Human Factors in Design, Engineering, and Computing, Vol. 238, 2026, 2570–2581
https://doi.org/10.54941/ahfe1008482

A Real-Time Edge-Based Multimodal Safety Assistance System for Forklift Operations

Thasnimol Valuthottiyil Shajahan[1], Babji Srinivasan[1,3], and Rajagopalan Srinivasan[2,3]

[1]Department of Applied Mechanics and Biomedical Engineering, Indian Institute of Technology Madras, 600036, India
[2]Department of Chemical Engineering, Indian Institute of Technology Madras, Chennai 600036, India
[3] American Express Lab for Data Analytics, Risk & Technology, Indian Institute of Technology Madras, 600036, India

ABSTRACT

Forklifts are essential to industrial and warehouse operations, but their use in confined and dynamic environments creates safety risks associated with restricted visibility, unsafe maneuvering, delayed operator response, and improper load handling. Conventional manual safety checks and road-vehicle telematics are not sufficient for forklift environments, which involve short-range interactions, narrow aisles, reversing Maneuvers, and continuously changing operating conditions. Existing forklift safety systems have demonstrated the usefulness of cameras, radar, LiDAR, and wireless identification for visibility improvement and collision avoidance. However, these systems commonly focus on individual sensing functions and provide limited integration of surrounding hazards, operator attention, and vehicle state. To address this limitation, this study proposes a real-time edge-based multimodal safety assistance system for forklift operations. The system integrates four-side LiDAR sensing, operator viewing-direction and attention information from an in-cabin monitoring module, and vehicle-state information including travel direction, gear-related state, and speed. These inputs are processed by an embedded edge device that applies context-aware decision logic and generates audio warnings. Validation in six short warehouse runs showed that multimodal fusion increased final-decision accuracy from 83.1% to 97.0% and reduced audio-warning candidates by 84.4% compared with module-only proximity warning logic. These initial results indicate the potential of the proposed system to improve warning relevance and support safer forklift operation under the evaluated warehouse conditions.

Keywords: Forklift safety, Multimodal safety assistance, Operator attention monitoring, Real-time edge system

INTRODUCTION

Forklifts are widely used in warehouse and industrial logistics for material handling, storage, loading, and internal transport. However, they frequently operate near workers, racks, pallets, loading structures, and other equipment

Received August 2, 2026; Revised September 27, 2026; Accepted October 2, 2026; Available online December 1, 2026

© 2026 The Authors. This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License.
For more information, see https://creativecommons.org/licenses/by-nc-nd/4.0/

within confined and continuously changing environments. Maneuvers such as reversing, turning, positioning, and travelling with a load can restrict visibility and increase the visual and cognitive demands on the operator. Consequently, collisions may occur even at low speeds when nearby hazards are not detected or interpreted in time.

Forklift safety depends not only on the presence of an obstacle but also on whether that obstacle is relevant to the current maneuver and whether the operator has attended to the corresponding hazard region. Situation awareness may decrease during demanding activities such as reversing, lifting, lowering, and travelling with a load, while misperception and habitual behaviour may further contribute to unsafe decisions (Choi et al., 2020; Arias-Portela et al., 2024). Conventional proximity-warning systems may generate an alert whenever an object enters a predefined distance region, even when the forklift is stationary, the object lies outside the active travel direction, or the operator has already observed it. In warehouse environments, where racks, walls, pallets, and loading surfaces may remain close to the vehicle for extended periods, such repeated alerts can reduce warning relevance and contribute to alarm fatigue.

Existing forklift and industrial-vehicle safety systems use LiDAR, radar, cameras, and wireless localization to detect nearby workers and obstacles, while operator-monitoring approaches use eye movement, head orientation, and physiological information to estimate attention and situation awareness (Lang, 2018; Wei et al., 2018; Neruda et al., 2025; Kang et al., 2023; Li et al., 2024; Kang et al., 2024). Field-of-view- and head-pose-based methods have also shown that operator-awareness information can improve warning relevance (Ray and Teizer, 2012; Chan et al., 2020). However, hazard sensing, operator monitoring, vehicle movement assessment, and warning selection are still commonly developed and evaluated as separate functions. As a result, the real-time integration of obstacle severity, movement relevance, and recent operator attention into a single forklift warning decision remains insufficiently explored.

