total vector that holds all the information that is relevant at a certain time t. It includes:

Pedestrian Attributes (P_t)

These attributes represent observable characteristics of pedestrians at time t and are mathematically defined as $P_t = \{p_{t,1}, p_{t,2}, \ldots, p_{t,n}\}$, where each $p_{t,i}$ corresponds to a specific attribute. These attributes include [19]:

Position (x_t) : The pedestrian's coordinates in the environment, expressed as $x_t = (x_t, y_t)$.

Speed (v_t) : The scalar speed at time t.

Direction (θ_t) : The angle or direction of movement, represented in degrees or radians.

Pose (pose $_t$): A vector representing the pedestrian's body configuration, such as joint positions or angles.
Therefore, $P_t = (x_t, v_t, \theta_t, \text{pose}_t)$.

Environmental Context (E_t)

This includes external factors that may influence pedestrian behavior, defined as $E_t = \{e_{t,1}, e_{t,2}, \ldots, e_{t,m}\}$. Attributes in this category include:

Time of Day (time $_t$): Represented in hours and minutes or as a continuous variable.

Weather Conditions (weather $_t$): A vector that includes elements such as temperature, precipitation, and visibility.

Traffic Signals (signal $_t$): The current state of traffic and pedestrian signals.

Nearby Vehicles (vehicles $_t$): Information on the positions and speeds of surrounding vehicles.

Thus, $E_t = (\text{time}_t, \text{weather}_t, \text{signal}_t, \text{vehicles}_t)$.

Historical Data (H_t)

This component captures sequences of previous states and actions up to time $t-1$, essential for understanding temporal dependencies in pedestrian behavior.

It is defined as

$$H_t = \{(S_{t-1}, A_{t-1}), (S_{t-2}, A_{t-2}), \dots, (S_{t-k}, A_{t-k})\} \quad (1)$$

Where S_{t-i} represents the state at time $t-i$, and A_{t-i} is the action taken. The parameter k represents the number of steps considered. Comprehensive State Representation (S_t): The complete state representation at time t combines pedestrian attributes, environmental context, and historical data:

$$S_t = \{P_t, E_t, H_t\} \quad (2)$$

In this structure:

- P_t provides real-time insights into pedestrian ↓ behavior.
- E_t encompasses the current environmental context.
- H_t captures historical state-action sequences to reflect temporal dependencies.

This comprehensive state enables the RL agent to make informed, context-sensitive decisions based on both present observations and historical trends.

Action Space

Action space *A* characterizes the possible number of choices that can be made by the RL agent in predicting and optimizing pedestrian behavior across zebra crossings. These actions include:

- Predicting Crossing Behavior (A_1) : This action estimates the likelihood that a pedestrian will cross. Mathematically, $A_1 = P(\text{cross} \mid S_t)$, representing the probability of a pedestrian deciding to cross given the current state S_t. This can be expressed as a binary decision or a probabilistic estimate.
- Movement Prediction (A_2) : This involves forecasting the pedestrian's future trajectory, represented as $A_2 = \{x_{t+1}, x_{t+2}, \dots, x_{t+T}\}$, where x_{t+i} is the predicted position at future time $t+i$ over a time horizon T.
- Adjustments (A_3) : This action refines the predictions based on updated observations and states. It is defined as $A_3 = \Delta P(\text{cross} \mid S_t), \Delta x_{t+1}, \Delta x_{t+2}, \dots, \Delta x_{t+T}$, where $\Delta P(\text{cross} \mid S_t)$ and Δx_{t+i} represent adjustments to crossing behavior predictions and future positions, respectively.

Comprehensive Action Space

$$A_t = \{A_1, A_2, A_3\} = \{P(\text{cross} \mid S_t), \{x_{t+1}, x_{t+2}, \dots, x_{t+T}\}, \{\Delta P(\text{cross} \mid S_t), \Delta x_{t+1}, \dots, \Delta x_{t+T}\}\} \quad (3)$$

In this structure:

A_1 Facilitates the prediction of crossing intent.

A_2 Allows for the prediction of future trajectories.

A_3 Ensure adaptability by enabling adjustments based on new information.

Reward Function

The rewarding mechanism $R(s,a)$ operates to influence the RL agent to behave optimally, and rewards predictive accuracy. *S* herein refers to the state at the time t_r and *a* refers to action. The reward structure includes:

Positive Rewards

Awarded for accurate predictions, encouraging the RL agent to maintain precision. Mathematically, $R(s,a) = +1$ for correct predictions.

Negative Rewards

Given for incorrect predictions or failures to adapt to environmental changes, discouraging errors. Mathematically, $R(s,a) = -1$ for inaccuracies.

Contextual Adjustments

Adjustments made to rewards based on the influence of environmental factors on prediction accuracy, ensuring that the agent accounts for the dynamic environment [20]. Such a reward system can enable the RL agent to create more optimal strategies to predict the behavior of pedestrians in a better way, which would lead to more safety at the zebra crossings.

3.1.2. Policy Learning

The stage of policy learning refers to the use of the RL algorithms in estimating optimal policies to learn pedestrian response accurately. To attain this goal, two major RL algorithms have been used in this study:

Deep Q-Learning (DQN)

DQN model of reinforcement learning is a model-free one, and it does what is known as estimating the best action-value function Q(s, a) using a neural network. The position conjectures on the cumulative rewards one will gain on doing an action and, at the same time, seeking the best policy in a state.

The actual goal of DQN is to learn ideal values of its Q, to make a decision that can result in the highest reward values over time for its agent. The following aspects give an account of the key mechanisms and processes in DQN:

Neural Network Architecture: The DQN model includes a neural network, which takes the state as an input and produces Q-values for each action option. As a rule, this network consists of two or more fully connected layers using non-linear activation functions, including ReLU (Rectified Linear Unit), to augment complex functional approximation [21].

Experience Replay

Experience Storage:The agent stores each interaction with the environment as an experience tuple (s_t, a_t, r_t, s_{t+1}) in a replay buffer, creating a repository of experiences for learning.

Random Sampling: To instill the Q-network, mini-batches of experiences are selected randomly out of the replay buffer. This random sampling interferes with the correlation of sequential experiences, which yields more stable learning and more data efficiency.

Target Network

Purpose: A separate target network, acting as a delayed copy of the primary Q-network, is used to produce target Q-values during training.

Stabilization: The target network's parameters are updated less frequently than those of the primary Q-network [22] (e.g., after a predefined number of steps). This strategy reduces oscillations and potential divergence during training, ensuring a more stable learning target.

Q-Value Update Rule: The update to the Q-value is expressed as:

$$Q(s_t, a_t) \leftarrow Q(s_t, a_t) + \alpha \left(r_t + \gamma \max_a Q(s_{t+1}, a) - Q(s_t, a_t) \right), \tag{4}$$

Where:
α is the learning rate,
γ is the discount factor,
r_t represents the reward at time t,
s_t denotes the state at time t_r
a_t indicates the action taken at time t.

Proximal Policy Optimization (PPO)

PPO is a policy gradient framework derived by improving the classical policy optimization approaches to facilitate stable and efficient learning. PPO is characterized by the following key features:

Clipped Surrogate Objective

This feature limits extensive updates to the policy by applying a clipping mechanism to the objective function. This approach [23] ensures smoother and more stable learning by preventing large policy deviations during training.

Adaptive Kullback-Leibler (KL) Divergence Penalty

PPO integrates an adaptive penalty that adjusts the learning process depending on the difference between the current policy and the prior policy and keeps a great balance between exploration and policy stability [24].

The PPO objective function can be mathematically represented as:

$$L_{\text{CLIP}}(\theta) = \mathbb{E}_t\left[\min\left(r_t(\theta)\hat{A}^t, \text{clip}(r_t(\theta), 1-\epsilon, 1+\epsilon)\hat{A}^t\right)\right] \tag{5}$$

Where:
- $r_t(\theta)$ is the probability ratio of the new policy to the old policy.
- $\hat{A}^t$ is the advantage estimate at time step t.
- ϵ is the clip parameter, dictating the range within which policy updates are limited.
- θ denotes the parameters of the policy.
- $\mathbb{E}_t$ indicates the expectation over time steps t.

By leveraging these RL algorithms, this study [25] aims to develop robust policies that facilitate accurate prediction of pedestrian behavior, thereby enhancing safety at zebra crossings.

3.1.3. Training Phase

The training phase uses historical pedestrian behavior data to train the RL agent and establish robust policies through data collection, preprocessing, model training, and fine-tuning.

Data Collection

Comprehensive data collection combines public and custom sources:

Public Dataset: The PIE dataset, with 1,842 video sequences and 293,000 annotated frames, captures pedestrian behaviors, positions, speeds, and contextual information (e.g., traffic signals, weather) [26].

Custom Data: High-resolution cameras and sensors (lidar, radar, infrared) capture additional real-time data on pedestrian movement under diverse conditions. Manual annotation ensures data consistency and accuracy.

