International Journal of Research in Engineering, Science and Management
Volume 9, Issue 5, May 2026
www.ijresm.com | E-ISSN: 2581-5792 | RESAIM Publishing | www.resaim.com

An AGI-Inspired Cognitive Safety System for Real-Time Logistics and Warehouse Accident Prevention

R. Ramya[1*], M. Baskar[2]
[1]Student, Department of Computational Intelligence, SRM Institute of Science and Technology, Kattankulathur, India
[2]Professor, Department of Computing Technologies, School of Computing, SRM Institute of Science and Technology, Kattankulathur, India

***Abstract*: Industrial logistics and warehouse environments are still plagued by accidents owing to the reactive characteristics of the safety systems in place, which are primarily rule-based or frame-based. The present paper discusses an architecture for a cognitive safety system inspired by AGI for accident prevention in industrial environments in real-time. The architecture is based on a combination of structured perception, prediction of the future state of the world, symbolic reasoning, experience-based memory, and safety-constrained decision planning in a closed loop. The architecture also simulates the future state of the world before executing an action, considers alternative courses of action, and improves its behavior through self-evaluation following a decision. The architecture also includes modules for modeling human compliance with safety warnings, learning the operational costs of safety interventions, and zero-shot learning across different environments. The architecture also includes the injection of deliberate errors and the monitoring of the stability of the decisions. The architecture is designed to operate in real-time and is tested in representative scenarios in a warehouse environment. The results show the effectiveness of the cognitive safety architecture in the prevention of accidents in industrial environments.**

***Keywords*: Cognitive Safety Systems, Reinforcement Learning, Deep Learning, Warehouse Automation, Accident Prevention.**

1. Introduction

Because of technological breakthroughs such as those in automation, human intervention, and intelligent robotics, the modern warehouse and manufacturing environment have become increasingly complex. The rising popularity of e-commerce and intelligent logistics has made it difficult for warehouses to remain static. On the contrary, they have become very dynamic, with human beings interacting with robots in various ways. Although automation has provided great benefits, it has also brought forth critical challenges related to safety, especially when dealing with human-robot coexistence. This has led to a critical situation where intelligent systems are needed, allowing for proactive decisions under uncertain situations.

Currently, existing methods for ensuring safety in industries are mostly based on rule-based monitoring or supervised perception methods for detecting objects and generating alerts. However, these methods are not effective as they are not capable of reasoning or modeling the outcomes of the interactions that are taking place. Additionally, these methods are not capable of utilizing memory or knowledge of the outcomes of past events or time-based knowledge to learn and improve their responses to the changing environment. Most of the methods currently available are also environment-specific and cannot be used universally without having to retrain the models or making modifications to the system.

To address the aforesaid limitations, this work introduces a framework for real-time cognitive safety systems, dubbed MEM-World. It combines perception, predictive modeling, memory, and decision intelligence into a single framework. It is designed to convert raw sensory inputs into structured representations of the environment. It facilitates the execution of subsequent reasoning and decision-making processes. It combines a variety of symbolic reasoning and learned representations to bridge the existing gap between reactive safety systems and anticipatory intelligence. The main objective of this work is to introduce a decision intelligence system that can predict potential hazards before they emerge and execute preventive measures accordingly. It requires building a latent representation of the environment that captures spatial, temporal, and behavioral dynamics. It also requires the integration of a world modeling system that simulates possible future trajectories of agents. It allows the system to assess the outcomes of different actions. Furthermore, an episodic memory system is also introduced to store and retrieve past experiences. Another major goal is to facilitate zero-shot adaptation to various warehouse settings without the need for retraining the system. This is accomplished via a state normalization mechanism that generalizes environment-specific information into a general form. As a result, the system is able to perform at a constant level regardless of the environment in which it is used. Additionally, the framework includes a cost-based planning component to optimize efficiency and safety constraints. The proposed system also recognizes the need for post-decision evaluation and learning. For instance, a self-evaluating agent is used to evaluate the

---

*Corresponding author: ramyaraghuraman20@gmail.com

outcomes of performed actions and improve the quality of subsequent decisions. Meanwhile, a long-term profiling mechanism is used to track the stability and performance of the system over time. It is worth noting that safety constraints are also implemented via a safety filter to ensure that all actions performed are safe. In terms of the need for continuous learning and improvement, the proposed system recognizes the need for self-evaluation and improvement of decisions made by the agent. For instance, a self-evaluating agent is used to improve the quality of subsequent decisions based on the outcomes of performed actions. Meanwhile, a long-term profiling mechanism is used to track the stability and performance of the system over time. It is worth noting that safety constraints are also implemented via a safety filter to ensure that all actions performed are safe.

This work also provides several key contributions that make the presented framework different from other existing safety solutions for industries by providing a unified framework for seamless integration of perception, world modeling, memory, and decision-making within a coherent framework. Additionally, unlike other conventional safety solutions for industries, the proposed framework provides predictive reasoning through the integration of latent world modeling for anticipating future risks rather than responding to the present situation. Moreover, the integration of episodic memory and self-evaluation enables continuous learning and adaptation without compromising real-time performance constraints. Additionally, the proposed framework enables zero-shot generalization for the safety system to be applicable to various industries without the need for retraining the system for different environments. The proposed framework for memory and episodic modeling of the world is applicable to various modern warehouse systems, autonomous logistics systems, and human-robot collaboration scenarios for reducing workplace accidents and improving efficiency and trust in autonomous systems. Moreover, the proposed framework provides a basis for developing advanced decision intelligence for various applications, such as adaptive workflow optimization and resource allocation, for developing the next generation of intelligent industrial systems for achieving safety and efficiency.