To address this gap, this study proposes a real-time edge-based multimodal safety assistance system for warehouse forklift operations. The system integrates four-side LiDAR sensing, operator-attention information from an in-cabin monitoring module, and vehicle-state information obtained from a vehicle-telemetry module. The module outputs are combined within one-second decision windows and processed by an edge-based fusion layer that evaluates each obstacle according to its proximity region, relevance to the current forklift movement, and the operator's recent attention. The main contribution is a forklift-specific decision system that distinguishes hazards requiring operator notification from valid detections that only need to be recorded. The system is evaluated through module-level and decision-level analyses in a preliminary warehouse study to determine whether multimodal fusion can improve final-decision accuracy and reduce repeated, non-actionable audio warnings while preserving alerts for movement-relevant hazards.

LITERATURE REVIEW

Forklift-safety research has mainly examined surrounding-hazard detection, operator-awareness assessment, and warning-generation methods. The following review considers these areas and their integration for context-aware warning selection.

Different sensing technologies have been used to improve hazard detection around forklifts and other industrial vehicles. Camera-based systems support person and object recognition but may be affected by worker posture, partial occlusion, lighting variation, and dependence on site-specific training data (Lang, 2018). LiDAR provides direct distance and spatial information, while combining LiDAR with camera data has been shown to reduce false-positive detections compared with LiDAR-only processing (Wei et al., 2018). Wireless identification and radar provide another approach: BLE can identify nearby workers, while 77 GHz radar can estimate obstacle distance and relative motion for proximity-based warning generation (Neruda et al., 2025). Distance and relative motion can also be used to estimate collision urgency through time-to-collision measures (Lee, 1976). However, these approaches mainly determine whether a hazard is present and how close or rapidly approaching it is, rather than whether it is relevant to the current forklift operation.

A second group of studies has examined forklift-operator attention and situation awareness. Virtual-reality and field investigations have shown that awareness may decrease during demanding activities such as reversing, lifting, lowering, positioning, and travelling with a load (Choi et al., 2020; Arias-Portela et al., 2024). Eye-tracking measures, including gaze location, fixation frequency, and viewing duration, have been used to evaluate whether the operator attends to task-relevant regions (Kang et al., 2023). Head orientation provides an additional indication of viewing direction, while physiological measures such as EEG have also been investigated for operator-state recognition (Li et al., 2024; Kang et al., 2024). These studies demonstrate that operator attention and awareness can be assessed under realistic conditions, but they mainly focus on operator-state recognition rather than combining this information with live obstacle and vehicle-state data for warning selection.

Research on selective warnings has shown that proximity alone is not sufficient to determine whether an alert is required. Incorporating field-of-view information into proximity-warning logic reduced redundant warnings compared with distance-only methods, showing that awareness information can improve warning relevance and reduce alarm fatigue (Chan et al., 2020). Head-pose estimation has also been used to define dynamic blind-spot regions around industrial equipment, allowing warning regions to vary with the operator's viewing direction rather than relying only on fixed geometry (Ray and Teizer, 2012). Warning presentation can further influence operator response, with multimodal feedback improving stopping performance in industrial-vehicle simulations (Garcia-Carrillo et al., 2024). However, effective warning generation still depends on whether the system can interpret the surrounding hazard, vehicle movement, and operator attention together.

Overall, previous studies have demonstrated effective hazard sensing and operator-state assessment, but the combined use of obstacle severity, vehicle movement, and operator attention for real-time forklift warning selection remains limited. This gap motivates the proposed integration of LiDAR proximity, vehicle-state, and operator-attention information.

