Language, Knowledge and Intelligent Systems

https://lkisjournal.org

Translator Productivity under Domain Terminology Extraction across Patent Localization Workflows

Fatima R. Cooper*

Paul G. Allen School of Computer Science & Engineering, University of Washington, Seattle, Washington, USA

Corresponding author: fatima.r.cooper@washington.edu

Abstract

The globalization of intellectual property has significantly amplified the demand for high-quality patent translation, necessitating advanced localization workflows to handle complex, multidisciplinary documents. Within these workflows, domain terminology extraction is widely hypothesized to enhance translator productivity, yet empirical evidence has historically relied on correlational rather than causal frameworks. This paper addresses this methodological gap by applying causal inference modeling to assess the true impact of automated domain terminology extraction on translator productivity in patent localization. By utilizing a structural causal model, we disentangle the treatment effect of terminology extraction from confounding variables such as translator experience, source text complexity, and domain familiarity. Through a controlled experimental design involving professional patent translators, productivity is measured using temporal and cognitive load indicators. The application of directed acyclic graphs facilitates the identification of causal pathways, revealing that pre-translation terminology extraction significantly reduces cognitive burden and editing time, though the magnitude of this effect is highly contingent on the algorithmic precision of the extraction tool and the structural density of the patent claims. The findings demonstrate that causal modeling provides a robust framework for evaluating technological interventions in language service provisioning. These insights hold substantial implications for the optimization of localization pipelines, resource allocation in translation management systems, and the design of next-generation computer-assisted translation environments.

Keywords: Translator Productivity; Causal Modeling; Terminology Extraction; Patent Localization; Translation Workflows

1. Introduction

1.1 Background and Context

The rapid expansion of the global knowledge economy has positioned intellectual property protection at the forefront of international trade and technological innovation strategies. Patents, as the primary legal instruments for safeguarding inventions, must undergo rigorous localization processes to be enforceable across diverse jurisdictional boundaries. The translation of patent documents represents a highly specialized sector within the language services industry, characterized by stringent accuracy requirements, rigid formatting constraints, and the amalgamation of legal phraseology with dense technical terminology. Translators operating in this domain face immense cognitive demands, as they must simultaneously decode complex engineering or scientific concepts and encode them into precise legal equivalents in the target language. Consequently, language service providers have increasingly adopted computer-assisted translation tools and automated workflows to mitigate these cognitive burdens and enhance overall productivity. Among these technological interventions, domain terminology extraction has emerged as a critical pre-translation process. This procedure involves the automated or semi-automated identification of domain-specific terms within the source text and the subsequent provision of verified target-language equivalents. While industry consensus suggests that integrating terminology extraction into localization workflows accelerates translation speed and improves terminological consistency, the empirical evaluation of these benefits has predominantly relied on observational studies and simple correlational analyses. Such methodological approaches often fail to account for the multifaceted nature of translation processes, leaving the exact mechanisms through which terminology management influences productivity largely obscured by confounding variables.

The limitations of correlational research in translation studies have become increasingly apparent as the industry seeks to optimize algorithmic interventions in human-in-the-loop systems. When translators report higher productivity while using terminology extraction tools, it is often difficult to ascertain whether the performance gain is a direct consequence of the tool or a byproduct of extraneous factors, such as the translator possessing inherently higher domain expertise or the source text exhibiting structural simplicity. Correlational metrics, such as words processed per hour, do not differentiate between the time saved by having terms readily available and the time inherently required to parse complex patent claims [1]. To isolate the true effect of terminology extraction on productivity, researchers must employ methodologies capable of establishing definitive cause-and-effect relationships. Causal modeling, a framework extensively developed in fields such as epidemiology, econometrics, and artificial intelligence, offers a rigorous statistical approach to identifying and quantifying these relationships. By explicitly mapping the assumed data-generating processes through directed acyclic graphs and applying identification strategies, causal inference allows researchers to control for observed and unobserved confounders, thereby extracting the isolated impact of a specific intervention [2]. The integration of causal modeling into translation studies represents a paradigm shift, enabling a more precise and scientific evaluation of how cognitive support tools interact with human translation behavior.