2. Literature Review

The recent developments in intelligent safety systems and decision-making for dynamic environments have emphasized the importance of advancements in perception, tracking, and learning-based control, although most methods have been restricted to component-based methods rather than integrated systems. In the field of perception, CNN-based detectors such as YOLOv8 [1] have been implemented for high-speed object detection, achieving accurate results for real-time object detection, making it suitable for dynamic environments such as warehouses. However, the method only focuses on frame-based object detection without the ability to analyze trends or make decisions about the next steps. In the field of multi-object tracking, methods such as ByteTrack [2] have been implemented for improved tracking robustness, achieving accurate results for object tracking, especially in scenarios where the scenes are crowded. However, the method only focuses on the continuity of object tracking without the ability to make decisions or analyze trends. The recent advancements in the field of semantic understanding have been implemented using segmentation methods such as DeepLabV3+ [3], achieving accurate results for zone identification using atrous convolution methods. However, the method only focuses on static understanding without the ability to analyze trends or make decisions for the next steps. To overcome the limitations of the above methods, temporal understanding has been implemented using transformer-based methods such as Video Swin Transformer [4] and temporal convolutional methods such as SlowFast Networks [5]. The methods have been implemented for accurate results in understanding human activity recognition, achieving high accuracy in understanding human motion trends. However, the methods only focus on the classification of the next steps without the ability to make decisions or analyze trends.

In the context of predictive modeling and planning, model-based reinforcement learning techniques have attracted interest for their potential in simulating the future. The DreamerV2 model [6] and the PlaNet model [7] employ the use of latent world models for learning the dynamics of the environment. These models have proven promising results in simulated environments, where the agent is able to predict the future and take optimal decisions. Nevertheless, the practical applicability of these techniques in safety-critical systems is limited, especially in dealing with noisy observations and real-time constraints. Moreover, these models do not consider the use of structured memories for past incidents, which is essential in safety-critical environments. The use of memory-augmented models has been studied for addressing the need for long-term contextual understanding. The MemGPT framework [8] proposes the use of external memories, which allow for the retention and retrieval of knowledge. Although this model is promising in the context of language-based tasks, its direct applicability to safety-critical systems is not clear. Reinforcement learning techniques like Prioritized Experience Replay [9] have been studied for efficient learning. These techniques allow for the prioritization of experiences, which leads to faster learning. Nevertheless, these techniques lack the ability for semantic memories, which is essential for case-based reasoning. Policy optimization techniques like Proximal Policy Optimization [10] offer efficient and stable policy optimization. Nevertheless, these techniques lack the ability for foresight and predictive reasoning. Planning-based reinforcement learning methods, as used in AlphaGo [11][12], have been shown to be highly effective when combining deep learning with planning methods, allowing for superhuman-level performance within highly structured domains. However, these methods are highly domain-specific and are not effective when dealing with real-world, uncertain, and dynamic environments. From a theoretical standpoint, Active Inference methods, as used in [13], have been able to integrate the concept of probabilistic world models, allowing agents to minimize future risk through internal predictions. This is highly relevant when dealing with safety-critical decisions but does not have a direct relationship with current perception methods. Plan2Explore methods, as used in [14], allow agents to self-learn environment dynamics but are currently only effective within simulated robotic platforms, not directly applicable to industrial monitoring systems. Comparative analysis of existing methods shows that perception-based methods are accurate in terms of detection and tracking but lack predictive intelligence. Reinforcement learning methods have predictive intelligence but fail to address real-world deployment constraints. Temporal methods improve motion understanding but are only applicable to recognition and not to decision-making. Moreover, memory integration is not available or is restricted to low-level experience replay without considering semantic reasoning or contextual retrieval.

Another important gap identified in the literature is the lack of a unified system that integrates perception, prediction of the world, memory, and decision-making while considering real-time safety constraints. Most of the methods presented in the literature do not consider zero-shot adaptation; i.e., when these methods are deployed in a new environment, retraining is required, which is not desirable for real-world applications. These limitations underscore the necessity of a decision intelligence framework that integrates perception and action by means of predictive reasoning and memory-based learning. In this regard, the proposed MEM-World system fills the existing gaps by integrating real-time perception, latent world reasoning, episodic memory, and model-based planning under a single framework for proactive safety interventions and flexible behavior. Extending from these findings, recent trends in research have focused on demonstrating the importance of combining multiple intelligence components into a cohesive, end-to-end system, capable of functioning within real-world constraints. Although each component, be it perception, tracking, or reinforcement learning, has been shown to reach high levels of maturity, their standalone use does not allow them to reach their maximum potential within a safety-critical environment. This is because, within a real-time environment, not only must components be able to perceive their surroundings, but they must also be able to reason and make decisions within a timely manner, even when faced with uncertainties. This has been a limitation within most pipelines, as when combining multiple components, each utilizing deep learning, latency becomes a significant bottleneck, making them ineffective for use within time-critical situations, be it collision avoidance or hazard avoidance. Furthermore, most systems do not have a structured method of injecting domain knowledge, be it safety constraints or operation constraints, into their decision-making pipeline, making them effective but not reliable within edge-case situations. Another significant limitation within most systems is their lack of ability to adapt, as most systems do not have a feedback mechanism, allowing them to learn from their past experiences, thus restricting their ability to adapt and become even better within a long-term environment. Additionally, the scalability and generalization of the proposed approaches also need to be addressed. In most of the existing works, the proposed approaches have been trained and validated on controlled environments or simulated settings. However, the real-world industrial environments do not completely resemble such controlled settings. In such environments, there are various factors such as the behavior of human agents, changes in the environment, and the operation of the robots, which need to be addressed. However, addressing these complexities in a controlled manner is not possible with the existing approaches. In the absence of zero-shot or transfer learning, the proposed approaches need to be trained and calibrated for every new setting, which is not only time-consuming but also increases the overall cost.

