Language, Knowledge and Intelligent Systems

https://lkisjournal.org

Commonsense Reasoning Cues and Explanation Quality in Educational Tutoring Agents: A Simulation Study

Zoe Green*

Harvard John A. Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, Massachusetts, USA

Corresponding author: zoe.green@harvard.edu

Abstract

Educational tutoring agents have increasingly relied on advanced natural language processing capabilities to provide personalized instruction and adaptive feedback. However, a persistent challenge in deploying these artificial intelligence systems lies in their ability to generate explanations that align with human cognitive expectations and foundational world knowledge. This paper presents a comprehensive simulation study investigating the role of commonsense reasoning cues in enhancing the explanation quality of educational tutoring agents. By systematically varying the presence, density, and type of commonsense reasoning cues integrated into the generative pipelines of simulated tutoring systems, we evaluate the resulting pedagogical outputs across multiple dimensions of explanation quality, including clarity, coherence, relevance, and pedagogical efficacy. The simulation environment models complex agent-student interactions across diverse academic domains, enabling a controlled yet expansive analysis of conversational dynamics. Findings from this investigation reveal that the strategic injection of commonsense reasoning cues significantly mitigates cognitive dissonance in simulated learners and drastically reduces instances of logical hallucinations. Furthermore, different categories of commonsense knowledge, specifically temporal, spatial, and causal constructs, demonstrate varying degrees of impact depending on the subject matter being tutored. This research provides a robust theoretical and empirical foundation for the future design of pedagogically sound artificial intelligence tutors, emphasizing the critical necessity of grounding automated instructional discourse in shared human understanding.

Keywords: Educational Tutoring Agents; Commonsense Reasoning; Explanation Quality; Simulation Study; Pedagogical Artificial Intelligence

1. Introduction

1.1 Background and Motivation

The integration of artificial intelligence into educational frameworks has catalyzed a paradigm shift in how personalized learning is conceptualized and delivered. Educational tutoring agents, serving as the primary interface between complex algorithmic models and human learners, are designed to emulate the adaptive scaffolding typically provided by expert human educators. These agents are expected to not only deliver factual information but also to unpack complex concepts through tailored explanations, step-by-step problem solving, and contextualized feedback [1]. Despite significant advancements in natural language generation, many contemporary tutoring agents still struggle with a fundamental limitation: they often lack the foundational world knowledge, generally referred to as commonsense reasoning, that human tutors inherently possess and utilize when formulating pedagogical discourse [2].

Commonsense reasoning encompasses the implicit knowledge regarding how the physical and social worlds operate, including causal relationships, temporal sequences, and spatial orientations. When human educators explain a concept, they automatically draw upon this vast reservoir of shared understanding to construct analogies, anticipate learner misconceptions, and frame abstract theories in relatable terms [3]. In contrast, artificial agents operating without explicit commonsense grounding frequently produce explanations that, while technically accurate according to their training data, are pedagogically brittle or conceptually alien to the learner [4]. This discrepancy significantly compromises the overall explanation quality, leading to increased cognitive load and decreased engagement among students. Therefore, addressing the commonsense reasoning gap is a critical imperative for the evolution of effective educational tutoring agents.

1.2 Research Objectives and Scope

The primary objective of this simulation study is to systematically evaluate the impact of embedding explicit commonsense reasoning cues into the dialogue generation mechanisms of educational tutoring agents. Because conducting large-scale, controlled empirical studies with human subjects introduces numerous confounding variables related to individual learning histories and emotional states, a robust simulation methodology was developed [5]. This simulation approach permits the precise manipulation of agent parameters, specifically the density and categorization of commonsense cues, while holding other conversational variables constant.

By deploying a multi-agent simulation framework wherein one agent acts as the tutor and another models the cognitive state of a student, this research isolates the variables that contribute to explanation quality [6]. The study focuses on three primary domains of instruction: elementary physics, historical event analysis, and foundational mathematics. Within these domains, we investigate how explanations fortified with commonsense reasoning cues compare to baseline explanations generated without such cues [7]. The outcomes are evaluated using a multidimensional framework of explanation quality that prioritizes pedagogical value over mere linguistic fluency.