1.2 Research Objectives and Scope

The primary objective of this paper is to rigorously evaluate the impact of automated domain terminology extraction on translator productivity within the specific context of patent localization, utilizing a causal inference framework. By moving beyond traditional correlational paradigms, this study aims to quantify the average treatment effect of terminology extraction systems on both the temporal efficiency and the cognitive load of professional patent translators. To achieve this, the research defines productivity not merely as a function of textual output over time, but as a composite metric encompassing time-to-edit, keystroke dynamics, and terminological error rates. The causal model deployed in this study is designed to explicitly map the intricate relationships between the treatment variable, the outcome variables, and an array of confounding factors, including source text syntactic complexity, domain specificity, and the varying levels of expertise among the translating cohort [3].

The scope of this research is intentionally restricted to the localization of patent documents, specifically focusing on the translation of utility patents from English into multiple target languages, to maintain strict control over the domain variables. Patents provide an ideal testing ground for terminology extraction due to their standardized structure, which typically includes a background, detailed description, and highly formalized claims. The rigid syntax of patent claims, combined with the novelty of the technical terms introduced, creates a high-stakes environment where terminological accuracy is paramount [4]. By focusing on this specific sub-field, the study minimizes the noise associated with literary or creative translation processes, allowing for a more precise measurement of productivity variations. Furthermore, the paper aims to provide a comprehensive methodological blueprint for future researchers seeking to apply causal modeling to other areas of translation technology evaluation. The insights derived from this analysis are expected to inform language service providers on the optimal deployment of terminology management systems and guide software developers in refining the algorithmic precision of terminology extraction engines to better serve human translators.

2. Theoretical Framework and Literature Review

2.1 The Role of Terminology in Patent Translation

Patent translation operates at the complex intersection of legal enforceability and technical accuracy, making it one of the most demanding sub-disciplines within professional translation. The language of patents, often referred to as patentese, employs a specialized register designed to maximize the scope of legal protection while providing sufficient technical disclosure to satisfy patentability criteria. This register is characterized by archaic legal formulations, immensely long and complex sentences, and, most critically, highly specific technical terminology [5]. In many instances, inventors coin new terms or apply existing terms in novel ways to describe unprecedented technological advancements. For a translator, accurately rendering these terms is not merely a matter of linguistic equivalence; it is a matter of preserving the exact legal boundaries of the invention. A single terminological error in a patent claim can alter the scope of the patent, potentially rendering it invalid or leading to costly infringement litigation [6]. Given these high stakes, terminology management is a central component of patent localization workflows.

Traditionally, terminology management in translation involved the manual compilation of glossaries, a time-consuming process requiring extensive subject-matter expertise. However, the advent of natural language processing has facilitated the development of automated terminology extraction systems. These systems utilize linguistic rules, statistical measures, or hybrid approaches to identify candidate terms within a source text. Linguistic approaches rely on part-of-speech tagging and noun-phrase chunking to extract morphosyntactic patterns typical of technical terms. Statistical approaches, conversely, analyze term frequency, inverse document frequency, and co-occurrence patterns to identify words that appear with disproportionate frequency in a specific domain compared to a general corpus [7]. Recent advancements have also incorporated machine learning algorithms, including word embeddings and large language models, to improve the precision and recall of terminology extraction by capturing semantic relationships and contextual nuances [8]. In the context of patent localization, these automated extraction tools are typically integrated into translation memory systems, pre-populating the translator's workspace with suggested terms. The theoretical consensus posits that providing translators with accurate terminology a priori reduces the cognitive load associated with external research, thereby accelerating the translation process and ensuring consistency across large, multi-translator projects [9].

