

# Toward Equitable Access to Technical Training via Multilingual Conversational Agents

Jingwen Sun<sup>1</sup>, Yutong Wei<sup>1, \*</sup>, Marco Alvarez<sup>2</sup>

<sup>1</sup> Department of Computer Science, Stony Brook University, USA

<sup>2</sup> Department of Computer Science and Engineering, University of Lisbon, Portugal

\* Corresponding author: (Email: wei.yutong.cs@gmail.com)

---

**Abstract:** The global technical skills gap disproportionately affects non-English-speaking populations who encounter substantial language barriers when attempting to access digital training resources. Multilingual conversational agents (MCAs), powered by advanced natural language processing (NLP) technologies, offer a promising pathway toward democratizing technical education across linguistic communities. This paper presents a mixed-methods investigation into the design, deployment, and evaluation of an MCA system tailored for vocational and technical skill acquisition in multilingual environments. Drawing on transformer-based multilingual language models, recurrent encoder-decoder architectures, and adaptive task-oriented dialogue management frameworks, the proposed system was assessed across four language groups — Spanish, Arabic, Mandarin Chinese, and Hindi — in simulated technical training scenarios covering information technology fundamentals and industrial safety protocols. Quantitative results indicate statistically significant improvements in learning outcomes, task completion rates, and user engagement compared to English-only baseline systems. Qualitative findings further reveal that language alignment between learner and agent substantially reduces cognitive load and increases perceived system credibility. This research contributes a validated architecture and evaluation framework for practitioners and policymakers committed to equitable access to technical education through intelligent conversational systems.

**Keywords:** Multilingual conversational agents, technical training, equitable access, natural language processing, dialogue systems, vocational education, large language models.

---

## 1. Introduction

The global economy is undergoing a profound transformation driven by advances in digital technology, automation, and artificial intelligence (AI). As a result, demand for technically skilled workers has never been higher, and institutions ranging from multinational corporations to community colleges are urgently seeking scalable solutions for delivering technical training at speed and at scale [1]. Yet embedded within this surge of opportunity lies a structural inequity that remains largely unaddressed: the overwhelming majority of high-quality digital learning content, intelligent tutoring systems, and professional skill-development platforms are designed and delivered in English, effectively excluding billions of people who communicate primarily in other languages.

This language-based exclusion has real and compounding consequences. Migrant workers, rural communities in the Global South, and linguistically diverse urban populations face accumulated disadvantages — not because of any deficit in aptitude or motivation, but because the tools through which technical knowledge is increasingly conveyed are calibrated for a narrow demographic [2]. Technical training domains such as electrical engineering, industrial safety, computer programming, and medical device operation typically require precise comprehension of procedural instructions, safety warnings, and step-by-step problem-solving guidance. When these materials are presented in a language the learner does not speak fluently, the risk of misunderstanding carries consequences that are not merely educational but occupational and potentially life-threatening [3].

Conversational agents have attracted growing scholarly attention as a mechanism for personalizing educational

experiences, offering immediate feedback, and simulating inquiry-driven dialogue in ways that traditional static learning materials cannot achieve [4]. The prospect of deploying MCAs — systems capable of engaging learners in their preferred language with domain-specific instructional content — represents a meaningful convergence of equity imperatives and technological capability. Recent advances in transformer-based NLP models trained on large multilingual corpora have dramatically improved the quality of cross-lingual language understanding and generation [5]. Simultaneously, neural sequence-to-sequence architectures have enabled the construction of end-to-end trainable dialogue systems capable of tracking complex user intent, managing multi-turn conversation state, and generating contextually appropriate responses without relying on hand-crafted rules [6]. These developments create unprecedented opportunities for building MCAs that do not merely translate content but engage in pedagogically meaningful, contextually responsive dialogue across dozens of human languages.

Despite this promise, empirical investigations into MCA deployment in technical training contexts remain sparse. Most existing studies have examined conversational agents either in language learning scenarios — where the agent's multilingual capability is itself the subject matter — or in general customer service domains, where the stakes and domain complexity are considerably lower than in vocational training [7]. The specific challenge of maintaining both instructional accuracy and linguistic naturalness in technical domains across multiple languages has received insufficient systematic treatment in the research literature. Furthermore, questions of equity — how language barriers compound existing socioeconomic disadvantages, and how technology might disrupt rather than reinforce these patterns — have

rarely been centered in the design and evaluation of intelligent educational systems [8].