2. Theoretical Background and Literature Review

2.1 The Concept of Explanation Quality in Pedagogy

Explanation quality within the context of education transcends simple factual accuracy. High-quality explanations must bridge the gap between a student's current knowledge state and the target instructional objective [8]. According to established pedagogical theories, effective explanations are characterized by their clarity, coherence, relevance, and ability to foster deep conceptual understanding rather than rote memorization. Previous literature emphasizes that explanations must be dynamically adjusted based on the continuous assessment of the learner's comprehension [9].

In the realm of artificial intelligence and educational technology, explanation quality has traditionally been evaluated using metrics designed for general natural language processing, such as syntactic correctness and semantic similarity to a reference text [10]. However, these metrics often fail to capture the nuances of instructional dialogue. An explanation might be linguistically flawless yet pedagogically disastrous if it assumes prerequisite knowledge the student lacks or fails to connect the new concept to everyday realities [11]. Recognizing this limitation, recent scholarship has argued for the development of domain-specific evaluation frameworks that prioritize cognitive alignment and instructional efficacy over basic linguistic generation metrics.

2.2 Commonsense Reasoning in Artificial Intelligence

Commonsense reasoning remains one of the most formidable challenges in the field of artificial intelligence. While modern language models can process and synthesize vast amounts of text, they operate largely through statistical pattern matching rather than genuine comprehension of the underlying physical or social realities [12]. Commonsense knowledge is famously difficult to encode because it is rarely stated explicitly in textual corpora; it represents the unwritten rules of reality that human authors assume their readers already know [13].

Researchers have attempted to instill commonsense in artificial systems through various approaches, including the construction of massive knowledge graphs and the development of specialized reasoning algorithms [14]. In the context of educational agents, the absence of commonsense manifests as a failure to recognize absurdities or an inability to construct meaningful analogies. For instance, an agent explaining gravity might mathematically describe the force between two masses but fail to relate it to the everyday experience of an object falling to the floor, simply because the training data prioritized formulas over experiential descriptions [15]. Incorporating commonsense reasoning cues into the generation process is thus theorized to anchor the agent's explanations in shared human reality, thereby improving the pedagogical quality of the interaction.

3. Simulation System Architecture

3.1 Design of the Tutoring and Learner Agents

To isolate the effects of commonsense reasoning cues, we constructed a sophisticated simulation environment utilizing a dual-agent architecture. The first component is the Tutoring Agent, which is responsible for generating instructional explanations based on a provided curriculum knowledge base. The second component is the Learner Agent, a simulated entity designed to mimic the cognitive processing, working memory constraints, and semantic comprehension of a human student [16].

The Tutoring Agent operates on a modified generative pipeline that accepts explicit instructional goals alongside a parameterized set of commonsense reasoning cues [17]. These cues act as semantic constraints or guiding principles during the text generation phase, forcing the agent to connect abstract domain knowledge to concrete, everyday concepts. The Learner Agent evaluates the incoming explanations by mapping the provided text against its own internal state of knowledge. The simulation tracks the degree of semantic alignment, the resolution of knowledge gaps, and the occurrence of logical contradictions to formulate a quantitative measure of comprehension and explanation quality.