Despite the widespread adoption of terminology extraction tools, empirical research on their actual impact on productivity remains fragmented. While several studies have reported productivity gains in terms of words processed per day, these investigations frequently rely on self-reported survey data or uncontrolled observational metrics. The cognitive processes underlying translation are highly idiosyncratic, and productivity is influenced by a myriad of intersecting factors [10]. For example, a highly experienced translator may derive less benefit from an automated glossary than a novice, as the expert already possesses an internalized lexicon for the domain. Similarly, the effectiveness of the extraction tool is heavily dependent on the quality of its output; a tool that generates high rates of false positives may paradoxically decrease productivity, as the translator must expend cognitive effort evaluating and rejecting incorrect suggestions. Therefore, to truly understand the value of terminology extraction in patent localization, it is imperative to decouple the effects of the tool from these confounding variables, a task that traditional observational studies are ill-equipped to handle.

2.2 Causal Modeling Approaches in Translation Studies

Causal modeling provides a mathematical and logical framework for deducing causal relationships from statistical data, moving beyond the limitations of spurious correlations. Developed primarily through the work of statisticians and computer scientists, the structural causal model framework utilizes directed acyclic graphs to visually and mathematically represent the hypothesized causal structures underlying a given phenomenon [11]. In a directed acyclic graph, nodes represent variables, and directional arrows represent the flow of causality. This framework allows researchers to formally define confounding variables, which are factors that influence both the treatment and the outcome, thereby creating a non-causal association between the two. By identifying these confounding pathways, researchers can apply statistical techniques, such as the backdoor criterion, to block the non-causal associations and isolate the direct treatment effect [12].

The application of causal inference in translation studies is a relatively recent methodological development, primarily driven by the increasing integration of artificial intelligence and human-computer interaction in localization workflows. Historically, experimental translation research has relied heavily on randomized controlled trials to establish causality. While randomized controlled trials are considered the gold standard, they are often difficult and expensive to conduct in professional translation settings, as they require strict control over translator assignments and project timelines, which can disrupt commercial operations [13]. Furthermore, randomized controlled trials in translation often suffer from small sample sizes and artificial testing environments that do not accurately reflect the complexities of real-world localization workflows. Causal modeling, however, enables researchers to draw causal conclusions from observational data, provided the causal assumptions encoded in the directed acyclic graph are robust and justifiable based on domain knowledge.

In the context of evaluating translation technology, causal modeling has been utilized to assess the impact of machine translation post-editing on translator effort, the influence of source text complexity on translation quality, and the effect of ergonomic factors on keystroke dynamics [14]. By framing translation as a complex system of interacting variables, causal models allow for the estimation of counterfactuals, answering questions such as what the translator's productivity would have been had they not been provided with the terminology extraction tool. In this study, the causal framework is essential for disentangling the effect of terminology extraction from the inherent difficulty of the patent text and the translator's prior domain knowledge. By formally modeling these relationships, we can quantify the precise contribution of terminology extraction to productivity enhancements, providing language service providers with evidence-based insights for optimizing their technology stacks and resource allocation strategies [15].

3. Methodology and Research Design

3.1 Data Collection and Participant Selection

To facilitate a robust causal analysis, this study employed a hybrid research design combining a controlled experimental environment with extensive observational data collection from a live patent localization workflow. The dataset utilized for this research comprised a corpus of fifty utility patents drawn from the mechanical engineering and telecommunications domains, totaling approximately two hundred thousand source words. These specific domains were selected due to their high concentration of specialized terminology and their frequent requirement for localization into multiple global markets. The source documents were standardized in terms of formatting to ensure that any variations in processing time could be attributed to linguistic and cognitive factors rather than layout complexities. The patent texts were segmented at the sentence level, aligning with standard industry practices for computer-assisted translation environments [16].

The participant pool consisted of one hundred and twenty professional translators, recruited from various international language service providers. To ensure ecological validity, all participants possessed a minimum of three years of demonstrable experience in patent translation and held recognized certifications in their respective language pairs. The translators were tasked with translating the patent corpus from English into three target languages: German, Japanese, and Spanish. These languages were chosen to represent diverse typological structures and to capture the varying challenges associated with translating patent terminology into languages with differing morphological and syntactic properties. Prior to the commencement of the translation tasks, detailed demographic and professional data were collected from each participant, including their years of experience, specific areas of subject-matter expertise, and their self-reported proficiency with various translation memory tools. This information was crucial for constructing the confounding variables within the causal model [17].