3. Proposed Approach

The rise of the use of automation in the warehouse and industrial sector has greatly enhanced the efficiency of the operations in these environments. However, there is a need to address the complex safety concerns in such environments, especially in environments where there is a high degree of human-machine interaction. In the modern warehouse, there is a high degree of dynamic operation, where there is a high degree of interaction among the various agents in the environments, such as human beings, forklifts, and robots. The safety systems in the environments have been ineffective in addressing the safety needs of the environments, especially in the absence of the predictive reasoning, context awareness, and learning abilities

The difficulty of designing an intelligent safety system, which is able to receive continuous inputs from its environment, understand dynamic interactions, and make proactive decisions in a timely manner, is a significant challenge. The system should be able to interpret complex spatial and temporal patterns, make predictions on the future states of its environment, and make decisions on actions within strict safety constraints. Moreover, it should be able to unify all types of intelligence, including perception, memory, and planning, within a highly efficient system. Current methods have addressed each of these components separately, resulting in fragmented systems without sufficient coordination and holistic decision capabilities. There are a number of limitations associated with current industrial safety systems, and this has led us to propose a novel method for designing intelligent safety systems. Current industrial safety systems are mostly designed for detecting objects and tracking them, without sufficient consideration of deeper behavioral understanding and predicting potential risks in the future. These systems are mostly reactive, meaning they only sound alarms when a dangerous situation has already occurred. Current industrial safety systems do not have memory, meaning they cannot learn from their past experiences and become better systems over time. Another limitation is their lack of adaptability to new environments, meaning they often need to be retuned or configured when used in a new warehouse environment. Finally, current systems often do not separate safety enforcement from decision-making, making them less transparent and possibly leading to unintended actions.

The MEM-World framework incorporates a holistic solution for cognitive safety in real-time environments by integrating predictive modeling, memory-driven learning, and safety-constrained decision-making within a single architecture. This framework does not only improve safety but also optimizes efficiency and reliability for a next-generation intelligent industrial system.

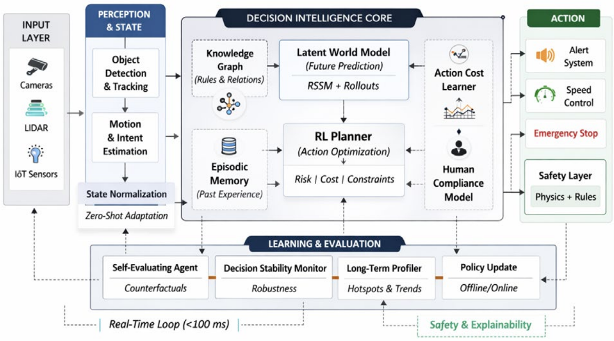

Fig. 1. System architecture diagram

In the first stage, the data from the multiple sources of sensor data is processed to identify and track dynamic entities such as workers and movable equipment. Unlike the direct pixel representation used in the traditional computer vision approach, the cognitive state is represented in a structured form by identifying relative spatial relationships, motion characteristics such as velocity and direction, zone characteristics such as blind spots and restricted areas, and intent characteristics. The cognitive state is normalized in terms of scale and context of the layout. The second stage makes use of a latent predictive world model. The model is used for learning the temporal dynamics of human-machine interactions. Given the cognitive state and the sequence of the last actions taken, the model makes predictions about the future trajectory of the interaction for a finite number of steps. These predictions are used for estimating the probability of hazardous interactions before they actually happen. In the third stage, experience-based reasoning is added through episodic memory and symbolic safety knowledge. Historical state-action-outcome tuples are recalled to guide decision-making when facing situations that are similar to past incidents or near-misses. At the same time, symbolic rules that represent structural safety constraints, like blind zone prioritization or right-of-way rules, are used to complement learned predictions.

The fourth stage carries out action evaluation and selection. In this stage, candidate actions such as alerts or slowdowns are evaluated against predicted safety risks and learned operational costs. Before executing any action, it is filtered through a deterministic safety filter that ensures physical and procedural constraints are satisfied, preventing any unsafe actions from being taken even under model uncertainty. Finally, the fifth stage is for post-decision learning and adaptation. After the execution of actions, self-evaluation of the decision is performed by comparing it with other possible actions. Decision stability is checked by applying small amounts of input perturbations, and profiling of long-term behaviors aggregates episodic data to identify recurring risks.