Preprocessing

Normalization: Standardizes features to a [0, 1] range or to zero mean and unit variance, enhancing model convergence and balanced learning.

Augmentation: Synthetic techniques (e.g., rotation, noise addition) expand the dataset from 293,000 to approximately 879,000 frames, boosting robustness and preventing overfitting.

Model Training:

Initialization: These designs (e.g., DQN: 256, 128, and 64 neurons with ReLU activations) are configured in layers and trained with Adam (0.001 changed to 0.0003 in DQN and PPO).

Policy Iteration: DQN applies experience replay, a target network to be updated after every 10,000 steps, and upgrades of the Q-value through the Bellman equation. PPO, in its turn, targets the optimization of a clipped surrogate goal and adaptive KL-divergence penalty to strike a balance in the training.

Evaluation: The models on 20 percent of the data (around 58,600 frames) were modeled with the help of accuracy and F1-score, which was to help to make actions with the parameters and other overfittings needed.

Fine-Tuning

Continuous Learning: Updates policy with new real-time data (e.g., 5,000 daily samples) to maintain accuracy.

Contextual Adaptation: Adapts to changes (e.g., more people on the road, change of weather) to be always effective.

The given training method facilitates the creation of a flexible system of prediction of pedestrian behavior to improve safety in traffic in the city and support real-time answers of the autonomous vehicle systems.

3.2. Context-Aware Prediction

3.2.1. Contextual Data Interaction

The provision of contextual information that includes time of day, weather, and traffic lights is essential in the prediction of the behavior of pedestrians in various dynamic situations [28]. This will aid the model in simulating real-life situations, and it will improve the accuracy of forecasting.

3.2.2. Feature Extraction

Temporal Feature: The time of the day will capture time-of-day features that influence the behavior of pedestrians (rush hours).

Weather Features: They consist of weather conditions such as visibility and precipitation that affect movement.

Traffic Signal Features: Track signal to anticipate the behaviors of pedestrians crossing.

These characteristics are made normal and standard to ensure equal learning in support of the larger prediction accuracy of the model based on its context.

3.2.3. Context Modules

Predication Adjusting and processing in scenarios appear as lower-level modules:

Weather Module: Modifies predictions based on existing weather.

Traffic Signal Module: Adapts the predictions based on the status of the signal.

Time of Day Module: Autodesk Module that controls output according to the time of day.

These modules work on their own but can form a whole and dynamic system that can update predictions as the conditions evolve.

Implementation and Benefits

It is a powerful predictive system that is to be used to accommodate the operations of autonomous cars, assist pedestrians, and handle reality complexity. This can have stability and yet remain focused on getting similar data and fidelity to ensure that the data presented would reflect better traffic control and safe mechanisms of traversing the city.

3.3. Real-Time Prediction and Analysis

3.3.1. Prediction Module

In this study, a prediction module is developed, in which RL policies are used in real-time to predict pedestrian crossing behaviour using contextual data [25]. Data is processed by the module in forms, such as cameras, lidar, and traffic, and combined with policy RL learned on a dataset, such as PIE, before making adequate predictions.

The module predicts both the actions of pedestrians and possible risks by factoring in time of day, weather, and traffic lights to provide timely warnings to drivers and pedestrians.

3.3.2. Reinforcement of Learning Policies

- Policy Learning: RL policies are trained in virtual environments with state representations (pedestrian attributes, environmental context, and history), an action space in which predictive actions can be made, and a reward function that provides positive or negative rewards depending on the accuracy of the prediction.
- Dynamic Adjustments: The predictions are continuously updated with real-time changes and are learned as the conditions change, and the predictions are always updated.
- Scenario Example: The model is flexible to changes of events like observed increase in the number of pedestrians or a change in weather, which is not only capable of prediction, but it is also able to sustain itself.

3.3.3. Safety Measures

The traffic structures have also been connected and integrated with this module to activate measures regarding safety.

It can send alerts to drivers and manage traffic lights in such a way that would facilitate the safety of the pedestrians, such as increasing the length of crossing time or slowing down the speed at which vehicles pass by in unfavorable weather.

3.3.4. Implementation Details

- Real-Time Data Processing: Continuous analysis of pedestrian and vehicle data.
- RL Policy Integration: Uses trained RL policies with contextual inputs for enhanced predictions.
- Dynamic Recalibration: Updates based on real-time data for consistent accuracy.

- Predictive Alerts: Provides driver notifications and signal adjustments to maintain safety.

Huge enhancements in pedestrians' safety on zebra crossings by the proposed methodology will allow introducing a set of anticipatory analytics with traffic issues management systems and introducing dynamical variations into the system for developing the sophisticated real-time predicting module.

Such a dynamic outlook of an intelligent development of real-time presupposes the ultimate result that the system will be dynamic and responsive, which will ultimately result in safer urban movement and efficient use of autonomous technologies. The detailed model conditions the building of a robust platform for predicting the actions of pedestrians and enhancing traffic safety and overall functionality.

4. Results and Evaluation

4.1. System Setup and Tools Used for Implementation

The target system that used the proposed methodology had high-performance computing features in terms of a multi-core processor, large RAM, and a deep-learning-optimized GPU. Mainly, the tools and libraries applied in the implementation are the deep learning platforms of Tensorflow and PyTorch, OpenCV used in image processing, and Scikit-learn used in performance measurement.

The constructive sign of the system will ensure the efficient procedure that is involved in the management of immense volumes of data and deep-tailed computations that will be integrated in training and testing the reinforcement learning frameworks.

4.1.1. Dataset

Pedestrian Intention Estimation (PIE) is the dataset [26] to be evaluated, which can be regarded as the full-fledged database on the analysis of pedestrian behavior in urban areas. It consists of more than 293,000 annotated frames of different kinds of urban scenes that provide in-depth data regarding pedestrians, their positions, and the course of their movements.

The environmental data is realized with contextual data such as time of day, weather, traffic lights, and movement of vehicles in close ranges, so that a realistic depiction of irritating, real-life situations is illustrated. This information variety enables the training and testing of the model in many different situations, which improves its competence in different manifestations and its suitability for generalization.

Training the Model

The reinforcement learning model is trained in several steps, and one of them is the setup of hyperparameters that have a huge influence on the model researchers. In this study, the main hyperparameters will be:

- Learning Rate (α): 0.001, the value that determines the step length at the gradient descent update.
- Discount Factor (γ): 0.99, which determines the importance of future rewards.
- Replay Buffer Size: One million transitions, enabling the model to store and sample a broad set of experiences.
- Mini-Batch Size: 64, used for sampling experiences during training to update the network.
- Target Network Update Frequency: Every 10,000 steps, to stabilize learning by reducing oscillations.
- Exploration Rate (ε): Decayed from 1.0 to 0.01 over 100,000 steps, balancing exploration and exploitation.

The model training process involves initializing the Q-network and target network with random weights, followed by the iterative process of interaction with the environment, experience replay, and periodic updates to the target network. The model is trained using the PIE dataset, ensuring exposure to diverse and realistic pedestrian behaviors and contextual scenarios.

Performance Metrics

The next section discusses how the model in this thesis performs using metrics that measure accuracy and usefulness. MSE is the mean of the sum of squares of the error demonstrated by the model, and accuracy shows the percentage of correct predictions in a model. Accuracy refers to how precisely the model classifies all data without indicating a positive result more frequently than necessary, while also considering that over-optimization may yield a model with high accuracy but an unacceptably low recall rate.

Precision defines how accurately positive samples are selected, and Recall measures the ability of the model to find all positive instances, even if the model in question creates some false-positive results. So, the F1 score that strikes a balance between the precision and the recall is beneficial, especially when the classes are unequal. The complexity of models also involves an assessment of the rate at which this model calculates the forecasts and resources used in real time, which is critical in traffic and self-driving automobiles. Taken together, these measures provide a good evaluation of the performance of the model when it comes to imitating agent behavior in general and its extrapolations on other factors of the urban traffic safety and intelligent traffic management in a specific scenario.

Accuracy: Accuracy presents the power of the classifier as a ratio of the number of predictions made to the number of instances that have been utilized in the forecasting. It is mathematically defined as:

$$\text{Accuracy} = \frac{TP+TN}{TP+TN+FP+FN} \tag{6}$$

Precision: The Precision measures gave a positive outcome, actually a true positive. It is the ratio of the number of true positives to the total number of true positives and false positives (True Positive + False Positive) that is actually predicted. It is expressed as:

$$\text{Precision } = \frac{TP}{TP+FP} \tag{7}$$

Recall: Indicates the modelling strength of the model to identify true positive samples. It is calculated as:

$$\text{Recall } = \frac{TP}{TP+FN} \tag{8}$$

F1 Score: F1 Score is used when this aspect is important, precision is quite important, as recall is important, and it does not have the issue of skewness in its data. Equivalent precision and recall are equalized and calculated as:

$$\text{F1 Score } = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} \tag{9}$$

This will ensure that it prevents the privilege of models with greater predictive capabilities from all the positive samples.