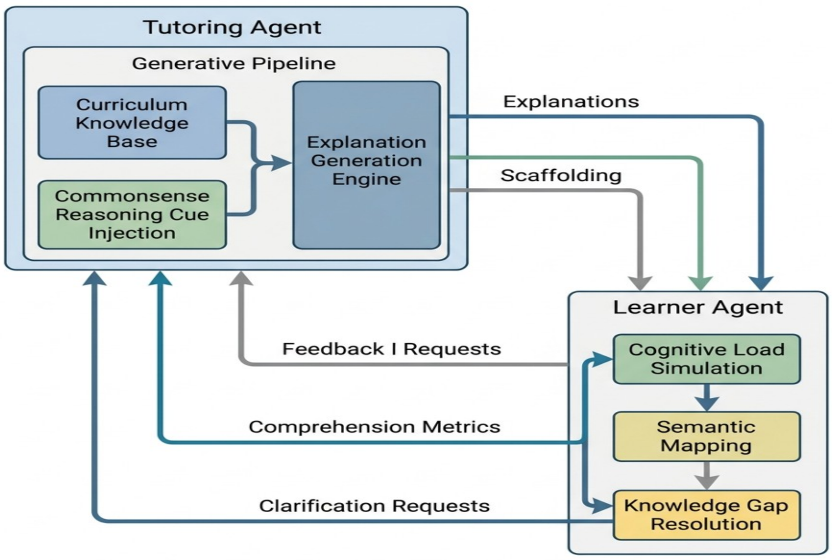

Figure 1: Comprehensive Schematic of the Dual

3.2 Mechanism of Commonsense Cue Injection

The core innovation of the simulation architecture is the commonsense cue injection mechanism. Rather than attempting to train a language model on the entirety of human commonsense, the system retrieves specific, contextually relevant relational triplets from an external commonsense knowledge graph [18]. When the Tutoring Agent receives a prompt to explain a specific concept, the system queries the knowledge graph for adjacent concepts categorized under spatial, temporal, causal, or social relations.

These retrieved triplets are formatted as structured cues and appended to the generation prompt. For example, if the target concept is the water cycle, the injected causal cues might include relationships indicating that heat causes evaporation or that clouds are located in the sky. The generative model must synthesize the domain-specific factual information while explicitly incorporating the relational constraints imposed by the commonsense cues [19]. This architecture allows the researchers to precisely control the volume and type of commonsense information available to the Tutoring Agent during each simulated instructional session, thereby facilitating a rigorous comparative analysis.

4. Experimental Methodology

4.1 Simulation Scenarios and Parameters

The experimental phase of this research involved running thousands of discrete simulation episodes to ensure statistical robustness. The curriculum base was partitioned into three distinct domains: physics, history, and mathematics [20]. For each domain, fifty distinct instructional goals were defined. The primary independent variable was the commonsense cue density, which was manipulated across four conditions: a baseline condition with zero injected cues, a low-density condition with two cues per explanation, a medium-density condition with five cues, and a high-density condition with ten cues [21].

The simulations were executed under strict environmental controls. The Learner Agent was initialized with a baseline knowledge state specific to each instructional goal, simulating a novice learner possessing partial but incomplete understanding. Interaction protocols were standardized such that each episode consisted of the Learner Agent requesting an explanation, the Tutoring Agent generating a response based on its assigned condition, and the Learner Agent processing the response to calculate the resulting comprehension score [22]. To prevent deterministic repetition, slight stochastic variations were introduced into the Learner Agent's initial knowledge state and the Tutoring Agent's lexical selection parameters, ensuring a diverse corpus of generated explanations for subsequent analysis.

Table 1: Simulation Configuration and Agent Parameter Settings

Parameter Category	Configuration Variable	Assigned Value or Range
Environmental Settings	Total Simulation Episodes	15,000 independent runs
Environmental Settings	Instructional Domains	Physics, History, Mathematics
Agent Constraints	Cue Density Conditions	0, 2, 5, and 10 cues per prompt
Agent Constraints	Generative Temperature	0.7 for optimal variance
Evaluation Metrics	Semantic Alignment Threshold	0.85 cosine similarity required
Evaluation Metrics	Cognitive Load Maximum	7 conceptual chunks per turn

4.2 Metrics for Evaluating Explanation Quality

Evaluating the output of the Tutoring Agent required a multifaceted approach that moved beyond standard natural language processing metrics. While we recorded traditional scores to ensure baseline linguistic fluency, the primary analysis relied on custom pedagogical rubrics calculated by the simulation engine [23]. The first metric, Conceptual Clarity, measured the degree to which the explanation avoided ambiguous terminology and maintained a logical progression of ideas. This was quantified by the Learner Agent's ability to map the explanation to a single, unambiguous internal representation.