METHODOLOGY

This study presents a real-time edge-based multimodal safety assistance system for forklift operations in warehouse environments. Four LiDAR modules provide surrounding-obstacle information, an in-cabin monitoring module estimates the operator state, and a Vehicle Telemetry Module provides vehicle-movement information. The module outputs are processed locally and transmitted to a common edge-processing unit. Within consecutive one-second decision windows, the edge unit combines the latest available information and evaluates obstacle proximity, movement relevance, and operator attention. Based on this context, the system assigns a final safety state and determines whether an operator-facing audio warning is required. This evaluation avoids unnecessary alerts when the forklift is stationary, the detected obstacle lies outside the active movement direction, or the operator has already attended to the corresponding hazard region.

System Architecture and Hardware Configuration

The overall system architecture is shown in Fig. 1. It consists of four LiDAR sensing modules, one operator-monitoring module, one Vehicle Telemetry Module, a common edge-processing unit, a speaker-based audio unit, and an event-logging component.

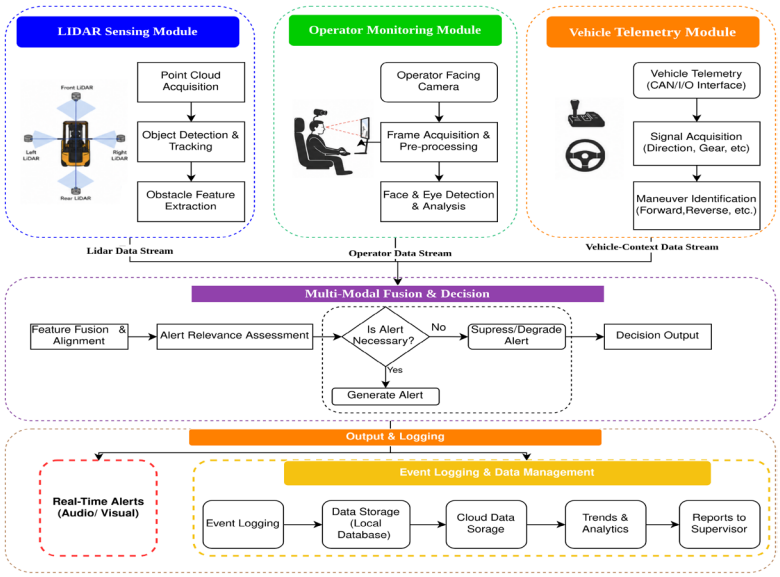


Figure 1: Architecture of the proposed edge-based multimodal forklift safety assistance system.

The four LiDAR units are positioned at the front, rear, left, and right sides of the forklift. Each LiDAR is connected to a Raspberry Pi Compute Module 4 (CM4)-based embedded node for local data acquisition and obstacle-state generation. Each sensor performs a complete 360° scan; however, only the angular sector relevant to its mounting position is used for obstacle evaluation. The active sector is configured separately for the front, rear, left, and right modules according to the sensor orientation and required coverage region.

The operator-monitoring module uses an operator-facing camera connected to a separate CM4-based embedded node installed inside the cabin. The camera captures the operator's face and head orientation. The local vision pipeline estimates face visibility, head pose, viewing direction, distraction state, and supplementary drowsiness-related indicators.

The Vehicle Telemetry Module acquires speed, gear-related state, travel direction, and movement information. These signals are used to represent stationary, forward, reverse, and turning conditions.

The outputs of the three modules are transmitted through wired Universal Asynchronous Receiver–Transmitter (UART) communication to the common edge-processing unit. The edge unit receives the structured state messages, combines them within the decision window, applies the multimodal safety rules, controls the audio output, and stores the corresponding event information.

Data-Stream Acquisition and Feature Extraction

Each module converts its raw measurements into a structured real-time state stream before transmitting the information to the edge-processing unit.

LiDAR sensing module: The four LiDAR units continuously acquire obstacle-distance measurements within their assigned spatial regions. Each sensor operates at approximately 10 Hz. For obstacle evaluation, measurements are restricted to the configured angular sector associated with the sensor position. Measurements outside the selected sector, outside the valid sensing range, or identified as invalid returns are excluded before the nearest obstacle distance is calculated. The angular sectors and distance limits are configurable and can be adjusted according to the sensor placement and operating environment. For side s at time t, the nearest valid obstacle distance is calculated as,

$$d_{min}^{s(t)}(t) = \min_{i \in V_s(t)} d_i^{(s)}(t) \tag{1}$$

where $d_i^{(s)}(t)$ is the i-th valid distance measurement obtained from side s at time t, and $V_s(t)$ is the set of valid LiDAR measurements satisfying the configured angular and distance criteria.