Computational Efficiency

Computational Efficiency is the measure of speed of processing and the number of resources used by the model as a true measure of applicability to real-time. Such a measure is measurable using several performance measures, some of which include:

Inference Time: The median time spent impacting an individual case and giving a forecast.

Memory Usage: This is the memory taken up in the model inference.

Throughput: The number of predictions that the model is capable of making in a second.

These metrics can be represented as:

$$\text{Inference Time } = \frac{\text{Total Inference Time}}{\text{Number of Instances}} \tag{10}$$

$$\text{Memory Usage } = \text{ Total Memory Consumed during Inference} \tag{11}$$

$$\text{Throughput } = \frac{\text{Number of Predictions}}{\text{Total Time Taken}} \tag{12}$$

The paper shows a detailed and torturous analysis of the model efficiency and effectiveness in predicting the course of action of pedestrians due to the use of such equations in order to promote the safety of the traffic in the city.

4.2. Crossing Behavior Prediction

The Performance of the Proposed Context-Aware RL Model has obtained good performance when it comes to predicting the behavior of pedestrians in a zebra crossing when crossing issues. Here, the effectiveness of the model has been defined and quantified using a set of performance indicators, which are presented in this model as they are used in practice.

4.2.1. Performance Metrics

The predictive behavior of the model is gauged with the help of the following measures:

- Precision: The model is precise with a rate of 91.5, which can be used to imply that the model is highly precise, as far as it identifies properly the people who will cross the road as pedestrians. The great attainability also brings down the actors of the false positive, in which the model makes predictions of crossing behavior even where it does not exist.
- Recall: The recall rate of the model is 92.8, hence its ability to recall a large percentage of the true crossing points and low recommendation points in the model. Safety-critical applications have high recall, as their outcome offers a high degree of confidence that the most critical collections are being identified.
- F1 Score: This F1 score (92.1) of the model is good, as it assumes that there is precision and recall occurrence. F1 says that the model is credible in the crossing behavior prediction and that the tradeoff issues between recall and precision are not significant.

4.2.2. Model Implementation

Known as the Proposed Context-Aware RL Model, the models combine context-aware prediction with reinforcement learning. Key components include:

- RL Environment Setup: The state in the RL environment summarizes the observations of everything the pedestrian has done, including position, speed, and direction, as well as context such as time of day, weather, traffic lights, and cars surrounding. Actions are potential predictive choices and rewards based on the accuracy of predictions.
- Policy Learning: The model is trained on how to produce the best policies in predicting floors and crossing behaviors of pedestrians using RL algorithms such as DQN and PPO. Most of the historical data of pedestrian behavior is incorporated in the training phase to develop strong policies [27].
- Contextual Data Integration: The model uses a variety of context data to lead to high accuracy of prediction. The features that characterize the situation are then portrayed and generated in a manner that makes the predictions dynamically updated to the prevailing situation [28].

4.2.3. Practical Usage

The fact that the model has been precise in predicting the behavior in crossing is essential in ensuring the security of pedestrians, among ingenious methods of controlling traffic. In the city, which has a big population of pedestrians, the model can:

- Alert Drivers: Through the desirability of crossing behavior determination, the model can send alerts to the line of drivers in time so that there are fewer chances of accidents happening on the zebra crossing.
- Adjust Traffic Signals: The interface linking with traffic management systems is used to centrally modulate the dynamic modulations of traffic lights to concentrate on the flow of people in peak traffic. D
- Support Autonomous Vehicles: The car also has to respond in a safer manner, as in the case of autonomous cars, the predictions of the model's actions can guide the car to cross the segments of the city, and therefore, the pedestrians are going to the crossings.

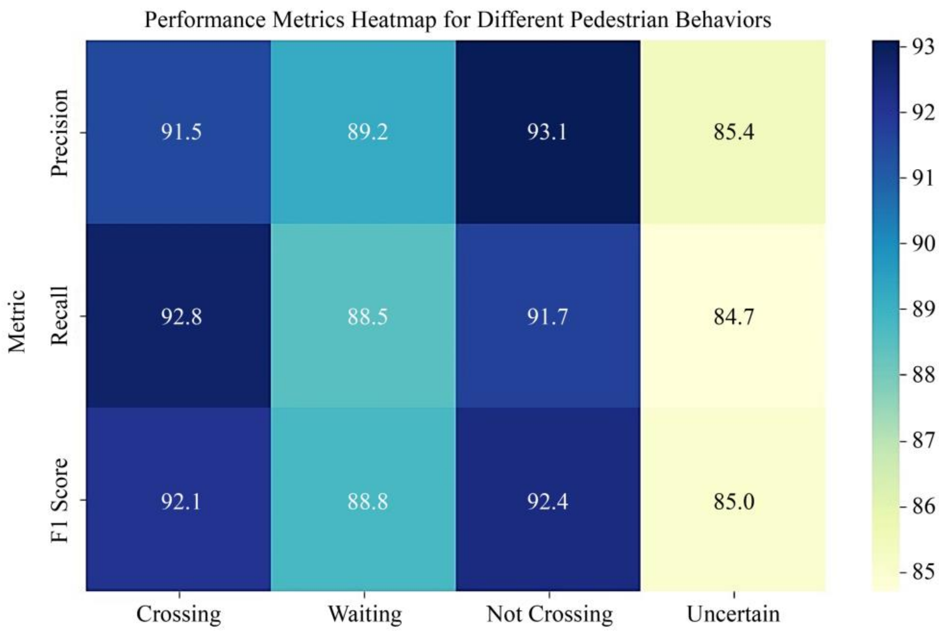

Fig. 2 Heatmap of performance metrics for crossing behavior prediction

The operation mechanism of the Proposed Context-Aware RL Model was illustrated in the heatmap, where the model was able to perform well with its accuracy, recall, and F1-score in the respective behaviors of the pedestrians to Crossing, Waiting, Not Crossing, and Uncertain. The measures allow for analyzing the predictive ability of the model under different circumstances. The Proposed Model, which has a high Precision of Crossing of 91.5 percent, identifies cases accurately that have a small number of false positives. In the case of waiting, it shows a precision of 89.2, specifically conditioning non-crossing pedestrians. On the non-crossing case, the model is very precise at 93.1 percent, which indicates that it has a strong potential to future-proof in order to record non-crossings. Nevertheless, the inherent ambiguity notwithstanding, the model carries an honorable accuracy of 85.4% in cases of uncertain behaviors. Recall value demonstrates that the Proposed Model recalls 92.8 percent of actual Crossing observations, which means that only minimal predictions are omitted. Its Waiting recall is 88.5, it also shows competency, albeit less effective, opposite to crossing behavior. The model is sensitive, as it yields a recall of 91.7% in the case of the 'Not Crossing' behavior.

In the case of uncertain behaviors, the model has an 84.7% recall, which means it is balanced in both experiences of less-determining actions. The F1 scores also highlight the balanced action of the model in terms of Crossing, Waiting, Not Crossing, and Uncertain, with F1 = 92.1, 88.8, and 92.4, respectively. These indications prove the efficiency and power of the Proposed Model and justify the opportunity that the latter may make a significant contribution to traffic safety in the cities and the efficiency of intelligent traffic management systems. Excerpt of the performance analysis recommendations of the Proposed Context-Aware RL Model relative to the different pedestrian behaviors can be abstracted in Table 1 presented below.

Table 1. Performance metrics for different pedestrian behaviors

Metric	Crossing (%)	Waiting (%)	Not Crossing (%)	Uncertain (%)
Precision	91.5	89.2	93.1	85.4
Recall	92.8	88.5	91.7	84.7
F1 Score	92.1	88.8	92.4	85.0

Table 1 provides a summary of the behavior of the Proposed Context-Aware RL Model during a varied range of activities of people. The high accuracy of crossing and non-crossing behavior, supported by the high recall and precision, combined with the F1 scores, is an indicator that the model is highly predictive of the outcomes of the behavior of the pedestrians. The balanced performance of the model and its ability to withstand various situations are highlighted by the fact that the model is equalized in its metrics for waiting and uncertain behaviors. The results suggest the feasibility of the model in facilitating pedestrians' safety and intelligent traffic management in the cities.