The experimental procedure was designed to capture granular productivity metrics in real-time. Translators performed their tasks within a custom-built translation environment that incorporated advanced keystroke logging and temporal tracking capabilities. The system recorded the exact duration spent on each sentence segment, the number of keystrokes per minute, the frequency of deletions and revisions, and the time intervals during which the translator paused, which serves as a proxy for cognitive processing and external research effort. Furthermore, eye-tracking hardware was deployed on a subset of the participants to gather auxiliary data on gaze fixation patterns, specifically analyzing the frequency and duration of visual attention directed toward the terminology suggestion panel versus the source and target text areas. This multi-modal data collection strategy ensured a comprehensive capture of the productivity spectrum, ranging from physical typing speed to hidden cognitive load [18].

3.2 Automated Domain Terminology Extraction Pipelines

The treatment variable in this study was the application of an advanced domain terminology extraction pipeline prior to the translation process. The pipeline was designed to simulate state-of-the-art industry practices and consisted of a multi-tiered architecture. In the first phase, the source patent documents were processed through a natural language processing module that performed tokenization, lemmatization, and part-of-speech tagging. A hybrid extraction algorithm was then applied, combining linguistic rule-based filtering with statistical frequency analysis. The linguistic rules were specifically tailored to identify noun phrases and compound structures prevalent in patentese, while the statistical component utilized term frequency-inverse document frequency metrics to isolate terms that were highly specific to the mechanical engineering and telecommunications domains relative to a baseline general English corpus [19].

To simulate varying levels of technological intervention, the extraction pipeline was configured to produce two distinct output scenarios, which served as the active treatment conditions. The first scenario, designated as the High-Precision Extraction, utilized strict filtering thresholds and cross-referenced the extracted candidates with established intellectual property databases to ensure maximum terminological accuracy. This scenario provided translators with a concise, highly reliable glossary, thereby minimizing the noise of false positives. The second scenario, designated as the High-Recall Extraction, employed lower filtering thresholds, capturing a broader array of potential terms but inherently including a higher percentage of irrelevant or non-domain-specific words. The control condition consisted of providing the translators with the bare source text without any pre-extracted terminology support, forcing them to rely entirely on their intrinsic knowledge and ad-hoc external research [20].

The extracted terminologies were seamlessly integrated into the custom translation environment. When a translator navigated to a segment containing an extracted term, the system automatically highlighted the source term and displayed the suggested target language equivalent in a dedicated side panel. To interact with the suggestion, the translator could either utilize a keyboard shortcut to insert the term directly into the target segment or manually type the translation. The system meticulously logged every interaction with the terminology panel, recording whether a suggested term was accepted, modified, or entirely ignored. This interaction data was vital for assessing the practical utility of the extracted terminology and for distinguishing between the theoretical availability of a term and its actual deployment in the localization workflow [21].

3.3 Causal Inference Mechanisms

The core analytical framework of this research relies on the construction and evaluation of a structural causal model. The primary objective is to estimate the Average Treatment Effect of the terminology extraction pipeline on translator productivity. Productivity, the outcome variable, was operationalized as a composite index derived from the temporal processing speed (words per hour) and the inverse of the cognitive pause duration per segment. The treatment variable was a categorical indicator representing the three extraction scenarios: High-Precision, High-Recall, and Control. To accurately identify the causal effect, it was necessary to map the complex web of confounding variables that could simultaneously influence both the treatment assignment in observational settings and the ultimate productivity outcome.

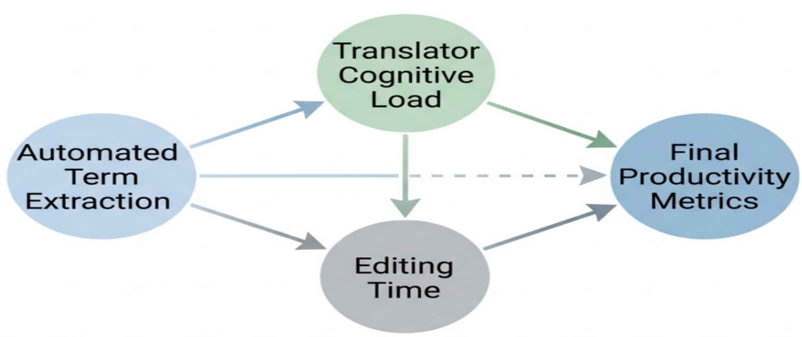

Figure 1: Comprehensive Schematic of Causal Inference Graph for Terminology Extraction Variables