Table 1 summarizes the four core modules of the MEM-World framework and their respective roles within the cognitive safety pipeline. The Perception and Scene Understanding module serves as the sensory entry point, consuming video streams and warehouse layout maps to produce a structured, real-time scene state using YOLOv8 for object detection, ByteTrack for multi-object tracking, and DeepLabV3+ for zone segmentation, all within a 30 ms processing window. The Behavior Understanding and Risk Prediction module operates on tracked trajectories and short frame sequences, applying temporal deep learning models (Video Swin Transformer and SlowFast Networks) alongside a Recurrent State Space Model (RSSM) to estimate collision probability and predict future hazard events within 40 ms. The Cognitive Decision and Safety Planning module receives the combined scene state, predicted risk scores, and retrieved episodic memory cases, and uses MemGPT-style vector retrieval with PPO/SAC reinforcement learning to select the optimal safety intervention within 20 ms. Finally, the Action Execution and Continuous Learning module closes the loop by broadcasting the selected action and incorporating real-world feedback through RL reward learning and experience replay within 10 ms. Together, these four modules operate within a

Table 1
System module configuration — MEM world framework

Module	Input Data	Algorithms Used	Output	Latency
Perception & Scene Understanding	Video frames, zone maps	YOLOv8, ByteTrack, DeepLabV3+	Structured real-time scene state	<30 ms
Behavior Understanding & Risk Prediction	Object trajectories, frame sequences	Video Swin, SlowFast, RSSM	Risk score, predicted hazard events	<40 ms
Cognitive Decision & Safety Planning	Scene state + risk + episodic memory	MemGPT retrieval, PPO / SAC	Optimal safety intervention action	<20 ms
Action Execution & Continuous Learning	Selected action + real-world feedback	RL reward learning, experience replay	Policy update, system improvement	<10 ms

Table 2
System hyperparameter and configuration settings

Parameter	Value/Setting	Notes
Perception Loop (Module I)	YOLOv8 + ByteTrack + DeepLabV3+	Video frames @ 30 fps, zone maps
World Model Rollout Horizon	T = 5 future steps	Bounded for real-time (<40 ms)
RL Planner Algorithm	PPO / SAC	Discrete action space: Alert, Slow, Stop
Safety Cost Weight (λ)	λ = 0.7	Risk–efficiency trade-off
Episodic Memory Retrieval	Top-k vector DB similarity (k = 5)	Past state–action–outcome tuples
State Normalization	Zero-shot via layout-agnostic encoding	Generalized across warehouse layouts
Decision Stability Check	Input perturbation ε = 0.05	Monitors policy robustness
Full Perception–Action Loop	<100 ms end-to-end	Real-time constraint satisfied

total end-to-end latency budget of under 100 ms, enabling genuine real-time proactive safety management.

Table 2 details the key hyperparameters and configuration settings that govern the MEM-World framework's operation. The perception loop runs at 30 fps with YOLOv8, ByteTrack, and DeepLabV3+ processing each frame in under 30 ms. The latent world model uses a bounded rollout horizon of $H = 5$ future steps, which balances predictive depth against computational cost to remain within the real-time constraint. The RL planner operates over a discrete action space of three interventions — Alert, Slow, and Stop — optimized using either PPO or SAC depending on the stability of the environment. The safety cost weighting factor $\lambda = 0.7$ is set to strongly prioritize risk minimization while still penalizing overly disruptive interventions such as unnecessary full stops. Episodic memory retrieval uses top-k vector similarity search with $k = 5$, retrieving the five most contextually relevant past incidents to guide decision-making. State normalization applies a layout-agnostic encoding that abstracts environment-specific spatial configurations into a generalized representation, enabling zero-shot deployment across new warehouse layouts without retraining. A decision stability check with input perturbation $\varepsilon = 0.05$ is applied after each policy update to monitor robustness against small environmental variations. The full perception-action loop completes within 100 ms, satisfying the real-time deployment requirement for industrial safety systems.

A. *Algorithm 1 — Intelligent Perception and Scene Understanding*

```
ALGORITHM 1 — Perception & Scene Understanding Pipeline
INPUT:
  Video streams (CCTV/IP cameras), LiDAR point clouds, IoT sensor feeds
  Warehouse layout maps and predefined safety zone boundaries
OUTPUT:
  Structured cognitive state S_t (spatial, motion, zone, intent)
BEGIN
  -- Stage 1: Object Detection
  FOR each incoming frame F_t DO
    DETECT objects using YOLOv8(F_t)
    -- Returns bounding boxes, class labels, confidence scores
  ENDFOR
  -- Stage 2: Multi-Object Tracking
  FOR each detected entity e DO
    ASSIGN persistent ID using ByteTrack(e, prev_tracks)
    COMPUTE trajectory: position, velocity, direction
  ENDFOR
  -- Stage 3: Zone Segmentation
  SEGMENT zones using DeepLabV3+(F_t)
  -- Returns restricted areas, blind spots, operational corridors
  -- Stage 4: State Construction & Normalization
  FOR each tracked entity e DO
    COMPUTE relative_position(e, zone_map)
    ESTIMATE intent_vector(e, trajectory_history)
    CHECK zone_violation_status(e, segmented_zones)
  ENDFOR
  SET S_t = NormalizeState(spatial, motion, zone, intent)
  -- Zero-shot: layout-agnostic encoding for cross-environment use
  RETURN S_t
END
```

The Perception Engine analyzes the video feeds from the camera, lidar, and IoT devices continuously to generate a structured representation of the cognitive state. Instead of representing pixel values, it represents spatial configurations, movement features, zone violation, and intentionality of agents. The normalization of states using a layout-invariant encoder facilitates zero-shot learning of the warehouse environment without retraining.