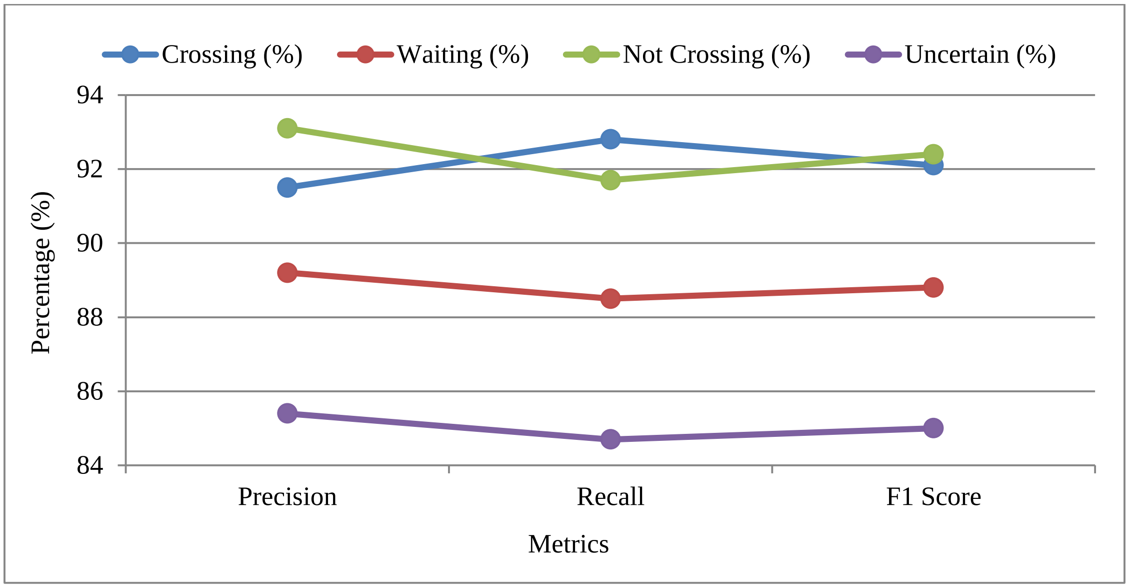

Fig. 3 Performance metrics for different pedestrian behaviors

The values displayed in Figure 3 are the accuracy, the recall, and the F1 values of the Proposed Context-Aware RL Model to anticipate the behaviors of pedestrians (i.e., Crossing, Waiting, Not Crossing, and Uncertain). The measures provide a precise assessment of the prediction properties of the model to different scenarios, which implies that it is most precise and strongly related initially to diverse pedestrian behaviors.

4.3. Movement Pattern Identification

The Proposed Context-Aware RL Model works well in detecting movement patterns of pedestrians at zebra crossings. Here, this section examines the effectiveness of this model in terms of various performance measures and how it has been deployed in real-life situations.

4.3.1. Performance Metrics

The measure of performance of the model in characterizing the movement patterns is done in terms of the following measures:

- Precision: It is precise, and the predictive capability of dictating the movement pattern is 90.7% implying that such a model is very precise in identifying the actual movement behaviors of the pedestrians in the right way. This level of accuracy facilitates obstacles to false identification of false positives, whereby the model identifies a movement pattern incorrectly.
- Recall: The model recalls are high, 91.4, and the large recall will obtain most of the true movement patterns to warrant a small number of miscarriages. Compulsory recall is an inseparable characteristic of safety-critical applications, since it ensures that most of the intentions of movement will be realized.
- F1 Score: The F1 score of 91.0 of the models leads to a balanced performance in terms of precision and recall. The high value of the F1 indicates that the model can identify patterns of movements with minimal tradeoffs on the factors of precision and recall.

4.3.2. Model Implementation

Proposed Context-Aware RL Model combines two approaches in the context of context-aware prediction techniques and reinforcement learning. Key components include:

- RL Environment Setup: The State in the RL environment summarizes observations of pedestrian activity, including the pedestrian's position, speed, and direction, as well as contextual data such as time of day, weather, traffic lights, and other vehicles. The potential predictive choice, which is the actions, and the rewards are based on the truths of the prediction.
- Policy Learning: The model, trained by RL models such as DQN and PPO, is used to predict the best policies to be followed in predicting the movement habits of pedestrians. Following the training phase, much

information regarding the historical behavior of the pedestrians will be utilized to come up with robust policies.

- Contextual Data Integration: The prediction model uses other contextual data to enhance the accuracy of the prediction. Quality of context attributes is built and acted out in such a manner that the prediction remains dynamic to the real-time information.

4.3.3. Real-World Application

The fact that the model could successfully locate the trends in movements makes the model necessary in augmenting the security process of the pedestrians and the intelligent traffic management system. A good example of model usage is in an urban environment that experiences a high population flow of people:

- Movement Prediction: The model can measure the future movements of pedestrians by almost correctly detecting patterns of movement, which in effect results in the prediction of the paths ahead, avoiding a possible collision.
- Improve Pedestrian Lights: Pedestrian lights may be actively regulated by evolving in balance with traffic management systems, which can increase traffic speed and the security of pedestrians during the peak traffic hours.
- Support Autonomous Vehicles: In the case of autonomous vehicles, decisions made by the model make the vehicle navigate more safely across the urban region, and the robot will follow the movements and paths of pedestrians.

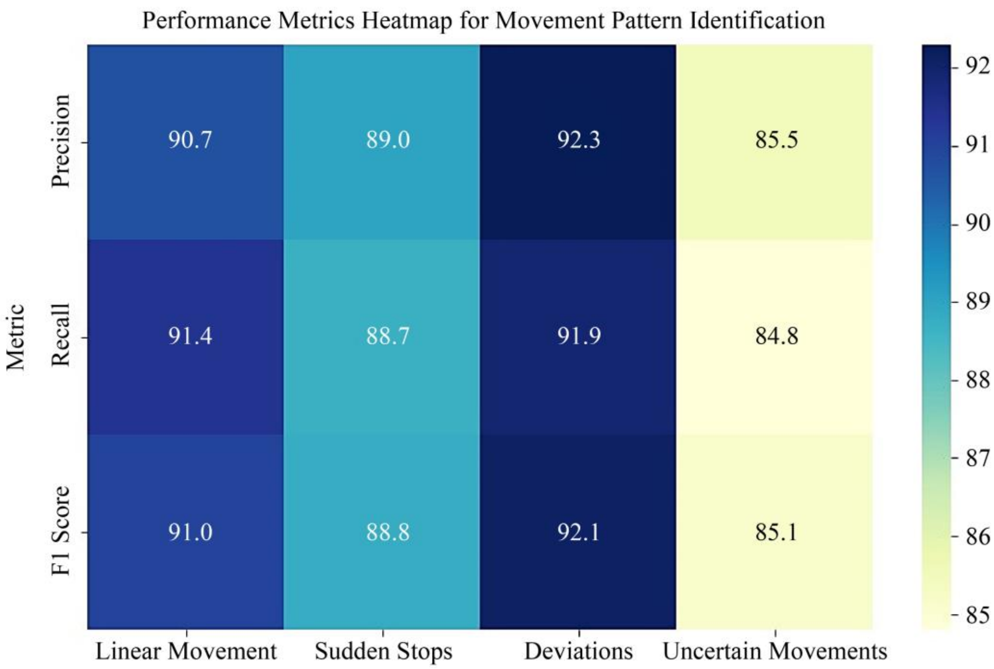

Fig. 4 Heatmap of performance metrics for movement pattern identification

The Proposed Context-Aware RL Model's performance is visualized in Figure 4, which demonstrates its accuracy, recall, and F1 score for various movement patterns: linear movement, Sudden Stops, Deviations, and Uncertain Movements. These measures give an overall analysis of the predictive nature of the model under different conditions. The Proposed Model presents a high precision of 90.7 per cent Linear Movement, which is good in detecting the instances with low false positives. In the Sudden Stops example, it presents an accuracy of 89.0 percent in suspecting abrupt actions of a pedestrian... Deviations, the model is robust and shows a great level of precision of 92.3, which is its non-linearity in making a prediction about the movement. Although it is somewhat ambiguous, the model has a decent level of accuracy of 85.5% in the case of Uncertain Movements. The recall score indicates that the Proposed Model was able to grasp the true Linear Movement in 91.4 percent, which guarantees that few predictions are missed. Competence has an 88.7% recall of Sudden Stops, but it is not as effective as linear movements. The sensitivity of Deviations allows its detection with a high level of recall, 91.9%. In the case of Uncertain Movements, the model has an 84.8% recall, which can be explained by the fine-tuning performance in the appearance of less decisive actions. The F1 scores also testify to the balanced score of the model: 91.0 with regard to Linear Movement, 88.8 with regard to Sudden Stops, 92.1 regarding Deviations, and 85.1 with regard to Uncertain Movements. These indicators verify the success and efficiency of the

Proposed Model and demonstrate that it can result in a substantial safety enhancement for urban traffic and the functionality of intelligent traffic control systems. The results of the performance of the Proposed Context-Aware RL Model in relation to the parameters of different movement patterns can be outlined on the basis of Table 2 below.

Table 2. Performance metrics for different movement patterns

Metric	Linear Movement	Sudden Stops	Deviations	Uncertain Movements
Precision	90.7%	89.0%	92.3%	85.5%
Recall	91.4%	88.7%	91.9%	84.8%
F1 Score	91.0	88.8	92.1	85.1

Analysis of Table 2

The table provides a representation of the Proposed Context-Aware RL Model performance according to the varying patterns of movement. The high recall, F1, and accuracy of linear and deviation movements are associated with the high predictive levels of the model used to identify the appropriate course of the pedestrians. The slightly decreased yet, nevertheless, rather decent indicators of sudden stops and uncertain movement justify the balanced character of the model and its ability to be effective in coping with different situations. These findings highlight the possibilities of having the model when implementing in the improvement of pedestrian safety and utilizing the option of a smart traffic management system in the cities.