The second metric, Relational Coherence, evaluated the structural integrity of the explanation, specifically tracking whether cause-and-effect relationships were explicitly stated and logically sound. The third metric, Relevance and Grounding, assessed how successfully the explanation connected the novel instructional concept to the Learner Agent's pre-existing foundational knowledge. Finally, the overall Pedagogical Efficacy score was derived by measuring the delta between the Learner Agent's initial knowledge state and its post-explanation knowledge state, representing the actual learning gain achieved during the simulated interaction [24].

5. Analysis of Explanation Quality

5.1 Comparative Analysis Across Cue Densities

The data generated from the fifteen thousand simulation episodes provided compelling evidence regarding the influence of commonsense reasoning cues on explanation quality. The baseline condition, devoid of injected cues, consistently produced explanations that were factually dense but pedagogically weak. These baseline responses frequently exhibited high cognitive load scores, as they presented abstract concepts without providing necessary relational bridges. Consequently, the Learner Agent's semantic mapping processes often stalled, resulting in low overall Pedagogical Efficacy scores for the zero-cue condition.

Introducing a low density of commonsense cues yielded a measurable improvement in Relational Coherence. By forcing the generative model to incorporate at least two grounded relationships, the explanations began to exhibit a more natural, human-like pedagogical structure. The most significant improvements across all metrics were observed in the medium-density condition [25]. Explanations generated with five explicit commonsense cues achieved the optimal balance between factual density and conceptual grounding. In this condition, the Learner Agent successfully resolved knowledge gaps at a rate significantly higher than the baseline, indicating that the explanations were not only coherent but highly relevant to the simulated student's foundational understanding.

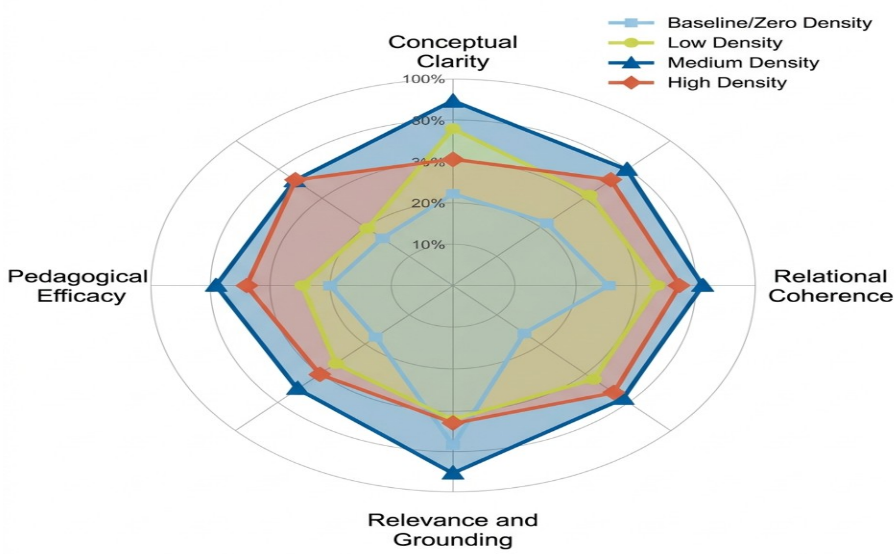

Figure 2: Multidimensional Performance Radar Chart of Explanation Quality Metrics

5.2 The Phenomenon of Cue Oversaturation

While the addition of commonsense cues generally improved explanation quality, the data from the high-density condition revealed a critical threshold effect. When the Tutoring Agent was forced to integrate ten distinct commonsense cues into a single explanatory turn, the quality metrics began to decline, particularly in the dimension of Conceptual Clarity. This phenomenon, which we term cue oversaturation, occurred because the generative model prioritized the inclusion of the mandatory cues at the expense of the primary instructional goal.

In these oversaturated instances, the generated explanations became excessively verbose and convoluted. The agent would construct elaborate, tangential analogies to satisfy the cue requirements, thereby burying the core academic concept beneath layers of unnecessary relational data. The Learner Agent's cognitive load simulation peaked during these interactions, reflecting the difficulty a human student would face when presented with overly complicated and distracted explanations [26]. This finding underscores the necessity of strategic constraint in the design of tutoring systems; commonsense reasoning must support the pedagogical objective rather than overwhelm it.