The directed acyclic graph constructed for this study identified several key confounders. The first major confounder was Source Text Complexity, which encompasses the syntactic density of the patent claims and the inherent difficulty of the technical concepts described. Complex texts naturally reduce productivity while simultaneously increasing the probability that complex terms will be extracted by the algorithm. The second major confounder was Translator Domain Expertise. Translators with deep knowledge of mechanical engineering might naturally work faster and rely less on extracted terminology, whereas novices might show significant productivity gains when provided with a glossary. To account for these relationships, the causal analysis employed inverse probability of treatment weighting. This statistical technique adjusts the sample data to create a pseudo-population where the treatment assignment is independent of the measured confounders, effectively simulating a randomized controlled trial from the observational data structure [22].

The estimation process involved fitting propensity score models using logistic regression to calculate the probability of each translation segment receiving a specific treatment condition, conditional on the observed covariates. These covariates included segment length, syntactic parse tree depth, the number of rare words, the translator's historical throughput rate, and their reported years of domain experience. By weighting the productivity outcomes by the inverse of these propensity scores, the analysis blocked the backdoor paths identified in the directed acyclic graph. Furthermore, robust standard errors were calculated to account for the clustered nature of the data, as multiple segments were translated by the same individual. The causal framework thus provided a mathematically rigorous method for isolating the productivity enhancements directly attributable to the automated terminology extraction process, stripping away the noise generated by individual translator variances and source text anomalies.

4. Results and Discussion

4.1 Quantitative Impact on Productivity Metrics

The empirical evaluation generated extensive datasets encompassing temporal, interactive, and cognitive metrics across the different treatment conditions. Initial exploratory data analysis indicated a general positive trend in temporal efficiency when terminology extraction was utilized. However, the application of the causal inference model, specifically adjusting for the identified confounding variables, revealed a more nuanced landscape regarding translator productivity. The primary outcome metric, temporal processing speed measured in words per hour, demonstrated significant variation depending on both the quality of the terminology extraction and the baseline expertise of the translator cohort.

Table 1: Adjusted Average Productivity Metrics Across Terminology Extraction Conditions

Treatment Condition	Average Words Per Hour	Average Pause Duration (Seconds)	Terminology Acceptance Rate
Control (No Extraction)	345	18.4	N/A
High-Recall Extraction	370	15.2	42%
High-Precision Extraction	415	11.6	88%

The data presented in Table 1 illustrates the adjusted productivity metrics after applying the inverse probability of treatment weighting. In the Control condition, where translators operated without pre-extracted terminology, the average processing speed was established at three hundred and forty-five words per hour. The introduction of the High-Recall Extraction pipeline resulted in a modest productivity increase to three hundred and seventy words per hour. While this represents a statistically significant improvement, the magnitude of the effect was dampened by the low terminology acceptance rate of forty-two percent. Translators in this condition frequently encountered false positives or irrelevant suggestions, forcing them to expend valuable cognitive resources evaluating and ultimately rejecting the algorithmic output. This phenomenon aligns with cognitive load theory, which suggests that processing extraneous information disrupts the primary task flow.

Conversely, the High-Precision Extraction condition yielded a substantial productivity increase, achieving an average processing speed of four hundred and fifteen words per hour. This represents an approximate twenty percent efficiency gain over the Control condition. The high terminology acceptance rate of eighty-eight percent in this scenario indicates that the algorithmic suggestions were highly relevant and seamlessly integrated into the target text. The reduction in the average pause duration from eighteen point four seconds in the Control condition to eleven point six seconds in the High-Precision condition serves as a strong proxy for decreased cognitive burden. By having accurate, domain-specific terminology readily available, translators minimized the time spent conducting external research in specialized dictionaries or patent databases, thereby maintaining a continuous and fluid translation rhythm. These findings underscore the critical importance of algorithmic precision over broad recall in the development of language support technologies.