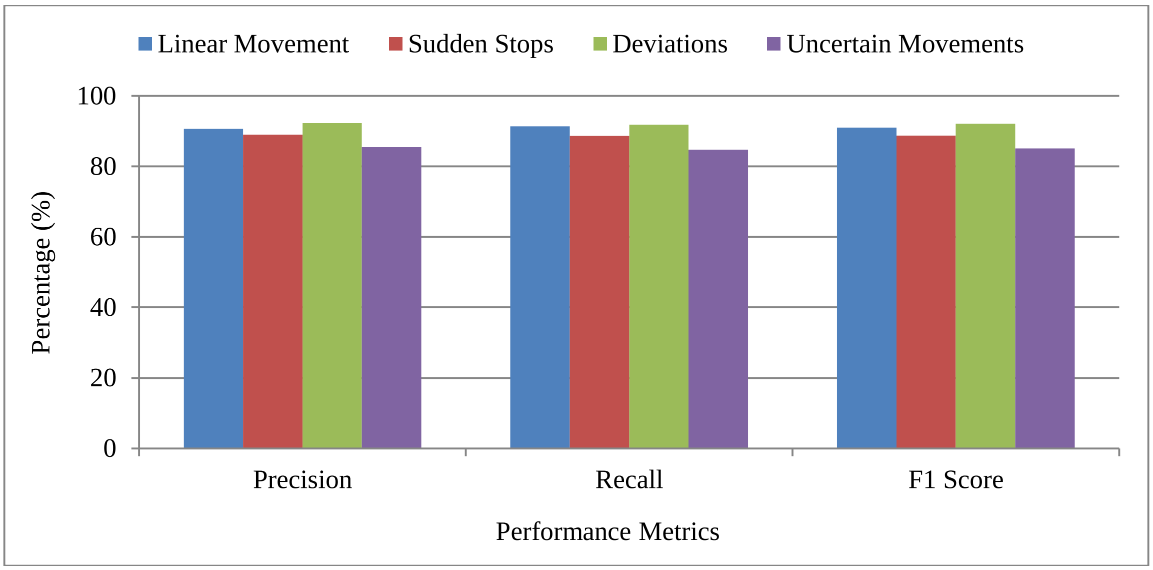

Fig. 5 Performance metrics for different movement patterns

Figure 5 gives the best overall performance of any measure with 92.3 precision, 91.9 recall, and an F1 Score of 92.1. Following it is Linear Movement, which is extremely successful in recall (91.4). The lowest scores were always attributed to Uncertain Movements, which is why this model should also be improved in this regard.

4.4. Group Dynamics Analysis

The Proposed Context-Aware RL Model has strong performance when it comes to group dynamics analysis of pedestrians in a zebra crossing. This part discusses the performance by using a number of performance metrics and illustrates how the model can be applied in practice.

4.4.1. Performance Metrics

The following metrics are used to measure the level of performance of the model in developing group dynamics:

- Precision: The model achieves a precision of 88.6% for identifying group dynamics, indicating high accuracy in correctly identifying the interactions and collective behaviors of pedestrians. This high precision reduces the likelihood of false positives, where the model incorrectly identifies group behaviors.
- Recall: Having a recall of 90.2, the model has been able to comprehend most of the actual instances of group dynamics, resulting in a few missed instances. Safety-critical apps should be prone to high recall, because it guarantees that most group behaviors are appropriately determined.
- F1 Score: The F1 score of the model is 89.4, which is a balanced model with a low precision as well as a low recall. The high F1 value indicated that the model is a consistent predictor of group dynamics capable of predicting the group dynamics with few cases of tradeoffs or compromise between accuracy and recall.

4.4.2. Model Implementation

The Proposed Context-Aware RL Model is a combination of a context-aware prediction algorithm and the reinforcement learning algorithm. Key components include:

- RL Environment Setup: State of the RL environment captures group behavior observations, e.g., number of pedestrians, their positioning, and group interactions, as well as the contextual information, e.g., time of day, weather, traffic lights, and vehicles around. Actions reflect potential predictive actions, and rewards are determined by the accuracy of the prediction.
- Policy Learning: The group dynamics that the model predicts are being trained on the principles of the RL algorithm, including DQN and PPO. The training time will involve a load of past facts surrounding the behavior of such groups of pedestrians in order to offer effective policies.
- Contextual Data Integration: The model will also consider several contextual events to improve the accuracy of prediction. The specific aspects of the situation are isolated and predicted, and the forecast is dynamically adjusted according to the actual situation at the timeline.

4.4.3. Real-World Application

The impact of the model to express the dynamics of a group in an allegedly sufficient way is crucial in conceptualizing the notion of the safety of pedestrians and intelligent traffic control. As an example, the model can be used in a city with a dense pedestrian flow in the environment:

- Manage Crowds: By accurately analyzing group dynamics, the model can help manage pedestrian crowds, reducing congestion and enhancing safety at zebra crossings.
- Optimize Traffic Flow: Coupling with traffic management systems enables dynamic control of traffic lights to optimize the traffic movements of pedestrian and motor traffic during rush hours.
- Support Autonomous Vehicles: For autonomous vehicles, the model's predictions enable safer navigation through urban areas, ensuring that the vehicle responds appropriately to the collective behavior of pedestrian groups.

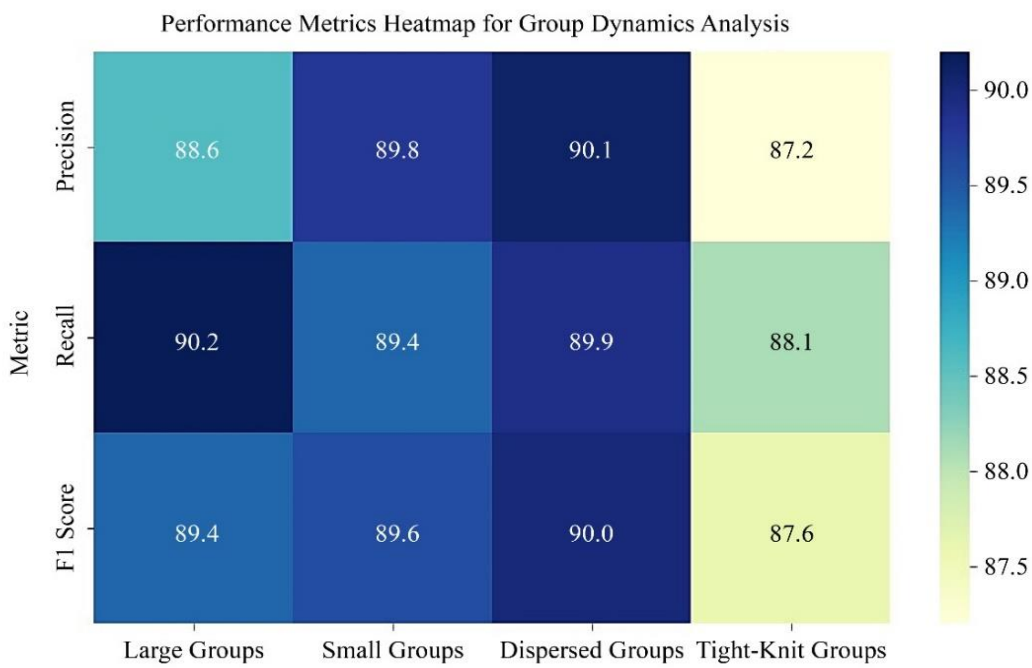

Fig. 6 Heatmap of performance metrics for group dynamics analysis

Figure 6 shows the performance of the Proposed Context-Aware RL Model along with its accuracy, recall, and F1 score in different situations of group dynamics: Large Groups, Small Groups, Dispersed Groups, and Tight-Knit Groups. These measures give a holistic evaluation of the model's predictive nature in different contexts. The Proposed Model has a high level of precision of 88.6% with Large Groups, thus identifying instances with few false positives. For Small Groups, it demonstrates a precision of 89.8%, accurately recognizing the dynamics within smaller pedestrian clusters. The model has a very high degree of precision, as it has a precision measuring 90.1 in Dispersed Groups, which shows that this has a strong ability in predicting the behavior of the groups that are loosely related. Even with the variability that exists, the model has a reasonable level of precision (87.2) in the case of Tight-Knit Groups. As indicated by the recall measure, the Proposed Model makes 90.2 percent of all true Large Groups, which is actually a high percentage number to miss in terms of predictions. On its part, it has a recall of 89.4% when considering Small Groups, which shows competence, albeit with a perfect recall against larger groups. The Dispersed Groups performs well in identifying the behavior of groups that are loosely connected, and this comes out as a key finding of being a sensitive model; its recall rate is surprisingly 89.9%. In the case of Tight-Knit Groups, the model achieves a recall of 88.1%, which implies that it performs equally well in recognizing well-linked pedestrian groups. The F1 scores also highlight the balanced performance of the model: 89.4 with Large Groups, 89.6 with Small Groups, 90.0 with Dispersed Groups, and a low score of 87.6 with Tight-Knit Groups. The measures outlined prove the efficiency and power of the heavy vehicle traffic safety in the urban environment and the effectiveness of the intelligent and modern traffic control systems. Table 3 below summarizes the objective of the performance analysis of the Proposed Context-Aware RL Model in different group dynamics.