The resulting state contains the detected side, nearest obstacle distance, detection status, and assigned proximity region. Configurable warning and danger thresholds classify each measurement as normal, warning, or high priority. These thresholds are installation-dependent and were configured for the forklift geometry and warehouse validation setup.

Operator-monitoring module: The operator-facing camera stream is processed through face detection, facial-landmark extraction, and head-pose estimation. Face visibility is evaluated first because reliable head-pose and viewing-direction estimation requires a valid face observation. When the face is occluded, outside the camera field of view, or not detected reliably, the operator state is marked as unverified rather than being interpreted as attentive or distracted.

When the face is visible, yaw, pitch, and roll are estimated and used to classify the operator's viewing direction into task-relevant regions, including front, left, right, rear-check, and downward. The estimated viewing direction is compared with the detected obstacle side to determine whether the operator has recently attended to the corresponding hazard region. Distraction is identified using configurable angular and temporal limits.

Supplementary drowsiness indicators, including eye closure, blink behaviour, PERCLOS, and yawning, are also calculated but are not evaluated separately in this study. The underlying monitoring method is described in Shajahan et al. (2024).

Vehicle-telemetry module: The Vehicle Telemetry Module receives speed, gear-related information, travel-direction signals, and available movement information through the vehicle interface. These inputs are used to classify stationary, forward, reverse, and turning conditions and determine whether a detected obstacle is relevant to the current movement.

Multimodal Fusion and Safety Decision Logic

The LiDAR, operator-state, and vehicle-state streams are generated independently and transmitted to the edge-processing unit. Exact timestamp synchronization is not applied in the current implementation. Instead, the most recent information from the three modules is combined within consecutive one-second decision windows. Each window therefore represents the obstacle condition, operator state, and vehicle state available during the same operating interval.

The fusion process follows a fixed sequence. First, the detected obstacle side is compared with the current forklift movement. Front detections are treated as movement-relevant during forward travel, rear detections during reverse travel, and detections on the corresponding side during turning. Detections outside the active movement direction are retained in the event log but normally do not generate audio unless they are classified as danger-region events. Second, the fusion layer evaluates the proximity region assigned by the relevant LiDAR module. Danger-region detections retain high priority, whereas warning-region detections are evaluated further using movement relevance and operator attention.

Third, the operator state is examined to determine whether the operator has recently attended to the corresponding hazard region. A movement-relevant warning-region obstacle is assigned a higher safety state when the operator is looking away, distracted, or the operator state is unverified. Recent attention to a non-critical warning-region obstacle may reduce the

event to a logging-only condition. However, operator attention does not suppress a danger-region alert.

The vehicle movement condition is also considered for persistent detections. When the forklift is stationary, repeated detections from nearby racks, walls, pallets, or loading structures may be stored without repeated audio output unless the obstacle remains within the danger region. During forward, reverse, or turning movement, detections aligned with the active movement direction remain eligible for warning generation.

Based on the combined information, the system assigns one of four final safety states: No Alert, Log Only, Warning, or Critical. No Alert indicates that no relevant hazard is present. Log Only indicates that a valid detection is retained without an audio warning. Warning indicates that a movement-relevant hazard requires operator notification, whereas Critical represents a high-priority hazard requiring immediate attention. These states are generated by the multimodal fusion process and are not direct outputs of the individual modules.

Real-Time Alerts, Event Logging, and Data Management

After the final safety state is assigned, Warning and Critical events are communicated through the speaker-based audio channel. No Alert and Log Only events are stored without operator-facing audio. To reduce repeated audio activation from the same persistent condition, the system maintains the active event state across consecutive decision windows. Audio is not repeated while the obstacle side, proximity region, vehicle movement state, and warning level remain unchanged. A new audio alert is generated when a new movement-relevant obstacle is detected, the obstacle enters a higher-risk proximity region, or a change in vehicle movement makes an existing obstacle relevant.