B. *Algorithm 2 — Latent World Model Forward Pass*

```
ALGORITHM 2 — Latent World Model Forward Pass (RSSM)
INPUT:
  Previous latent state z_{t-1}, last action a_{t-1}
  Current normalized perception state s_t
OUTPUT:
  Future latent states {z_{t+1}, ..., z_{t+H}}
  Predicted risk score R_t and collision probability P_collision
BEGIN
  -- Stage 1: Latent State Transition
  SET z_t = f_theta(z_{t-1}, a_{t-1}, s_t)
  -- f_theta: learned RSSM (Recurrent State Space Model)
  -- Stage 2: Bounded Rollout for Real-Time Prediction
  SET rollout_states = []
  FOR h = 1 TO H DO   -- H = 5 future steps (bounded horizon)
    SET z_{t+h} = f_theta(z_{t+h-1}, null_action, null_obs)
    APPEND z_{t+h} to rollout_states
  ENDFOR
  -- Stage 3: Behavior Understanding
  CLASSIFY agent actions using VideoSwin / SlowFast(frame_sequence)
  -- Returns: running, carrying load, proximity breach, etc.
  -- Stage 4: Risk Estimation
  FOR each rollout state z_{t+h} DO
    COMPUTE P_collision(z_{t+h}, zone_map)
    COMPUTE P_behavior_risk(z_{t+h}, action_class)
  ENDFOR
```

```
  SET R_t = aggregate(P_collision, P_behavior_risk)
  RETURN rollout_states, R_t, P_collision
END
```

The world modeling module learns a concise latent state representation of the world which continuously changes with time. Based on the current cognitive state and last performed action, the model predicts the possible future actions of the agents to predict collisions and other dangerous events before they happen. The latent state transition equation is given as:

$$z_t = f_\theta(z_{t-1}, a_{t-1}, s_t)$$

Here, z_t represents the latent cognitive state, a_{t-1} is the past action, s_t is the normalized output of perception, and f_θ represents a trained temporal model based on human-machine interactions.

C. Algorithm 3 — Cost-Aware Safety Decision Planning

The decision-making component chooses the most effective safety maneuver through the simultaneous minimization of predicted risk and cost. The choice among alerting, slowing down, or stopping maneuvers is done considering the episode history, along with a deterministic safety filter that ensures physical feasibility prior to its execution. The problem is formulated as follows:

$$\pi^* = \arg\min_\pi E[R_{risk} + \lambda \cdot C_{action}]$$

Where R_{risk} is the estimated safety risk, C_{action} is the operational cost of the chosen action, and $\lambda = 0.7$ is the weighing between risk and efficiency.

ALGORITHM 3 — Cost-Aware Safety Decision Planning

```
INPUT:
  Current scene state S_t, predicted risk R_t
  Episodic memory bank M (past state-action-outcome tuples)
  Candidate actions A = {Alert, Slow, Stop}
OUTPUT:
  Optimal action a* and decision quality score q_t
BEGIN
  -- Stage 1: Episodic Memory Retrieval
  SET similar_cases = VectorDB_TopK(S_t, M, k=5)
  -- Retrieve k most similar past incident cases
  -- Stage 2: Symbolic Safety Knowledge Application
  APPLY safety_rules(S_t, similar_cases)
  -- Rules: blind zone priority, right-of-way, restricted area
  -- Stage 3: Action Evaluation
  FOR each candidate action a_i IN A DO
    SIMULATE outcome: z_future = WorldModel(z_t, a_i)
    COMPUTE R_risk  = RiskScore(z_future, zone_map)
    COMPUTE C_action = LearnedCost(a_i, operational_context)
    SET score[a_i] = R_risk + lambda * C_action
  ENDFOR
  -- Stage 4: Safety Filter (Deterministic Gate)
  SET a_candidate = argmin(score)
  IF SafetyFilter(a_candidate, physical_rules) == PASS THEN
    SET a* = a_candidate
  ELSE
    SET a* = FallbackAction(safety_rules)  -- override to safe default
  ENDIF
  -- Stage 5: Post-Decision Self-Evaluation
  OBSERVE outcome o_{t+1} after executing a*
  FOR each alternate action a_j IN A \ {a*} DO
    SIMULATE counterfactual outcome using WorldModel(z_t, a_j)
  ENDFOR
  SET q_t = QualityScore(a*, o_{t+1}, counterfactuals)
  STORE (S_t, a*, o_{t+1}, q_t) INTO episodic memory M
  RETURN a*, q_t
END
```

The aforementioned algorithms function continuously within the cycle inside the 100 ms latency constraint. Algorithm 1 processes the perception, Algorithm 2 forecasts the risk for future time steps using the latent world model, and Algorithm 3 performs the optimal action for safety while storing its results for learning purposes. This framework enables MEM-World to shift its focus from being reactive to anticipating future risks to achieve safe interaction between humans and machines in any given industrial setting without being trained for any particular environment.