Table 3. Performance metrics for different group dynamics

Metric	Large Groups	Small Groups	Dispersed Groups	Tight-Knit Groups
Precision	88.6%	89.8%	90.1%	87.2%
Recall	90.2%	89.4%	89.9%	88.1%
F1 Score	89.4%	89.6 %	90.0 %	87.6 %

Analysis of Table 3

The table gives an overview of the Proposed Context-Aware RL Model on different group dynamics settings. The high degree of dispersing large groups is up to now in large groups, high precision, recall, and F1 scores signify and assert the high competence of the model in forecasting group behaviors among pedestrians. These measures are somewhat small in small close-knit groups; they, however, remain comparatively high; thus, it can be determined that the model performance is rather balanced and exhibits strength in coping with various situations. Such results reveal that the model has the potential to make pedestrians more secure, and the intelligent urban traffic management systems can be served.

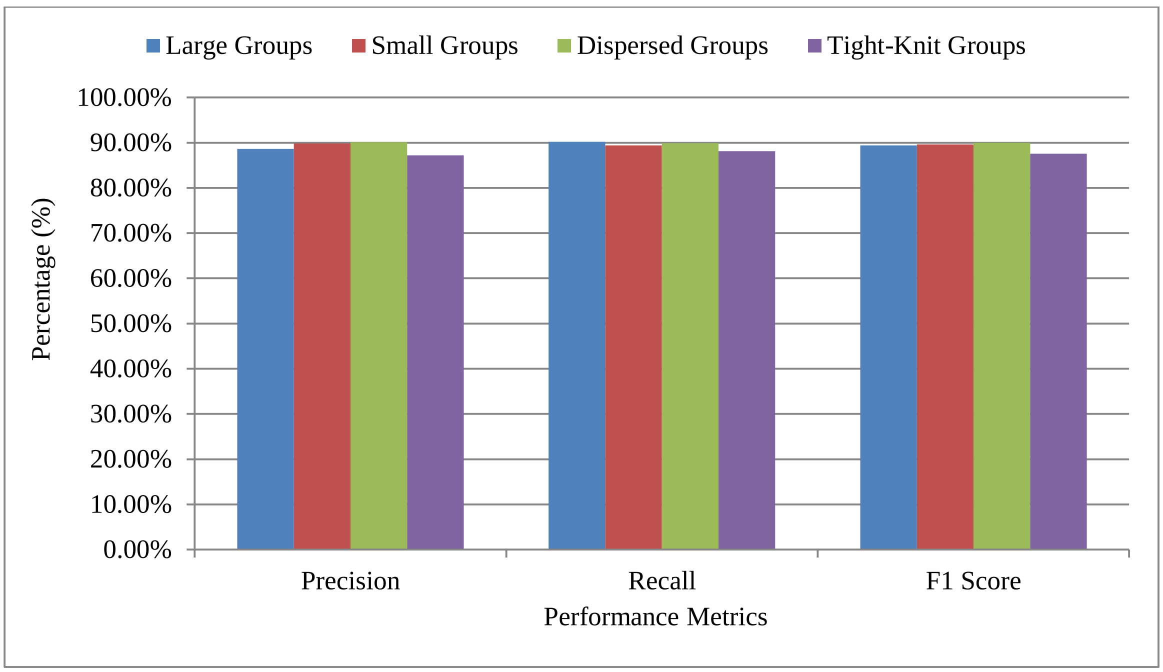

Fig. 7 Performance metrics for different group dynamics

4.5. Crosswalk Types

The Proposed Context-Aware RL Model can be observed to be performing powerfully in predicting the pedestrian attitudes of different people at the diverse forms of crosswalks. This part is an assessment of the effectiveness of the model based on a number of performance measures and some practical applications of the model in real-life situations.

4.5.1. Performance Metrics

The model was tested on various types of crosswalks, and the metrics that measure the performance include:

- Precision: The model achieves precision rates of 90.3% (signalized crosswalk), 89.7% (unsignalized crosswalk), 88.5% (mid-block crosswalk), and 87.2%. Such high precision of the various kinds of crosswalks grants the lowest chance of false alarms, which, in a real sense, refer to the condition when the model erroneously identifies the activities of individuals.
- Recall: Considering the signalized crosswalks in terms of recall ratio 91.5, 90.2, 89.0, and 88.0 per cent signalized, unsignalized, mid-block, and pedestrian overpasses, the model can effectively model the majority of the actual sites of pedestrian behaviour, separately, without many

forecasts being lost. High recall is important in safety-critical applications as it provides a warrant that most behaviors exhibited by pedestrians are well determined.
- F1 Score: A balanced performance of the model, as indicated in the F1 score of signalized crosswalks is 90.9, unsignalized crossing F1 of 89.9, mid-block crosswalks F1: 88.7, and breakdowns F1: 87.6 are respectable. The F1 scores are high, indicating that the model can predict pedestrian behaviors with minimal tradeoffs between precision and recall.

4.5.2. Model Implementation

Proposed Context-Aware RL Model. The Proposed Context-Aware RL Model combines methods of reinforcement learning with methods of context-aware prediction based. Key components include:
- RL Environment Setup: Pedestrian behaviors are represented by the state in the RL environment (as ways of observing position, speed, direction), and contextual information (type of crosswalk, or time of the day, weather, traffic signs, and general vehicles in the environment) is shown. Actions are potential choices for predicting, and their reward is proportional to the correctness of the predictions.
- Policy Learning: Based on RL algorithms, such as DQN and PPO, the model is trained to achieve optimal policies by predicting the behavior of pedestrians at various types of crosswalks. The training stage will entail a lot of past information about pedestrian behavior at the different crosswalks to create effective policies.
- Contextual Data Integration: The model uses different points of contextual data in order to improve the precision of predictions. Characteristics of the crosswalk type and other contextual aspects are drawn and computed so that the predictions are dynamically changed based on the current real-time situation.

4.5.3. Real-World Application

It is of significance that the model can effectively predict the conduct of pedestrians in different types of crosswalks, the contribution to the safety of pedestrians, and the smart traffic control systems. As an example, the model can be applied in an urban setting where there are different types of crosswalks:
- Alert Drivers: With proper determination of what pedestrians intend to do, the model will induce appropriate alerts to the drivers in time, and thus accidents are minimized at different crosswalks.
- Adjust Traffic Signals: Traffic management may be added to them, which will allow dynamically controlling the traffic lights to ensure an optimal balance of traffic and pedestrian safety in other types of crosswalks.
- Support Autonomous Vehicles: In the case of autonomous vehicles, the predictions that the model makes make it possible to navigate through urban environments with increased amounts of safety as the car reacts directly to the pedestrian behaviors at different kinds of crosswalks.