For each decision window, the system stores a time-stamped record containing the LiDAR state, operator state, vehicle state, final safety decision, and corresponding decision context. These records support system validation, event review, and subsequent refinement of the decision rules. The logged events are stored locally and can be synchronized with cloud storage to support trend analysis, performance reporting, and supervisor review.

RESULTS AND DISCUSSION

The experimental evaluation examined whether the proposed system could convert individual module outputs into context-relevant forklift safety decisions. This was an initial system-level validation rather than an exhaustive field-performance study. Six short warehouse validation runs were conducted under stationary, forward, reverse, and turning conditions. The LiDAR, operator-monitoring, and vehicle-state outputs were recorded together with an external reference video, which was used to annotate obstacle position,

vehicle movement, operator attention, and the expected safety response. The evaluated scenarios are summarized in Table 1.

Table 1: Warehouse validation scenarios and evaluation purpose.

Scenario	Obstacle Condition	Vehicle State	Operator State	Purpose
Stationary	Rack/wall	Stationary	Attentive/not-detected	Evaluate suppression of repeated alerts
Forward	Front/side obstacle	Forward	Front/away	Check path relevance
Reverse	Rear/side obstacle	Reverse	Rear/away	Check rear risk
Turning	Turn side obstacle	Turn	Turn side/away	Check turn-side risk

The LiDAR streams were processed at approximately 10 Hz, the operator-facing camera at approximately 4 frames/s, and the vehicle-state stream at approximately 10 Hz. Across the six runs, the analysis included 327 annotated LiDAR events, 102 labelled vehicle-state observations, and 166 one-second multimodal decision windows. Performance was calculated as the percentage of observations that agreed with the reference annotation:

$$Accuracy(\%) = \frac{N_{correct}}{N_{total}} \times 100 \tag{2}$$

where $N_{correct}$ is the number of correctly identified observations and N_{total} is the total number of evaluated observations. This metric was applied to LiDAR event detection, vehicle-state classification, final safety decisions, and audio-warning decisions. For LiDAR, this metric corresponds to the event-detection rate with respect to the annotated events; false-positive detections were not quantified separately because persistent returns from racks and walls made independent false events difficult to delimit. For audio-warning decisions, each window was labelled according to whether operator-facing audio was expected, and the output was counted as correct when the system audio decision matched this label.

For the operator-monitoring module, usable coverage was defined as the percentage of processed frames for which valid face, head-pose, and viewing-direction information was available. Frames in which the face was occluded, outside the camera field of view, or not detected reliably were retained as an unverified state rather than being interpreted as operator inattention. The reduction in audio-warning candidates was calculated as

$$Alert\,reduction(\%) = \frac{N_{before} - N_{after}}{N_{before}} \times 100 \tag{3}$$

where N_{before} is the number of audio-warning candidates generated by module-only proximity logic and N_{after} is the number retained after multimodal decision processing.

Table 2: LiDAR event-detection performance by sensor position.

Sensor Position	Correct Detections	Missed Detections	Annotated Events	Detection Rate
Front	30	3	33	90.9%
Left	34	3	37	91.9%
Right	72	6	78	92.3%
Rear	165	14	179	92.2%
Overall	301	26	327	92.0%

Module-Level Analysis

The module-level analysis examined the consistency of the three modules. LiDAR performance was evaluated by comparing the detected side and event interval with the reference video annotation. The four sensor positions achieved detection rates between 90.9% and 92.3%, with an overall rate of 92.0% (301/327), as shown in Table 2. Missed detections were mainly associated with obstruction by the mast, forks, or load, while persistent returns from racks and walls made some event boundaries difficult to identify.