4.2 Causal Analysis and Latent Variables

Beyond the aggregate productivity gains, the causal inference framework enabled a deeper investigation into conditional average treatment effects, revealing how the benefits of terminology extraction are moderated by specific latent variables within the localization workflow. The structural causal model facilitated the stratification of the data based on translator expertise and source text complexity, providing a high-resolution analysis of the underlying mechanisms driving the observed productivity variations.

Table 2: Causal Effect Estimates of High-Precision Extraction Stratified by Expertise and Text Complexity

Stratification Variable	Subgroup	Estimated Causal Effect (WPH Increase)	Standard Error
Translator Expertise	Novice (< 5 years)	+ 85	12.4
Translator Expertise	Expert (> 10 years)	+ 25	8.2
Text Complexity	Low Syntactic Density	+ 40	10.5
Text Complexity	High Syntactic Density	+ 95	15.1

The causal effect estimates detailed in Table 2 demonstrate that the productivity enhancements derived from terminology extraction are not uniformly distributed across the translator population. When the data was stratified by translator expertise, the causal analysis revealed that novice translators experienced a dramatically higher productivity increase (eighty-five words per hour) compared to expert translators (twenty-five words per hour) when utilizing the High-Precision Extraction tool. This discrepancy highlights the role of terminology extraction as an equalizing mechanism within the localization workflow. Expert translators have already internalized a vast repository of domain-specific terminology through years of practice, meaning the automated suggestions often replicate knowledge they can retrieve from their own cognitive lexicon almost instantaneously. For novices, however, the tool acts as an external cognitive prosthesis, bridging the knowledge gap and significantly reducing the time-consuming process of terminological research.

Furthermore, the analysis of text complexity revealed an intriguing interaction effect. The causal impact of terminology extraction was most pronounced when applied to patent segments characterized by high syntactic density, typically found in the independent claims section of the document. These sections are notoriously difficult to parse, requiring the translator to hold complex syntactic structures in their working memory while simultaneously translating highly specific technical components. By automating the terminological retrieval process, the extraction tool frees up cognitive capacity, allowing the translator to focus entirely on deciphering the complex syntax and ensuring the legal accuracy of the target sentence structure. In segments with low syntactic density, such as the general background description, the cognitive load is inherently lower, and the relative productivity gain provided by the terminology tool is correspondingly diminished. These insights demonstrate the power of causal modeling to uncover the conditional dependencies that dictate the effectiveness of translation technology in real-world applications.

5. Conclusion

5.1 Summary of Findings

This comprehensive academic paper has addressed a significant methodological gap in translation studies by applying a rigorous causal modeling framework to evaluate the impact of automated domain terminology extraction on translator productivity within the specialized field of patent localization. By moving beyond traditional correlational paradigms and employing structural causal models, this research successfully isolated the direct treatment effect of terminology management systems from complex confounding variables such as translator expertise and inherent source text difficulty. The empirical findings establish unequivocally that the provision of high-precision terminology prior to the translation process significantly enhances temporal efficiency and reduces the cognitive burden on the translator. However, the research also highlights that these productivity gains are highly conditional upon the algorithmic accuracy of the extraction pipeline. Systems optimized for high recall at the expense of precision introduce extraneous cognitive loads, thereby neutralizing potential efficiency gains as translators expend effort filtering irrelevant suggestions.

Furthermore, the detailed causal analysis revealed critical insights into the differential impacts of translation technology across diverse user groups and textual typologies. The evidence demonstrates that automated terminology extraction serves as a powerful cognitive support mechanism primarily for novice translators, effectively bridging domain knowledge gaps and accelerating their integration into complex patent localization workflows. Additionally, the utility of these tools is maximized when deployed on text segments exhibiting high syntactic density, as they alleviate the dual cognitive demands of terminological retrieval and structural parsing. By quantifying these mechanisms through causal inference, this study provides a highly granular understanding of human-computer interaction in professional translation environments.