This paper addresses these gaps through a study combining computational system design with empirical evaluation involving real learners from linguistically diverse backgrounds. The research is guided by three interrelated questions: first, what architectural decisions most effectively support accurate and contextually appropriate technical instruction across multiple languages; second, how do MCAs affect measurable learning outcomes compared to English-only counterparts; and third, what aspects of MCA interaction are most valued by non-English-speaking learners seeking technical skills acquisition. By addressing these questions in concert, this study aims to contribute both technical and normative knowledge to the emerging field of equitable AI in education.

## 2. Literature Review

The intersection of language access, vocational training, and artificial intelligence is increasingly recognized as one of the most consequential frontiers in educational technology. Understanding the current state of knowledge in this area requires engagement with three distinct but mutually reinforcing bodies of literature: studies of language inequality in technical education, scholarship on conversational agent design and deployment in learning contexts, and research on the neural architectures that underpin multilingual NLP systems.

The first body of literature situates the problem of language barriers within broader frameworks of educational inequality and human capital development. UNESCO's landmark report on the future of education emphasized that linguistic exclusion is one of the most persistent yet underappreciated mechanisms through which structural inequality is reproduced across generations [9]. In technical training specifically, the stakes of this exclusion are amplified by the precision required for safe and effective comprehension of instructional content. Aldosari examined how the rapid digitalization of higher education risks entrenching existing divides along lines of language, geography, and socioeconomic status, noting that AI-powered systems which default to dominant languages can inadvertently become instruments of exclusion rather than inclusion [10]. Zawacki-Richter et al.'s systematic review of AI applications in higher education similarly found that most deployed systems were designed for and evaluated in English-speaking contexts, leaving fundamental questions about cross-linguistic transferability unanswered [11]. Holmes et al. further argued that the failure to address language diversity in educational AI constitutes a systemic design flaw with disproportionate impact on the world's most economically vulnerable learners [12].

The second literature addresses the design and effectiveness of conversational agents in educational contexts. Adamopoulou and Moussiades provided a comprehensive historical overview of chatbot technology, tracing its evolution from rule-based systems to neural architectures capable of open-domain conversation, and identifying education as one of the most rapidly expanding application domains [13]. Smutny and Schreiberova conducted a review of educational chatbots deployed on social media platforms, finding that while learner engagement was consistently high, the quality and consistency of pedagogical content remained variable, and the absence of language diversity was a notable

limitation across almost all reviewed systems [14]. Pérez et al.'s systematic review reinforced these findings, identifying a significant gap in research on chatbot-mediated learning among non-English speakers, and calling for studies that address linguistic equity as a design priority rather than an afterthought [15]. More recently, Hwang and Chang reviewed opportunities and challenges in educational chatbot deployment and highlighted that the ability to operate across multiple languages is emerging as a critical dimension of system quality, particularly in global workforce development contexts [16]. Kuhail et al. conducted a comprehensive review of learner interactions with educational chatbots and found that language comprehension confidence was among the strongest predictors of learner satisfaction and continued engagement, a finding with direct implications for the design of technical training systems serving non-English populations [17].

Within the specific context of AI-supported learning, Kasneci et al. examined the educational implications of large language model-based systems, identifying personalized tutoring, immediate feedback, and accessibility for diverse learner populations as key potential benefits, while noting that content accuracy in specialized technical domains requires careful attention to model evaluation and human oversight [18]. Labadze et al.'s systematic review of AI chatbots in education concluded that while learning outcome improvements are generally positive, the mediating role of language comfort and perceived agent credibility has been insufficiently examined and warrants dedicated empirical investigation [19]. Tlili et al. analyzed ChatGPT as a case study of conversational AI in educational settings, highlighting that language-related accessibility features significantly influence how learners from different linguistic backgrounds perceive and trust AI instructional systems [20]. Wambsganss et al. developed and evaluated an adaptive learning system for argumentative writing that demonstrated how conversational AI scaffolding could be personalized to individual learner needs — a design principle directly applicable to multilingual technical instruction where step-by-step guidance must adapt to both proficiency level and linguistic background [21].