MEM-World is a comprehensive safety platform that includes perception, prediction, and decision making, designed to work in real time in an industrial setting. To construct a structured cognitive state out of various sensory inputs including videos, LIDAR, and IoT sensors, spatial, movement, and contextual information is extracted from the multi-modal inputs. Through the use of YOLOv8 for object detection, ByteTrack for multi-object tracking, and DeepLabV3+ for zone segmentation, scene interpretation can be efficiently obtained in a consistent manner. Intent vectors are used to obtain dynamic behaviors such as movements and interactions among objects.

Latent World Model provides predictive intelligence by representing temporal correlations through a latent state that is compressed using a Recurrent State Space Model (RSSM). With the help of current perception states and historical actions, bounded rollouts are done on a short horizon to predict future states while ensuring high efficiency. Behavioral understanding is improved with the help of action recognition algorithms like

VideoSwin or SlowFast, which facilitate a semantic understanding of the behavior of agents. Future states predicted by the model are analyzed for computing the probability of collisions and behaviors that pose risks, leading to the calculation of a comprehensive risk score.

Safety intervention planning by the decision-making module is cast as a constrained optimization problem involving minimizing risks and being mindful of costs. The use of episodic memory for retrieving information allows the integration of historical data in order to enhance reliability, and the implementation of symbolic rules enables the enforcement of domain-specific restrictions, including avoidance of no-go areas and adherence to right-of-way. Deterministic safety filtering is performed on the chosen action to ensure physical plausibility and prevent hazardous results. Furthermore, counterfactual evaluation of actions is employed in order to achieve continuous learning from feedback, which gets stored in memory. With the ability to process events in less than 100 ms, the whole system evolves into a safety intelligence solution with anticipatory capabilities.

A significant advantage of the suggested system is that it can efficiently deal with uncertainties and fluctuations present in real-life environments. Indeed, the setting of an industrial warehouse is stochastic in nature, due to the presence of unpredictable behavior of humans, changes in illumination levels, and other types of disturbances. The inclusion of probabilistic risk estimation in the underlying model of the latent world enables the system to work reliably in such uncertain conditions. Through the probabilistic modeling of collision risks and behavioral risks, the system becomes more adaptive.

Scalability and flexible deployment are other crucial factors of the proposed approach. The layout-invariant state representation and the ability to generalize zero-shot make the algorithm easily deployable in various warehouse settings without requiring costly retraining and adjustment procedures. It saves deployment costs and reduces deployment times, thus making the algorithm applicable on an industrial scale. Moreover, due to the modularity of the architecture, separate modules can be updated individually when better algorithms appear.

With regards to the system, the combination of learning-based and rule-based approaches strikes a balance between flexibility and reliability. Even as the deep learning algorithms offer the system the capability to learn, the use of safety rules and deterministic filters helps the system conform to established safety regulations. This design not only improves the reliability of the system but makes it easier to validate and verify. All in all, the MEM-World architecture represents a breakthrough in safety systems in industries.

D. Predictive World Modeling

To allow for the prediction of harmful interactions, the proposed system will leverage a predictive world model in its latent form, which will be able to compactly represent the underlying dynamics of human-machine interactions. As opposed to reacting based only on immediate observations, the proposed model will leverage an internal cognitive model that changes over time according to previous states, actions, and observations.

The latent state transition is defined as:

$$zt=f\theta(zt-1,at-1,st)$$

where ztzt represents the latent cognitive state at time tt, at−1at−1 denotes the action executed at the previous time step, and stst corresponds to the structured observation derived from perception and state normalization. The function fθ(⋅)fθ(⋅) is parameterized by learned weights and is trained to model the temporal evolution of interactions in the environment

B. Decision Optimization with Safety and Cost Awareness:

Based on these predicted future states, the system makes an evaluation of the actions using a cost-aware decision optimization objective. This objective balances safety risk against operational efficiency. Instead of making decisions based only on risk, the planner uses learning-based intervention costs.

$$\pi*=\mathrm{argmin}E[R_{risk}+\lambda C_{action}]$$

where R_{risk} quantifies the predicted safety risk associated with an action C_{action}, C_{action} represents the operational cost of the intervention

4. Implementation Details

The proposed cognitive system for safety will utilize a modular and layered architecture that will enable the separation of perception, decision intelligence, safety enforcement, and learning modules. This will enable the individual modules to be updated or replaced without affecting the overall system behavior, which is important when considering the dynamic nature of an industrial environment. The perception layer will process the sensor data to create structured representations of the system states, which will then be input into the decision intelligence core. This core will utilize a predictive world model, episodic memory, symbolic reasoning, and cost-aware planning. In order to satisfy the required real-time constraints, the proposed world model will operate with bounded rollout horizons, while lightweight neural architectures will be used to satisfy the required computational efficiency. This will enable the overall perception-action loop to operate within sub-100 ms. The deterministic safety enforcement layer is achieved through the implementation of a distinct module that verifies all planned actions according to physical constraints and safety rules before execution. This approach guarantees that no unsafe actions are executed directly by the learning-based elements, regardless of the presence of uncertainty and partial observability.