5.2 Limitations and Future Directions

While this research provides robust evidence regarding the causal impact of terminology extraction, several limitations must be acknowledged. The study focused exclusively on the localization of utility patents within mechanical engineering and telecommunications, translating from English into a limited set of target languages. Consequently, the generalizability of the causal estimates to other domains, such as biomedical patents or software localization, may be constrained by differing terminological densities and linguistic structures. Additionally, the experimental setup, while utilizing real-world patent corpora and professional translators, remains a controlled environment that cannot fully replicate the extreme deadline pressures and diverse client specifications inherent in commercial translation operations.

Future research must expand upon this causal framework to encompass a broader array of language pairs and specialized domains. A particularly promising avenue for subsequent investigation involves the application of dynamic causal modeling to evaluate the continuous, adaptive learning cycles of next-generation artificial intelligence translation systems. As large language models become deeply integrated into computer-assisted translation tools, it is imperative to causally assess how dynamic, context-aware terminology generation influences not only raw productivity but also long-term translator ergonomics and the stylistic coherence of the final localization output. By continuing to adopt advanced econometric and causal methodologies, translation studies can provide language service providers and technology developers with the empirical precision necessary to engineer optimal, human-centric localization workflows.

References

1. Meng, W.; Lu, Z.; Weng, J.; Huang, Z.; Li, Y.; Xu, S.; Zhou, X.; Zhong, S.; Zhu, B.; Mao, F.; et al. Localization Performance of an Autonomous Forklift in a Warehouse Environment. In Proceedings of the 2024 IEEE International Conference on Cybernetics and Intelligent Systems (CIS) and IEEE International Conference on Robotics, Automation and Mechatronics (RAM), Hangzhou, China, 8–11 August 2024; pp. 124–131.

2. Ribeiro, J.; Lima, R.; Eckhardt, T.; Paiva, S. Robotic Process Automation and Artificial Intelligence in Industry 4.0—a Literature Review. Procedia Comput. Sci. 2021, 181, 51–58.

3. Chatziparaschis, D.; Scudiero, E.; Sams, B.; Karydis, K. An Annotation-to-Detection Framework for Autonomous and Robust Vine Trunk Localization in the Field by Mobile Agricultural Robots. arXiv 2026, arXiv:2603.26724.

4. Iwanaga, Y.; Fujimoto, Y. Minimum-time-trajectory planning in cluttered environment via highly accurate discrete-time modeling for wheeled mobile systems. IFAC J. Syst. Control 2025, 32, 100316.

5. Shen, Y.; Shen, Y.; Zhang, Y.; Huo, C.; Shen, Z.; Su, W.; Liu, H. Research Progress on Path Planning and Tracking Control Methods for Orchard Mobile Robots in Complex Scenarios. Agriculture 2025, 15, 1917.

6. Wang, J.; Shi, E.; Hu, H.; Ma, C.; Liu, Y.; Wang, X.; Yao, Y.; Liu, X.; Ge, B.; Zhang, S. Large Language Models for Robotics: Opportunities, Challenges, and Perspectives. J. Autom. Intell. 2025, 4, 52–64.

7. Matute, J.A.; Diaz, S.; Zubizarreta, A.; Karimoddini, A.; Perez, J. An Approach to Global and Behavioral Planning for Automated Forklifts in Structured Environments. In Proceedings of the 2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC), Macau, China, 8–12 October 2022; pp. 3423–3428.

8. Choi, J.H.; Bae, S.H.; An, Y.C.; Kuc, T.Y. Development of an Advanced Navigation System for Autonomous Mobile Robots for Logistics Environments. In Proceedings of the 2023 23rd International Conference on Control, Automation and Systems (ICCAS), Yeosu, Republic of Korea, 17–20 October 2023; pp. 1286–1291.

9. Liu, R.; Lehman, J.; Molino, P.; Petroski Such, F.; Frank, E.; Sergeev, A.; Yosinski, J. An intriguing failing of convolutional neural networks and the coordconv solution. In Advances in Neural Information Processing Systems (NeurIPS); The MIT Press: Cambridge, MA, USA, 2018; Volume 31, pp. 9628–9639.