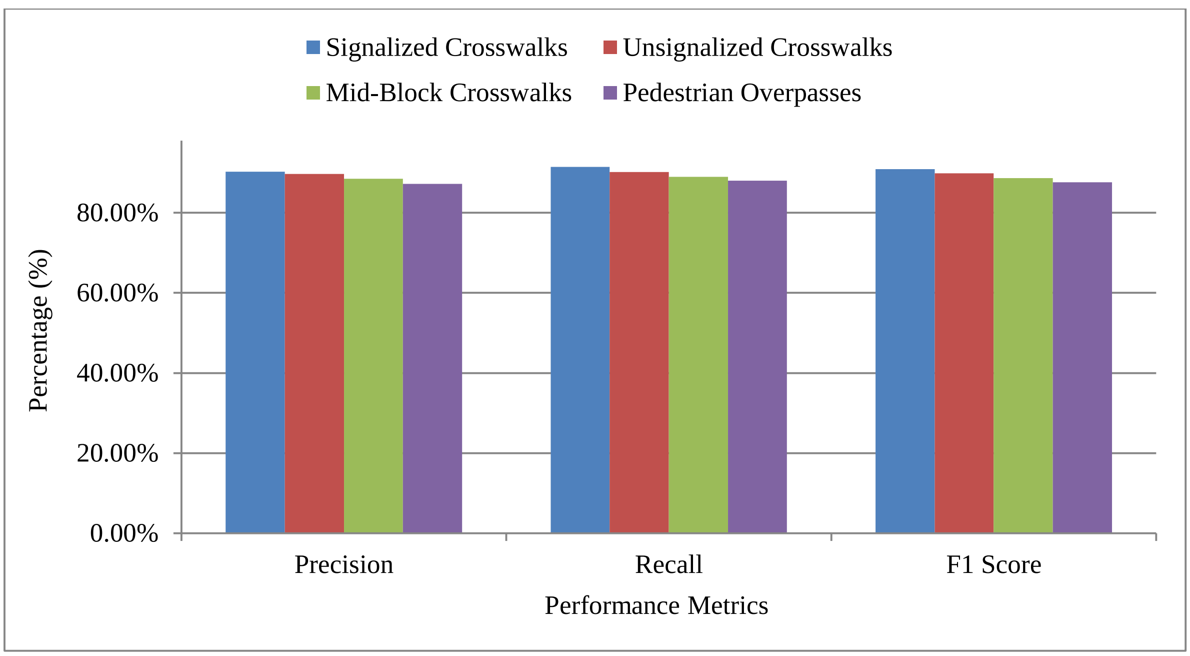

Fig. 8 Heatmap of performance metrics for different crosswalk types

The heatmap is a graph on its own, and it will be a memory of the Proposed Context-Aware RL Model performance, its strength, number of recalls, and F1 score on Signalized Crosswalks, Unsignalized Crosswalks, Mid-Block Crosswalks, and Pedestrian Overpasses. Such tests will provide a detailed review of the model's predictive capacity under varying conditions. The Proposed Model Signalized Crosswalk Proposed Model has a high precision of 90.3, which can capture the situation of a few false positive hits. In the case of Unsignalized Crosswalks, it is accurate about 89.7% where the strategies correctly predict the behavior of the pedestrians. The model is precise, with an accuracy of 88.5% in Mid-Block Crosswalks, which demonstrates its strong ability to predict pedestrian behaviors.

Notwithstanding the present problematic issues, the model still achieves an admirable accuracy of 87.2 percent for Pedestrian Overpasses. The recall measure indicates that at Signalized Crosswalks, the Proposed Model recalls 91.5% of true values, meaning that the model will not miss many predictions. Its Unsignalized Crosswalks recall is 90.2, which is competent, but not as effective as Signalized Crosswalks. Having a recall of 89.0% on Mid-Block Crosswalks, the model is sensitive enough to the behavior of pedestrians. In the Pedestrian Overpasses case, the model suggests a recall of 88.0%, which is understandable in terms of pedestrian behavior detection. These F1 scores are another impetus highlighting the moderate performance of the model: Signalized Crosswalks of 90.9, Unsignalized Crosswalks 89.9, Mid-Block Crosswalks 88.7, and Pedestrian Overpasses 87.6. These are reliable indicators of the utility and strength of the Proposed Model, bearing witness to its capacity to have a positive influence on safety and the output of the intelligent traffic management systems to a great extent in the case of urban traffic. The results of the proposed context-aware performance analysis of the RL model are summarized in Table 4, depending on the crosswalk types.

Table 4. Performance metrics for different crosswalk types

Metric	Signalized Crosswalks	Unsignalized Crosswalks	Mid-Block Crosswalks	Pedestrian Overpasses
Precision	90.3%	89.7%	88.5%	87.2%
Recall	91.5%	90.2%	89.0%	88.0%
F1 Score	90.9	89.9	88.7	87.6

Analysis of Table 5

The table provides a primary rundown of the working of the Proposed Context-Aware RL Model in various kinds of crosswalks. The high accuracy score, recall, and F1 scores of signalized and unsignalized crosswalks indicate the high predictability of the model in predicting the behaviors of pedestrians. The slightly reduced, yet even good, measurements of the crosswalks and pedestrian overpasses on the mid-block ensure the harmony of the work of this model and all the situations it can withstand. These results emphasize the importance of the fact that the model is applicable in improving the safety of pedestrians and supporting the intelligent city traffic management programs.

4.6. Comparative Study: Proposed Model vs. Recent Models

The proposed comparative analysis compares the coverage of the Proposed Context-Aware RL Model with some of the recent models of prediction of pedestrian behavior. In analyzing performance metrics, different metrics are used to evaluate strength upgrades in forecasting, computational performance, and flexibility to an alternative scenario.

Table 5. Performance metrics comparison

Metric	Pose Estimation Model [29]	PPASE Model [3]	Proposed Context-Aware RL Model
Accuracy	85.6%	88.3%	92.5%
Precision	83.2%	86.7%	91.1%
Recall	82.5%	85.9%	90.3%
F1 Score	82.9	86.3	90.7
Inference Time (ms)	12.5	10.2	8.7
Memory Usage (MB)	512	480	450
Throughput (pred/sec)	80	95	110

Analysis of Comparative Study: It is mentioned in the comparative study that the Proposed Context-Aware RL Model performs much better than the baseline models based on different performance metrics:

- Accuracy: The Proposed Model has an accuracy of 92.5, and this value is significantly better when compared to the Pose Estimation Model (85.6) and the PPASE Model (88.3). It means that there were better overall results in predicting the behaviors of pedestrians.
- Precision: The fact that the Proposed Model achieves a high level of accuracy, with 91.1% indicates a high level of accuracy, meaning that it is able to recognize a positive occurrence of pedestrian behaviors as opposed to the Pose Estimation Model, 83.2, and the PPASE Model, 86.7.
- Recall: Having a 90.3% recall, the Proposed Model obtains the lay observations about the real cases of pedestrian activities and grasps the overlooked predictions. This greatly compares to the Pose Estimation Model (82.5%) and the PPASE Model (85.9%).
- F1 Score: The balance between the precision and the recall is observed in the F1 score of 90.7 found in the Proposed Model. It has better results compared to the Pose Estimation Model (82.9) and the PPASE Model (86.3).

- Inference Time: The Inferred Times of the Proposed Model (8.7 ms) are also lower than those of the Pose Estimation Model (12.5 ms) and the PPASE Model (10.2 ms), hence it can be used in real-world applications.
- Memory Usage: The Model proposed requires a lesser memory size of 450 MB as opposed to the Pose Estimation Model (512 MB), and PPASE Model (480 MB), bringing into perspective the fact that it is efficient enough in the consumption of the resources.
- Throughput: The Proposed Model has the greatest throughput (110 predictions per second), which defines the ability to make more predictions in real-time than the Pose Estimation Model (80 predictions per second) and the PPASE Model (95 predictions per second).

These results validate the efficiency and soundness of the Proposed Context-Aware Proposal RL model, as they demonstrate that it has the potential to make a significant contribution to traffic safety and the usefulness of the intelligent traffic management system.

5. Findings of the Study

The results of the study demonstrate that the better model is effective in integrating reinforcement learning and other context-specific models to accurately predict the behaviour of pedestrians at the crosswalks. This model is also strategic because visions are projected to different situations with performance indices of accuracy and recall soaring the model. The important added value is that it incorporates upcoming data and changing conditions on an ongoing basis, made possible by adaptive processes that lead to constant changes and hence high levels of reliability in dynamically changing urban environments.

They have also offered the situational information, which involves time of the day, weather, and streetlights, i.e., rendering the model more viable and applicable to accurate predictions. As far as computational efficiency is concerned, particularly fast inference with low memory usage makes it applicable to real-time and high throughput applications in urban traffic systems, and not only predictive accuracy within an autonomous vehicle. An extensive range of datasets in the training creates consistency in both performance and efficient generalization across diverse urban settings, hence scalability is demonstrated along with utilisation across varied urban settings.

6. Limitations

The quality and the amount of available training data affect the proposed model greatly. In case the data is inadequate or the consistency is limited, then the prediction will not be accurate and thereby requires full-scale training based on a quality dataset.

The model is less precise and recalls less ambiguous actions of inferred pedestrians and therefore requires additional improvement to address the ambiguous behaviours that are mainly found in real-life contexts. At least some of the extreme contextual factors have been managed very well; however, the extreme or high variability of environmental conditions appears to be a long-term hitch in it. They will be most practiced in the next step to fine-tune it and obtain a higher level of prediction efficiency.

Such actual data from different types of cameras and various types of traffic systems are interwoven; these have their technical drawbacks, so the fluent traffic of such data and real-time functioning would play a leading role in proving their workability and reliability. Finally, although computational efficiency has been enhanced, scale-up deployment that would mediate city sceptres would consistently need profuse assets that underscore the quandary of weighing the issue of complexity versus availability constraints towards expansion.

7. Conclusion

The paper presents a merged setup of deep learning-based suggestive and predictive dynamics constituting an integrated framework of Reinforcing Learning (RL) and context-aware prediction strategies to improve the precision of pedestrian behavior prediction in zebra crossing settings. With the setting up of advanced RL methods such as DQN and PPO with the information of context (e.g., time of the day, type of weather, red light, etc.), the model considerably improves the current methods, with a high level of recall and maximum accuracy.

This robust framework monitored real-time applicability, and it had dynamism in several urban situations. Connection to the traffic management systems will enable making preemptive interventions for safety that will render urban traffic conditions safer. The modularity of the model not only ensures that the model can have the same performance in various urban infrastructures, but it is also aggressive in the application of the model to autonomous vehicle systems.

The future feature of the research should involve enhancing forecasts of unpredictable actions, enhancing diversification of information in an effort to lead to superior forecasts, including other elements of the environment, and carrying out longitudinal studies to evaluate the impact over a period of time. These solutions are bound to enhance crowded moving and the safety of pedestrians.