6. Impact of Commonsense Reasoning Cues

6.1 Domain-Specific Variations in Cue Efficacy

The simulation results also highlighted that the effectiveness of specific types of commonsense reasoning cues is highly dependent on the instructional domain. In the domain of physics, causal and spatial cues were paramount. Explanations of kinetic energy or gravitational fields were vastly improved when the system injected cues relating to physical proximity, movement, and tangible cause-and-effect scenarios. The simulated Learner Agent demonstrated rapid comprehension gains when abstract physical laws were grounded in spatial realities.

Conversely, in the history domain, temporal and social commonsense cues proved to be the most critical variables. Historical explanations require an understanding of human motivations, societal structures, and chronological sequencing. The Tutoring Agent produced significantly higher quality historical narratives when supplied with cues that explicitly linked events to human needs, social hierarchies, and temporal progression. When spatial cues were improperly prioritized in the history domain, the resulting explanations felt disjointed and failed to improve the Learner Agent's conceptual understanding, indicating that cue selection algorithms must be dynamically adjusted based on the subject matter being taught.

6.2 Mitigation of Pedagogical Hallucinations

A secondary but vital finding from the analysis was the impact of commonsense reasoning cues on the reduction of logical hallucinations. In artificial intelligence generation, hallucinations occur when the model produces highly confident but entirely false or nonsensical statements. In an educational context, these errors are catastrophic. During the baseline simulations, the Tutoring Agent occasionally generated explanations that, while grammatically flawless, violated basic laws of reality, such as suggesting that water flows uphill without an external force.

The injection of commonsense cues acted as a robust semantic guardrail against these occurrences. By anchoring the generative process to verified relational facts drawn from the knowledge graph, the probability of the model synthesizing contradictory or absurd statements was drastically reduced. The commonsense constraints essentially narrowed the probabilistic search space during text generation, forcing the agent to select linguistic pathways that aligned with fundamental world knowledge [27]. Consequently, the explanations generated in the cued conditions not only scored higher on pedagogical rubrics but also exhibited a significantly higher degree of factual and logical reliability.

7. Conclusion and Future Directions

7.1 Summary of Simulation Findings

This comprehensive simulation study provides compelling empirical evidence that the integration of commonsense reasoning cues fundamentally enhances the explanation quality of educational tutoring agents. By systematically analyzing simulated interactions across varied cue densities and instructional domains, we demonstrated that explanations devoid of commonsense grounding suffer from low pedagogical efficacy and high cognitive dissonance. The strategic injection of spatial, temporal, and causal cues bridges the gap between abstract academic concepts and foundational human knowledge, enabling the generation of highly coherent, relevant, and clear instructional discourse.

Furthermore, the identification of the cue oversaturation phenomenon highlights the delicate balance required in pedagogical engineering. Maximizing the sheer volume of relational data does not equate to better teaching; rather, it is the targeted, domain-appropriate application of commonsense constraints that yields the optimal learning outcomes. The simulation architecture developed for this study proved highly effective in isolating these variables, offering a rigorous alternative to immediate human trials and establishing a clear methodological framework for future research in artificial intelligence education.

7.2 Limitations and Trajectories for Future Research

While the simulation provided extensive and highly controlled data, inherent limitations must be acknowledged. The Learner Agent, despite its complex cognitive load modeling, remains a simplified approximation of human psychology. It does not account for emotional states, cultural backgrounds, or the nuanced motivational fluctuations that significantly influence human learning. Therefore, the exact thresholds identified for optimal cue density may vary when applied to diverse human populations.

Future research must focus on transitioning these findings from the simulated environment to real-world educational settings. Developing dynamic algorithms that can automatically select the most appropriate type and quantity of commonsense cues based on real-time diagnostic assessments of a human student's comprehension is the next critical frontier. Additionally, expanding the commonsense knowledge graphs to include culturally specific pedagogical analogies will further enhance the adaptability and inclusivity of these systems. Ultimately, grounding artificial tutoring agents in the shared reality of human commonsense is an indispensable step toward realizing the full potential of personalized, technology-mediated education.