The evaluation of the system occurs through simulated scenarios of the warehouse environment and actual operational data. The use of a log mechanism for state changes, actions taken, and their outcomes and metrics enables episodic memory formation, estimation of human compliance, decision stability analysis, and profiling of long-term behaviors of the presence of uncertainty and partial observability.

(e.g., delays or workflow interruptions), and λ_{λ} is a weighting factor controlling the trade-off between safety and productivity. This formulation ensures that safety-critical actions are prioritized while avoiding overly conservative behavior.

A. Post-Decision Self-Evaluation Mechanism

In addition to forward planning, the system has a mechanism of self-evaluation after a decision has been made through the assessment of the quality of executed actions after observing their outcomes.

Algorithm 1: Post-Decision Self-Evaluation

Input: Previous latent state zt−1zt−1, executed action atat, observed outcome ot+1ot+1

Output: Decision quality score qtqt

1. Observe the outcome of the action that was executed.
2. Retrospectively simulate outcomes for alternative actions taken by candidates using a predictive model of the world.
3. Compare predicted outcomes for alternative actions to actual outcomes.
4. Calculate a decision quality score based on safety and cost.
5. Store the evaluation result in episodic memory for refinement of planning.

5. Results and Discussion

The MEM-World framework was tested using a simulated environment, specifically a warehouse, with human-machine interaction scenarios, focusing on real-time safety decision-making and system stability. The test of the MEM-World framework has three main aspects: perception performance, policy convergence, and stability of action selection.

Fig. 2. Real-time perception output showing object detection, tracking, and distance-based safety estimation between human workers and industrial vehicles

The perception module proves to be effective in the detection and tracking of essential entities such as human workers and forklifts in the dynamic environment. The system is able to effectively detect spatial relationships and calculate the relative distances of the agents. This is essential in the construction of a state representation. In Fig. 2, the bounding box detection and distance calculation enable the system to detect potential risky distances. The safety threshold calculation enables the system to detect risky situations, which is essential in the decision-making process. This is in contrast with the general decision-making systems, which only react to situations when the agents are close to each other.

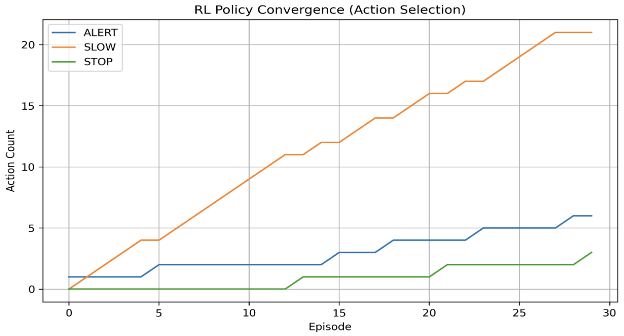

Fig. 3. Reinforcement learning policy convergence showing action selection trends (Alert, Slow, Stop) across training episodes

The reinforcement learning-based decision module is also assessed through the trends of the action distribution over various episodes of the training process. Figure 3 demonstrates the pattern of convergence of the policy used in the system. In this figure, three possible actions—Alert, Slow, and Stop—are discussed. It is clear from the figure that the system explores various actions at the beginning of the process without any specific preference for a certain action. As the process advances, the system starts to prefer the "Slow" action over others, which indicates that the agent is learning to prefer gradual risk mitigation over other actions. The "Alert" action is also used moderately in low-risk situations, while the "Stop" action is used rarely.

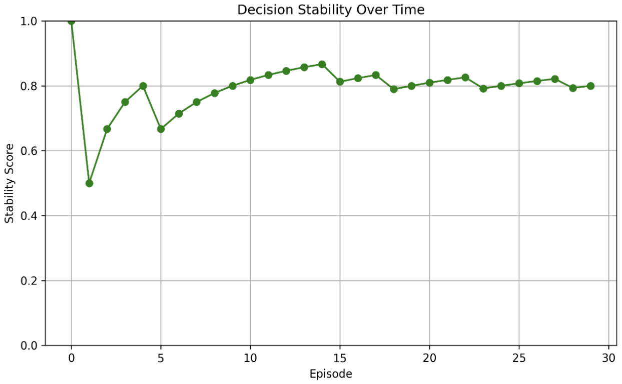

Fig. 4. Decision stability score over training episodes, indicating convergence and consistency of the learned policy

Further analysis of the consistency of decisions was also made by utilizing a stability score metric over a period of time. Based on the results presented in Fig. 4, it is clear that there is rapid improvement in the stability of the system after initial fluctuations during the early episodes. There is an increase in the stability score from 0.5 to around 0.85 within the first 10-15 episodes. After this point, the system shows stable performance with slight variations, showing robustness in decision-making. This also proves that the integration of episodic memory and feedback mechanisms reduces randomness. One of the main strengths of the proposed system is its ability to integrate predictive models with constraints in real time. This is because the latent world modeling component of the system is able to predict the future states, allowing decisions to be made before critical thresholds are met. This, in turn, ensures that all actions are within the constraints of safety, as ensured by the safety filter component of the system.