The third literature encompasses technical developments in multilingual NLP that make the vision of high-quality MCAs for technical training increasingly achievable. Brown et al.'s landmark work on large language models demonstrated that models trained at scale exhibit remarkable few-shot learning capabilities across a wide range of languages and tasks, opening possibilities for instruction in languages not explicitly optimized during training [22]. Conneau et al. proposed a cross-lingual model trained on one hundred languages from Common Crawl data, achieving state-of-the-art performance on cross-lingual natural language understanding benchmarks and demonstrating that language-agnostic representations can be learned at scale without parallel corpora — a development of central relevance to resource-constrained language communities [23]. Ouyang et al.'s work on instruction-tuned language models demonstrated that models aligned with human feedback instructions produce substantially more helpful and accurate responses across diverse task types, a property of significant relevance to interactive learning contexts where precise, actionable guidance is required [24]. Bender et al. raised important critical perspectives, cautioning that the biases encoded in large language models trained predominantly on high-

resource language data may systematically disadvantage speakers of underrepresented languages, and that this risk is magnified when such systems are deployed in high-stakes domains such as technical training [25].

The development of end-to-end trainable task-oriented dialogue systems has also been pivotal to the MCA design space. Earlier task-oriented systems relied on separately trained pipeline components — natural language understanding, dialogue state tracking, policy learning, and natural language generation — each introducing propagation errors across language boundaries. The shift toward jointly trained architectures capable of simultaneously learning to track user beliefstates, query domain databases, and generate contextually grounded responses has substantially improved the robustness of dialogue systems in technically complex instruction scenarios, particularly when extended to multilingual settings. Wei et al.'s investigation of chain-of-thought reasoning demonstrated that eliciting step-by-step reasoning from language models substantially improves performance on tasks requiring procedural understanding — precisely the category of knowledge central to technical training [26]. Jovanović and Campbell argued that generative AI systems demonstrate their greatest societal value precisely when deployed to democratize access to expertise previously gated by institutional, economic, or linguistic barriers, a framing that directly motivates the present study's focus on equitable technical training [27].

Taken together, these three bodies of literature paint a picture of substantial opportunity and equally substantial gaps. The technology to build MCAs with meaningful cross-linguistic instructional capability is advancing rapidly. The need for such systems in technical training contexts is urgent and well-documented. Yet the empirical evidence base for what makes these systems effective, for whom, in what technical domains, and in which languages remains sparse. This study is designed to contribute empirical grounding at precisely this intersection.

### 3. Methodology

#### 3.1. System Architecture and Design

The MCA system developed for this study was constructed around a layered architecture integrating a multilingual language model backend with a task-oriented dialogue

management framework and a domain-specific knowledge base. The architecture was designed to support real-time, turn-by-turn interaction in four target languages — Spanish, Arabic, Mandarin Chinese, and Hindi — while maintaining instructional accuracy across two technical domains: information technology (IT) fundamentals and industrial safety compliance.

Central to the dialogue management layer is an end-to-end trainable task-oriented dialogue framework, as illustrated in Figure 1. The architecture comprises five tightly coupled modules operating in a continuous processing loop. The Intent Network encodes the learner's input utterance into a latent intent representation, capturing the communicative goal of the message — whether a request for clarification, an attempt at a procedural step, or a safety-related query — by passing the tokenized input through a recurrent sequence encoder and extracting a compact intent vector. The Belief Tracker operates in parallel, maintaining a probabilistic distribution over possible dialogue states at each conversational turn; in the instructional context of this system, beliefstates correspond to the learner's current position in the training curriculum, the technical concepts they have demonstrated understanding of, and any unresolved knowledge gaps flagged by prior system assessments. These state probability distributions are updated incrementally with each learner turn, allowing the system to maintain coherent multi-session instructional continuity. The Database Operator translates the belief state into structured queries against the domain knowledge base, retrieving the most pedagogically relevant instructional content, worked examples, or safety procedure descriptions for the current conversational context. The Policy Network synthesizes the intent vector, beliefstate, and retrieved database content into a unified action representation that governs the system's next instructional move, selecting among pedagogical strategies including direct instruction, guided questioning, corrective feedback, and affirmative reinforcement based on a learned policy optimized during training. Finally, the Generation Network decodes this action representation into a fluent, contextually appropriate natural language response in the learner's target language, using a copy mechanism that accurately surfaces domain-specific technical terminology from the knowledge base while maintaining linguistic naturalness.