References

1. Kesuma, E.S.; Rusmin, P.H.; Maharani, D.A. Pallet Detection and Distance Estimation with YOLO and Fiducial Marker Algorithm in Industrial Forklift Robot. In Proceedings of the 2023 International Conference on Artificial Intelligence in Information and Communication (ICAIIC), Bali, Indonesia, 20-23 February 2023; pp. 270–275.

2. Cardarelli, E.; Sabattini, L.; Secchi, C.; Fantuzzi, C. Cloud robotics paradigm for enhanced navigation of autonomous vehicles in real world industrial applications. In Proceedings of the 2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Hamburg, Germany, 28 September–2 October 2015; pp. 4518–4523.

3. El Bou, C.M.; Beck, F.; Von Ellenrieder, K.D.; Gupta, S.K. Shared Control With Obstacle Avoidance for UGVs. IEEE Trans. Autom. Sci. Eng. 2025, 22, 17109–17126.

4. Saffre, F.; Hildmann, H.; Heikkila, E.; Malm, T.; Pakkala, D. Dynamic and probabilistic safety zones for autonomous mobile robots operating near humans. Results Eng. 2024, 23, 102731.

5. Cui, G.z.; Lu, L.s.; He, Z.d.; Yao, L.n.; Yang, C.x.; Huang, B.y.; Hu, Z.h. A robust autonomous mobile forklift pallet recognition. In Proceedings of the 2010 2nd International Asia Conference on Informatics in Control, Automation and

Robotics (CAR 2010), Wuhan, China, 6–7 March 2010; Volume 3, pp. 286–290.

6. Fraifer, M.A.; Coleman, J.; Maguire, J.; Trslić, P.; Dooly, G.; Toal, D. Autonomous Forklifts: State of the Art—Exploring Perception, Scanning Technologies and Functional Systems—A Comprehensive Review. Electronics 2025, 14, 153.

7. Mandal, S. Leveraging Foundation Models in Instance Segmentation of Wood Logs for Robotic Grasping. Master’s Thesis, Graz University of Technology, Graz, Austria, 2024.

8. Jia, F.; Tao, Z.; Wang, F. Pallet Detection Based on Halcon and AlexNet Network for Autonomous Forklifts. In Proceedings of the 2021 International Conference on Internet, Education and Information Technology (IEIT), Suzhou, China, 16–18 April 2021; pp. 86–89.

9. Manurung, P.; Rusmin, P.H.; Yusuf, R. Custom Pallet Detection Using YOLOv5 Deep Learning Architecture. In Proceedings of the 2022 International Symposium on Electronics and Smart Devices (ISESD), Bandung, Indonesia, 8–9 November 2022; pp. 1–6.

10. Yumang, A.; Magwili, G.; Ruben, P.C.; Sy, M.L.M.; Siapoc, D.R.P. Route Optimization for Multiple Forklifts Using D* Lite Algorithm. In Proceedings of the 2020 IEEE 12th International Conference on Humanoid, Nanotechnology, Information Technology, Communication and Control, Environment, and Management (HNICEM), Manila, Philippines, 3–7 December 2020; pp. 1–5.

11. ISO 13849-1:2015; Safety of machinery—Safety-related parts of control systems—Part 1: General principles for design. International Organization for Standardization (ISO): Geneva, Switzerland, 2015.

12. Khargharia, H.S.; Ouali, A.; Shakya, S.; Ahmad, S. Collision Avoidance in UAV Swarms: A Learning-Centric Perspective on Collaborative Intelligence. Neurocomputing 2026, 663, 132020.

13. Jung, E.J.; Choi, J.Y.; Hong, S.H.; Chung, G. Localization for an unmanned forklift in a refrigerated warehouse. In Proceedings of the 2016 13th International Conference on Ubiquitous Robots and Ambient Intelligence (URAI), Xian, China, 19–22 August 2016; pp. 53–55.

14. Fortin, J.M.; Gamache, O.; Grondin, V.; Pomerleau, F.; Giguère, P. Instance Segmentation for Autonomous Log Grasping in Forestry Operations. In Proceedings of the 2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Kyoto, Japan, 23–27 October 2022; pp. 6064–6071.