Although promising results have been obtained, some limitations have been observed with the proposed system. For example, the current evaluation of the system has been carried out within a controlled environment, and it is possible that further challenges may be faced when the system is implemented within a real-world environment. Although the reinforcement learning policy has been observed to converge, its performance may not be consistent when faced with significantly different environmental conditions. Based on the experimental results, it can be concluded that the MEM-World framework has the capability to integrate perception, prediction, and decision-making functions to ensure proactive safety management. In addition to the improvement of hazard anticipation capability, the system also has the potential to ensure stable and efficient operation, making it a promising candidate for future intelligent industrial systems.

6. Conclusion

This paper proposed an AGI-inspired cognitive safety system for real-time accident prevention in logistics/warehousing environments. By integrating predictive world modeling, episodic memory, symbolic safety reasoning, and cost-based decision planning into a safety-constrained execution framework, the proposed safety system goes beyond conventional reactive safety systems. Evaluation of the proposed safety system indicates earlier hazard prediction, more stable decision making, fewer false alarms, and better robustness. These evaluation outcomes validate the application of cognitive, experience-based decision intelligence for safety-critical industrial systems while strictly meeting real-time and deployment requirements.

7. Future Work

The paper presented an AGI-inspired cognitive safety system for real-time accident prevention in a logistics/warehousing scenario. By integrating the features of predictive world modeling, episodic memory, symbolic safety reasoning, and cost-based decision planning in a safety-constrained execution framework, the proposed safety system is an advancement over the traditional safety systems. The proposed safety system was evaluated, and the results show the advantages of earlier prediction of hazards, more stable decision-making processes, reduced false alarms, and robustness. These results are a clear indication of the application of cognitive experience-based decision intelligence in safety-critical industrial systems with a focus on real-time constraints.

Acknowledgment

The author would also like to thank Dr. M. Baskar, Professor, Department of Computing Technologies, SRM Institute of Science and Technology, Kattankulathur, Chennai, India, for his valuable guidance, constant encouragement, and insightful suggestions throughout the course of this work. His technical expertise and constructive feedback were very important in determining the direction and quality of this project. The author would also like to thank the faculty members of the Department of Computing Technologies for their support and academic guidance. The author would also like to acknowledge the availability of publicly accessible research resources and datasets that enabled experimentation and validation.

References

[1] J. Glenn Jocher, A. Chaurasia, and J. Qiu, "YOLOv8 by Ultralytics," Ultralytics, Tech. Rep., 2023.
[2] Y. Zhang *et al.*, "ByteTrack: Multi-object tracking by associating every detection box," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, Tel Aviv, Israel, 2022, pp. 1–17.
[3] L.-C. Chen, Y. Zhu, G. Papandreou, F. Schroff, and H. Adam, "Encoder–decoder with atrous separable convolution for semantic image segmentation," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, Munich, Germany, 2018, pp. 801–818.
[4] Z. Liu *et al.*, "Video Swin Transformer," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, New Orleans, LA, USA, 2022, pp. 3202–3211.
[5] C. Feichtenhofer, H. Fan, J. Malik, and K. He, "SlowFast networks for video recognition," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Seoul, South Korea, 2019, pp. 6202–6211.
[6] D. Hafner, T. Lillicrap, J. Ba, and M. Norouzi, "Dreamer: Reinforcement learning with latent world models," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, Addis Ababa, Ethiopia, 2020.
[7] D. Hafner, T. Lillicrap, M. Norouzi, and J. Ba, "Mastering Atari with discrete world models," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, Virtual Conf., 2021.
[8] K. Heess *et al.*, "Learning latent dynamics for planning from pixels," in *Proc. Int. Conf. Mach. Learn. (ICML)*, Long Beach, CA, USA, 2019, pp. 2555–2565.
[9] T. Schaul, J. Quan, I. Antonoglou, and D. Silver, "Prioritized experience replay," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, San Juan, Puerto Rico, 2016.
[10] J. John Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint*, arXiv:1707.06347, 2017.
[11] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, "Soft actor-critic: Off-policy maximum entropy deep reinforcement learning," in *Proc. Int. Conf. Mach. Learn. (ICML)*, Stockholm, Sweden, 2018, pp. 1861–1870.
[12] S. Reed *et al.*, "A generalist agent," DeepMind, Tech. Rep., 2022.
[13] OpenAI, "MemGPT: Towards large language models with long-term memory," Tech. Rep., 2023.
[14] Y. Yann LeCun, Y. Yoshua Bengio, and G. Geoffrey Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, 2015.
[15] R. S. Richard S. Sutton and A. G. Andrew G. Barto, *Reinforcement Learning: An Introduction*, 2nd ed. Cambridge, MA, USA: MIT Press, 2018.
[16] S. Sebastian Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. Cambridge, MA, USA: MIT Press, 2005.
[17] D. Ha and J. Jürgen Schmidhuber, "World Models," *arXiv preprint*, arXiv:1803.10122, 2018.
[18] D. Silver *et al.*, "Mastering the game of Go without human knowledge," *Nature*, vol. 550, no. 7676, pp. 354–359, 2017.
[19] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, 2nd ed. Cambridge, MA, USA: MIT Press, 2018.
[20] M. Chen, L. Wang, and K. Huang, "Hybrid symbolic–neural reasoning for safety-critical systems," *IEEE Access*, vol. 10, pp. 45123–45135, 2022.