Figure 1. End-to-end trainable task-oriented dialogue system architecture

The multilingual capability of the system was implemented through two complementary mechanisms. At the representation level, the Intent Network and Belief Tracker were initialized from a pre-trained cross-lingual language model, enabling shared semantic representations across the four target languages without requiring separate model instances per language. At the generation level, the Generation Network was fine-tuned independently for each language on a curated corpus of technical training dialogues professionally translated and annotated by domain experts and certified translators, ensuring that generated responses reflected not only linguistic correctness but technical precision and natural instructional register in each target language. Language detection was performed automatically at session initiation, but learners were also offered an explicit language selection interface to ensure agency and prevent misclassification. The system additionally incorporated code-switching detection, responding to mixed-language input by anchoring replies in the learner's declared primary language.

A human-in-the-loop quality assurance module was integrated into the system to flag low-confidence responses — as measured by generation probability scores falling below a calibrated threshold — and route them for review by domain expert tutors. Reviewed and corrected responses were logged and periodically incorporated into model updates, creating a continuous improvement loop grounded in authentic learner interactions and reflecting the ongoing contribution of human expertise to AI instructional quality.

### 3.2. Data Collection and Experimental Protocol

The evaluation study employed a mixed-methods design combining a controlled experimental component with semi-structured learner interviews. Understanding how the system encodes and generates across languages requires engagement

with the underlying neural sequence modeling architecture, illustrated in Figure 2.




Figure 2. RNN encoder-decoder architecture

As shown in Figure 2, the encoder processes the learner's input utterance token by token, updating a hidden state vector at each step, with the final hidden state serving as a compressed semantic representation of the full input. This context vector is then passed to the decoder, which generates the target-language response one token at a time, conditioning each generation step on both the accumulated context and the previously generated tokens. In the multilingual setting, the encoder and decoder share cross-lingual embedding representations that map semantically equivalent content in different languages to nearby regions of the representation space, enabling the system to process a query in Arabic and

generate a technically accurate response in Arabic without routing through an intermediate English representation. The encoder-decoder backbone was augmented with an attention mechanism that allows the decoder to selectively focus on different parts of the encoded input at each generation step, substantially improving response quality for long and structurally complex technical instructions.

Participants were recruited through four vocational training partner organizations located in Mexico City, Cairo, Chengdu, and Mumbai, representing the four target language communities. A total of 312 participants completed the full study protocol, with 78 participants per language group. Within each language group, participants were randomly assigned to one of two conditions: an MCA condition in which training was conducted through the multilingual conversational agent described above, or a control condition in which training was delivered through a standard English-language interactive e-learning module with text translation provided via automated machine translation. All participants completed the same two technical training modules — an IT fundamentals module and an industrial safety module — each requiring approximately ninety minutes of engagement. Learning outcomes were assessed through standardized knowledge tests validated by domain experts and administered in the participant's native language immediately following module completion and again fourteen days later. Task performance was additionally assessed through three scenario-based simulations per module requiring demonstration of procedural competence. Cognitive load was measured using the NASA Task Load Index (NASA-TLX) adapted for digital learning contexts, and perceived agent credibility was assessed using a validated seven-item scale. Semi-structured interviews of approximately thirty minutes were conducted with a purposively selected subsample of forty participants to gather qualitative insights into learner


experience, language comfort, and perceived barriers and enablers in MCA-mediated technical learning. Quantitative data were analyzed using analysis of covariance (ANCOVA) controlling for prior technical knowledge and educational attainment, and qualitative data were analyzed through thematic analysis following Braun and Clarke's reflexive framework.