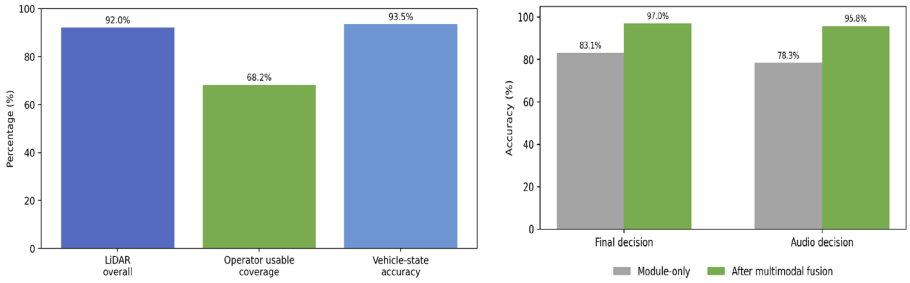


Figure 2: (a) LiDAR event-detection rate, operator-monitoring usable coverage, and vehicle-state classification accuracy; (b) final-decision and audio-decision accuracy before and after multimodal fusion.

The operator-monitoring module provided usable information in 68.2% of frames; the remaining 31.8% were classified as unverified because of occlusion, operator movement, or loss of face visibility. The Vehicle Telemetry Module achieved 93.5% classification accuracy across the labelled stationary, forward, reverse, and turning observations. The lower operator-monitoring coverage remains a practical limitation, although the explicit unverified state allowed this uncertainty to be considered by the fusion layer. The comparison of LiDAR event-detection rate, operator-monitoring usable coverage, and vehicle-state classification accuracy is shown in Fig. 2(a).

Multimodal Decision Analysis

The multimodal analysis evaluated each one-second decision window as a complete operating context rather than as an isolated sensor detection. For each window, the expected safety response was determined from the annotated obstacle side and proximity region, vehicle movement, and operator attention. For comparison, the module-only baseline treated each valid proximity detection as an independent alert candidate without considering vehicle movement or operator attention, so that each proximity region was mapped directly to a safety state and audio decision.

The module-only baseline agreed with the reference safety decision in 138 of the 166 windows, corresponding to a final-decision accuracy of 83.1%. After multimodal processing, 161 of the 166 windows were classified correctly, increasing final-decision accuracy to 97.0%. Audio-decision accuracy similarly increased from 78.3% to 95.8% after multimodal processing. The final-decision and audio-decision comparison is presented in Fig. 2(b).

The improvement was mainly observed when proximity alone did not represent the complete hazard context. These cases included persistent rack or wall detections while the forklift was stationary, obstacles outside the active movement direction, and warning-region obstacles already observed by the operator. In these situations, the multimodal logic retained the detections in the event log while avoiding unnecessary operator-facing audio. The fusion process preserved warnings for movement-relevant hazards by prioritizing front obstacles during forward travel, rear obstacles during reverse travel, and turn-side obstacles during turning. Warning priority increased when the operator was looking away or unverified.

Audio-warning candidates decreased from 90 to 14, corresponding to an 84.4% reduction within the evaluated runs. By design, this reduction does not arise from suppressing danger-region alerts, since operator attention and movement direction do not suppress danger-region events in the fusion logic. Overall, multimodal processing distinguished hazards requiring immediate notification from persistent or non-relevant detections requiring only logging within the evaluated scenarios.

Limitations

The validation comprised six short runs (166 one-second decision windows, 327 annotated LiDAR events, and 102 vehicle-state observations) conducted on a single forklift in one warehouse. The results should therefore be regarded as preliminary evidence for the tested setup rather than a demonstration of operational safety benefit. The operator-monitoring module provided usable information in only 68.2% of frames, so a considerable proportion of windows relied on the unverified operator state, which increases warning priority and may limit the achievable alert reduction. In addition, the thresholds and decision rules were configured for the tested scenarios and have not been validated across different forklift geometries, loads, lighting conditions, or operators.

CONCLUSION

The proposed system provided initial evidence that combining LiDAR proximity, vehicle-state, and operator-attention information can improve context-aware forklift warning decisions. In the warehouse validation, multimodal fusion increased final-decision accuracy from 83.1% to 97.0% and audio-decision accuracy from 78.3% to 95.8%, while reducing audio-warning candidates by 84.4%. At the module level, the system achieved a 92.0% LiDAR event-detection rate, 68.2% operator-monitoring usable coverage, and 93.5% vehicle-state classification accuracy. Future work will include longer trials across multiple operators, forklift configurations, and warehouse conditions, improved operator-monitoring coverage, and evaluation on other industrial vehicles.