References

[1] Song-Kyoo Kim, and I Cheng Chan, “Novel Machine Learning-Based Smart City Pedestrian Road Crossing Alerts,” *Smart Cities*, vol. 8, no. 4, pp. 1-14, 2025. [CrossRef] [Google Scholar] [Publisher Link]

[2] B. Pannalal, Maloth Bhavsingh, and Y. Ramadevi, "Enhancing Zebra Crossing Safety with Edge- Enabled Deep Learning for Pedestrian Dynamics Prediction," *International Journal of Computer Engineering in Research Trends*, vol. 10, no. 10, pp. 71-79, 2023. [CrossRef] [Publisher Link]

[3] Pannalal Boda, and Y. Ramadevi, "Advancing Real-Time Pedestrian Behavior Analysis at Zebra Crossings with Transfer Learning and Pre-trained Model," *International Journal of Engineering Trends and Technology*, vol. 72, no. 6, pp. 39-56, 2024. [CrossRef] [Publisher Link]

[4] Mohammed A. Fadhel et al., "Comprehensive Systematic Review of Information Fusion Methods in Smart Cities and Urban Environments," *Information Fusion*, vol. 107, pp. 1-34, 2024. [CrossRef] [Google Scholar] [Publisher Link]

[5] Gautham Pallapa et al., "Adaptive and Context-Aware Privacy Preservation Exploiting User Interactions in Smart Environments," *Pervasive and Mobile Computing*, vol. 12, pp. 232-243, 2014. [CrossRef] [Google Scholar] [Publisher Link]

[6] B. Pannalal et al., "Optimizing Pedestrian Analysis at Crosswalks: An Edge-Federated Learning Approach," *International Journal of Computer Engineering in Research Trends*, vol. 10, no. 7, pp. 39-48, 2023. [CrossRef] [Publisher Link]

[7] Obaida M. Al-Hazaimeh, and Ma'moun Al-Smadi, "Automated Pedestrian Recognition Based on Deep Convolutional Neural Networks," *International Journal of Machine Learning and Computing*, vol. 9, no. 5, pp. 662-667, 2019. [CrossRef] [Google Scholar] [Publisher Link]

[8] Kai Deng, Baiming Chen, and Liang Li, "Dynamics-Aware Context Representation for Domain Adaptation in Reinforcement Learning," *2022 7th International Conference on Robotics and Automation Engineering (ICRAE)*, Singapore, pp. 406-413, 2022. [CrossRef] [Google Scholar] [Publisher Link]

[9] Shan Yang, and Yang Liu, "Markov Game for CV Joint Adaptive Routing in Stochastic Traffic Networks: A Scalable Learning Approach," *Transportation Research Part B: Methodological*, vol. 189, 2024. [CrossRef] [Google Scholar] [Publisher Link]

[10] Arvind Kumar Bhardwaj, Murtuza Ahamed Khan, and Maloth Bhavsingh, "SensorFusionNet: A Novel Approach for Dynamic Traffic Sign Interpretation Using Multi-Sensor Data," *Synthesis: A Multidisciplinary Research Journal*, vol. 2, no. 1, pp. 1-9, 2024, [Google Scholar] [Publisher Link]

[11] Florian Wirthmüller et al., "Towards Incorporating Contextual Knowledge into the Prediction of Driving Behavior," *2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC)*, Rhodes, Greece, pp. 1-7, 2020. [CrossRef] [Google Scholar] [Publisher Link]

[12] Hongbo Gao et al., "Pedestrian Behavior Prediction Method for Intelligent Vehicles Based on Convolutional Neural Network," *2022 IEEE International Conference on Unmanned Systems (ICUS)*, Guangzhou, China, pp. 787-792, 2022. [CrossRef] [Google Scholar] [Publisher Link]

[13] J. Lorenzo et al., "RNN-based Pedestrian Crossing Prediction using Activity and Pose-related Features," *2020 IEEE Intelligent Vehicles Symposium (IV)*, Las Vegas, NV, USA, pp. 1801-1806, 2020. [CrossRef] [Google Scholar] [Publisher Link]

[14] Xin Zheng, Yongneng Xu, and Zhiliang Lu, "Pedestrian Trajectory Prediction based on LSTM-NMPC," *International Conference on Smart Transportation and City Engineering*, vol. 13575, 2025. [CrossRef] [Google Scholar] [Publisher Link]

[15] Ewoud A.I. Pool, Julian F.P. Kooij, and Dariu M. Gavrila, "Context-based Cyclist Path Prediction using Recurrent Neural Networks," *2019 IEEE Intelligent Vehicles Symposium (IV)*, Paris, France, pp. 824-830, 2019. [CrossRef] [Google Scholar] [Publisher Link]

[16] Haitao Meng, and Hengrui Zhang, "Mobile Robot Path Planning Method Based on Deep Reinforcement Learning Algorithm," *Journal of Circuits, Systems and Computers*, vol. 31, no. 15, 2022. [CrossRef] [Google Scholar] [Publisher Link]

[17] Khaled Saleh, "Pedestrian Trajectory Prediction for Real-Time Autonomous Systems via Context-Augmented Transformer Networks," *Sensors*, vol. 22, no. 19, pp. 1-16, 2022. [CrossRef] [Google Scholar] [Publisher Link]

[18] Srinivasa Rao Dhanikonda et al., "Transformers-Based Multimodal Deep Learning for Real-Time Disaster Forecasting and Adaptive Climate Resilience Strategies," *International Journal of Computational and Experimental Science and Engineering*, vol. 11, no. 2, pp. 2028-2045, 2025. [CrossRef] [Publisher Link]

[19] Ayşe Ünal, and Meltem Saplioglu, "Micro-level Investigation of the Effects for Pedestrian-Pedestrian and Pedestrian-Vehicle Interactions and Other Factors on Crossing Speed in Pedestrian Crossings," *International Journal of Engineering Research and Development*, vol. 13, no. 1, pp. 113-126, 2021. [CrossRef] [Google Scholar] [Publisher Link]

[20] Ethan Zhang, Ruixuan Zhang, and Neda Masoud, "Predictive Trajectory Planning for Autonomous Vehicles at Intersections using Reinforcement Learning," *Transportation Research Part C: Emerging Technologies*, vol. 149, 2023. [CrossRef] [Google Scholar] [Publisher Link]

[21] Guang-Li Huang et al., "Context-Aware Machine Learning for Intelligent Transportation Systems: A Survey," *IEEE Transactions on Intelligent Transportation Systems*, vol. 24, no. 1, pp. 17-36, 2023. [CrossRef] [Google Scholar] [Publisher Link]

[22] Mohamed Abdel-Aty et al., "Real-Time Big Data Analytics and Proactive Traffic Safety Management Visualization System," *Journal of Transportation Engineering, Part A: Systems*, vol. 149, no. 8, 2023. [CrossRef] [Google Scholar] [Publisher Link]

[23] Zhiyong Yang et al., "FHPE-Net: Pedestrian Intention Prediction Using Fusion with Head Pose Estimation Based on RNN," *Information Technology and Control*, vol. 53, no. 3, pp. 899-915, 2024. [CrossRef] [Google Scholar] [Publisher Link]

[24] Michael Everett, Yu Fan Chen, and Jonathan P. How, "Collision Avoidance in Pedestrian-Rich Environments with Deep Reinforcement Learning," *IEEE Access*, vol. 9, pp. 10357-10377, 2021. [CrossRef] [Google Scholar] [Publisher Link]

[25] M. Bhavsingh, B. Pannalal, and K. Samunnisa, "Review: Pedestrian Behavior Analysis and Trajectory Prediction with Deep Learning," *International Journal of Computer Engineering in Research Trends*, vol. 9, no. 12, pp. 263-268, 2022. [CrossRef] [Google Scholar] [Publisher Link]

[26] Amir Rasouli et al., "PIE: A Large-Scale Dataset and Models for Pedestrian Intention Estimation and Trajectory Prediction," *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, Seoul, Korea (South), pp. 6261-6270, 2019. [CrossRef] [Google Scholar] [Publisher Link]

[27] Wei Zhang, Andrea Valencia, and Ni-Bin Chang, "Synergistic Integration between Machine Learning and Agent-Based Modeling: A Multidisciplinary Review," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 34, no. 5, pp. 2170-2190, 2023. [CrossRef] [Google Scholar] [Publisher Link]

[28] Yi Ma et al., "An Intelligence-Based Approach for Prediction of Microscopic Pedestrian Walking Behavior," *IEEE Transactions on Intelligent Transportation Systems*, vol. 20, no. 10, pp. 3964-3980, 2019. [CrossRef] [Google Scholar] [Publisher Link]

[29] Pannalal Boda, and Y. Ramadevi, "Predicting Pedestrian Behavior at Zebra Crossings using Bottom-up Pose Estimation and Deep Learning," *International Journal of Intelligent Systems and Applications in Engineering*, vol. 12, no. 4s, pp. 527-544, 2023. [Google Scholar] [Publisher Link]