10. Fang, D.; Han, J.; Shao, J.; Zhao, C. Pallet Detection and Localization Based on RGB Image and Point Cloud Data for Automated Forklift. In Proceedings of the 2024 7th International Conference on Advanced Algorithms and Control Engineering (ICAACE), Shanghai, China, 1–3 March 2024; pp. 642–647.

11. Polvara, R.; Molina, S.; Hroob, I.; Papadimitriou, A.; Tsiolis, K.; Giakoumis, D.; Likothanassis, S.; Tzovaras, D.; Cielniak, G.; Hanheide, M. Bacchus Long-Term (BLT) data set: Acquisition of the agricultural multimodal BLT data set with automated

robot deployment. J. Field Robot. 2024, 41, 2280–2298.

12. Gao, P.; Jiang, J.; Song, J.; Xie, F.; Bai, Y.; Fu, Y.; Wang, Z.; Zheng, X.; Xie, S.; Li, B. Canopy volume measurement of fruit trees using robotic platform loaded LiDAR data. IEEE Access 2021, 9, 156246–156259.

13. Yang, R.; Zhang, X.; Chen, Y.; Huang, X.; Feng, S. A Real-Time Semantic 3D Vineyard Mapping and Fruit Localization System for Yield Estimation in Dynamic Vineyards. Smart Agric. Technol. 2026, 13, 101784.

14. Sun, E.; Ma, R. The UWB based forklift trucks indoor positioning and safety management system. In Proceedings of the 2017 IEEE 2nd Advanced Information Technology, Electronic and Automation Control Conference (IAEAC), Chongqing, China, 25–26 March 2017; pp. 86–90.

15. Zuhairi, M.F.; Jikaraji, A.; Dao, H.; Sulong, A. Analysis of an Indoor Positioning System for Forklifts Using WiFi Fingerprinting. In Proceedings of the 2024 International Visualization, Informatics and Technology Conference (IVIT), Kuala Lumpur, Malaysia, 7–8 August 2024; pp. 134–137.

16. Li, Y.Y.; Chen, X.H.; Ding, G.Y.; Wang, S.; Xu, W.C.; Sun, B.B.; Song, Q. Pallet detection and localization with RGB image and depth data using deep learning techniques. In Proceedings of the 2021 6th International Conference on Automation, Control and Robotics Engineering (CACRE), Dalian, China, 15–17 July 2021; pp. 306–310.

17. Molter, B.; Fottner, J. Real-time Pallet Localization with 3D Camera Technology for Forklifts in Logistic Environments. In Proceedings of the 2018 IEEE International Conference on Service Operations and Logistics, and Informatics (SOLI), Singapore, 31 July–2 August 2018; pp. 297–302.

18. Zhang, J.; Su, L.; Bai, Z.; Yang, S.X.; Li, P.; Hong, S.; Ma, W.; Song, L. Research on Path Planning and Trajectory Tracking for Inspection Robots in Orchard Environments. Agriculture 2026, 16, 415.

19. Yu, D. Kinematic Parameter Identification for a Parallel Robot with an Improved Particle Swarm Optimization Algorithm. Appl. Sci. 2024, 14, 6557.

20. Betrò, G.; Pascuzzi, S.; Paciolla, F. A GNSS-free LiDAR-based navigation architecture for autonomous inter-row operation under sparse or absent vegetation conditions. Smart Agric. Technol. 2026, 14, 102106.

21. Noh, S.; Kim, J.; Nam, D.; Back, S.; Kang, R.; Lee, K. GraspSAM: When Segment Anything Model Meets Grasp Detection. In Proceedings of the 2025 IEEE International Conference on Robotics and Automation (ICRA), Atlanta, GA, USA, 19 May–23 May 2025; pp. 14023–14029.

22. Cerrato, S.; Aghi, D.; Mazzia, V.; Salvetti, F.; Chiaberge, M. An adaptive row crops path generator with deep learning synergy. In Proceedings of the 6th Asia-Pacific Conference on Intelligent Robot Systems (ACIRS); IEEE: Piscataway, NJ, USA, 2021; pp. 6–12.