REFERENCES

Arias-Portela, C.Y., Mora-Vargas, J., Caro, M. and Salinas-Navarro, D.E. (2024), "Situational awareness errors in forklift logistics operations: A multiphase eye-tracking and think-aloud approach," *Logistics*, Vol. 8, No. 4, Article 124. DOI: 10.3390/logistics8040124.

Chan, K., Louis, J. and Albert, A. (2020), "Incorporating worker awareness in the generation of hazard proximity warnings," *Sensors*, Vol. 20, No. 3, Article 806. DOI: 10.3390/s20030806.

Choi, M., Ahn, S. and Seo, J.O. (2020), "VR-based investigation of forklift operator situation awareness for preventing collision accidents," *Accident Analysis & Prevention*, Vol. 136, Article 105404. DOI: 10.1016/j.aap.2019.105404.

Garcia-Carrillo, D., Garcia, R., Pañeda, X.G., Mourao, F., Melendi, D., Corcoba, V. and Paiva, S. (2024), "Testing driver warning systems for off-road industrial vehicles using a cyber-physical simulator," *Journal on Multimodal User Interfaces*, Vol. 18, pp. 179–194. DOI: 10.1007/s12193-024-00435-y.

Kang, Y., Zhou, X., Chen, W. and Li, X. (2023), "Investigating the relationship between eye movements and the situation awareness of forklift operators for accident prevention," *International Journal of Occupational Safety and Ergonomics*, Vol. 29, No. 4, pp. 1477–1485. DOI: 10.1080/10803548.2022.2157113.

Kang, Y., Liu, F., Chen, W., Li, X., Tao, Y. and Huang, W. (2024), "Recognizing situation awareness of forklift operators based on eye-movement and EEG features," *International Journal of Industrial Ergonomics*, Vol. 100, Article 103552. DOI: 10.1016/j.ergon.2024.103552.

Lang, A. (2018), "Evaluation of an intelligent collision warning system for forklift truck drivers in industry," in Duffy, V.G. (Ed.), *Digital Human Modeling: Applications in Health, Safety, Ergonomics, and Risk Management*, Lecture Notes in Computer Science, Vol. 10917, Springer, Cham, pp. 610–622. DOI: 10.1007/978-3-319-91397-1_50.

Lee, D.N. (1976), "A theory of visual control of braking based on information about time-to-collision," *Perception*, Vol. 5, No. 4, pp. 437–459. DOI: 10.1068/p050437.

Li, X., Kang, Y., Chen, W., Liu, F., Jiao, Y. and Luo, Y. (2024), "Recognizing the situation awareness of forklift operators based on EEG techniques in a field experiment," *Frontiers in Neuroscience*, Vol. 18, Article 1323190. DOI: 10.3389/fnins.2024.1323190.

Neruda, M., Friel, L., Mandelaire, T., Vojtech, L. and Unucka, J. (2025) 'Data fusion of BLE and 77 GHz radar for forklift-worker collision avoidance', in *2025 IEEE World Forum on Public Safety Technology (WF-PST)*, pp. 54–59. doi: 10.1109/WF-PST65083.2025.00015.

Ray, S.J. and Teizer, J. (2012), "Coarse head pose estimation of construction equipment operators to formulate dynamic blind spots," *Advanced Engineering Informatics*, Vol. 26, No. 1, pp. 117–130. DOI: 10.1016/j.aei.2011.09.005.

Shajahan, T.V., Srinivasan, B. and Srinivasan, R. (2024), "Real-Time Fatigue Monitoring System in Diverse Driving Scenarios," in 2024 IEEE Space, Aerospace and Defence Conference (SPACE), Bangalore, India, pp. 124–127. DOI: 10.1109/SPACE63117.2024.10668230.

Wei, P., Cagle, L., Reza, T., Ball, J. and Gafford, J. (2018), "LiDAR and camera detection fusion in a real-time industrial multi-sensor collision avoidance system," *Electronics*, Vol. 7, No. 6, Article 84. DOI: 10.3390/electronics7060084.