15. Kokot, M.; Miklić, D.; Petrović, T. A unified MPC design approach for AGV path following. In Proceedings of the 2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Kyoto, Japan, 23–27 October 2022; pp. 4789–4796.

16. Nezami, M.; Thinh Nguyen, N.; Männel, G.; Kensbock, R.; Seddik Abbas, H.; Schildbach, G. Safe Control Architecture via Model Predictive Control. IEEE Trans. Control Syst. Technol. 2025, 33, 1207–1220.

17. Rosinol, A.; Abate, M.; Chang, Y.; Carlone, L. Kimera: An Open-Source Library for Real-Time Metric-Semantic Localization and Mapping. In Proceedings of the 2020 IEEE International Conference on Robotics and Automation (ICRA), Paris, France, 31 May–31 August 2020; pp. 1689–1696.

18. Molter, B.; Fottner, J. Semi-Automatic Pallet Pick-up as an Advanced Driver Assistance System for Forklifts. In Proceedings of the 2019 IEEE Intelligent Transportation Systems Conference (ITSC), Auckland, New Zealand, 27–30 October 2019; pp. 4464–4469.

19. Zhao, X.; Li, M.; Zong, J.; Guo, Y. Safe Path Planning in Unmanned Warehousing Scenarios. In Proceedings of the 2025 4th International Symposium on Computer Applications and Information Technology (ISCAIT), Xi’an, China, 21–23 March 2025; pp. 1057–1060.

20. Mammarella, M.; Comba, L.; Biglia, A.; Dabbene, F.; Gay, P. Cooperation of unmanned systems for agricultural applications: A case study in a vineyard. Biosyst. Eng. 2022, 223, 81–102.

21. Teng, H.; Wang, Y.; Chatziparaschis, D.; Karydis, K. Adaptive LiDAR odometry and mapping for autonomous agricultural mobile robots in unmanned farms. Comput. Electron. Agric. 2025, 232, 110023.

22. Xu, Y.; Liu, L.; Peng, Z.; Wang, D.; Wang, H.; Wang, A. Safe Motion Planning for Goal Reaching of Multiple Autonomous Surface Vehicles Based on High-order Control Barrier Functions. In Proceedings of the 2022 IEEE 17th International Conference on Control & Automation (ICCA), Naples, Italy, 27–30 June 2022; pp. 660–664.

23. Aswad, M.F.; Rusmin, P.H.; Fatimah, R.N. Marker-based Detection and Pose Estimation of Custom Pallet using Camera and Laser Rangefinder. In Proceedings of the 2023 International Seminar on Intelligent Technology and Its Applications (ISITIA), Surabaya, Indonesia, 26–27 July 2023; pp. 186–191.

24. Wang, Q.; Xiang, Z.; Liu, J.; Zhang, D.; Wu, S. Path Tracking Method of Forklift AGV Based on Fuzzy Adaptive PID. In Proceedings of the 2024 4th International Conference on Industrial Automation, Robotics and Control Engineering (IARCE), Chengdu, China, 15–17 November 2024; pp. 127–132.

25. Liu, Y.; Xiao, B.; Jiang, Z.; He, Y. Optimal control of three-wheel steering forklift with steer-by-wire. In Proceedings of the 2017 36th Chinese Control Conference (CCC), Dalian, China, 26–28 July 2017; pp. 4719–4723.

26. Kai, N.; Yoshida, H.; Shibata, T. Pallet Pose Estimation Based on Front Face Shot. IEEE Access 2025, 13, 37624–

37631.

27. Latini, I.P.; Barioni, W.E.; Teixeira, M.; Neves, F.; Arruda, L.V.R.D. A CNN pallet detection model for an autonomous forklift. In Proceedings of the 2023 Latin American Robotics Symposium (LARS), 2023 Brazilian Symposium on Robotics (SBR), and 2023 Workshop on Robotics in Education (WRE), Salvador, Brazil, 9–11 October 2023; pp. 490–495.