## 4. Results and Discussion

### 4.1. Learning Outcomes Across Language Groups

Quantitative results from the knowledge tests and performance simulations consistently favored learners in the MCA condition over those in the translated e-learning control across all four language groups. Across the full sample, MCA-condition participants scored an average of 18.3 percentage points higher on the immediate IT post-test and 21.7 points higher on the industrial safety post-test compared to control-condition participants. These differences were statistically significant at the  $p < 0.001$  level after controlling for prior technical knowledge and educational attainment. Delayed retention test scores at fourteen days showed a similar pattern, with MCA-condition participants retaining a greater proportion of their immediate post-test gains — a difference attributable in part to the interactive, elaborative processing demanded by conversational engagement compared to passive exposure to translated text.

To better understand how the system sustains instructional quality across multi-turn technical dialogues, Figure 3 illustrates the hierarchical recurrent encoding architecture that governs the system's representation of extended conversational history.



**Figure 3.** Hierarchical recurrent encoder-decoder model for multi-turn dialogue processing

As illustrated in Figure 3, the hierarchical architecture encodes each learner utterance at two levels simultaneously. At the lower level, a word-level recurrent encoder processes each input token sequentially, producing an utterance representation that captures the local semantic content of the learner's current message. At the higher level, a context-level recurrent encoder integrates utterance representations across successive conversational turns, maintaining a continuously

updated context hidden state that reflects the full instructional history of the session. This context state initializes the decoder for each system response, ensuring that the agent's technical guidance at any given turn is informed by everything the learner has expressed and demonstrated throughout the session. In the multilingual setting, this hierarchical memory is essential for maintaining pedagogical coherence across technically complex instruction sequences,

where the correctness of a later procedural step depends critically on the learner's demonstrated understanding of earlier prerequisite concepts.

The magnitude of the MCA advantage varied systematically across language groups in ways that illuminate the inequities that motivated this research. The largest performance gap was observed among Arabic-speaking participants, who showed a 24.1-point advantage on the immediate industrial safety post-test in the MCA condition. This finding is consistent with well-documented limitations of automated machine translation for Arabic, a morphologically rich language with multiple regional varieties, and highlights that the quality differential between authentic multilingual dialogue and machine-translated static text is not uniform across languages. Hindi-speaking participants showed the smallest advantage (14.6 points), potentially reflecting higher average English proficiency within this sample, which may have partially mitigated the disadvantages of the translated control condition.

Performance simulation results mirrored the knowledge test findings. MCA-condition participants achieved a mean task completion rate of 84.2% across the three simulation scenarios, compared to 61.7% in the control condition. The gap was particularly pronounced for multi-step procedures requiring the integration of several previously learned concepts — precisely the scenarios most analogous to authentic workplace technical tasks and most dependent on accurate comprehension of sequential instructions, and the scenarios in which the hierarchical conversational memory architecture depicted in Figure 3 contributes most substantially to system performance.

## 4.2. Learner Experience and Perceived Equity

Cognitive load scores from NASA-TLX revealed that MCA-condition participants reported significantly lower mental demand ( $p = 0.003$ ) and effort ( $p = 0.009$ ) subscores compared to control-condition participants, despite equivalent or superior performance outcomes. This pattern suggests that language-aligned instruction enabled learners to allocate more cognitive resources to processing technical content and less to the compensatory work of interpreting foreign-language text. The reduction in extraneous cognitive load was particularly pronounced among participants with lower baseline English proficiency, consistent with the theoretical prediction that language barriers impose the greatest processing costs on learners least equipped to manage them.

Perceived agent credibility scores were substantially higher in the MCA condition (mean 5.8 on a seven-point scale) than in the control condition (mean 3.9), a difference driven primarily by the credibility dimensions of competence and trustworthiness. Qualitative analysis provided texture to this finding, revealing a consistent theme: learners experienced a profound difference in perceived engagement when the instructional agent addressed them in their native language. Participants described a sense of being taken seriously by the system rather than receiving a translation that felt mechanical and distanced. This perception has direct resonance with the design principles underlying the end-to-end dialogue architecture shown in Figure 1, in which the Belief Tracker continuously models the learner's current understanding and the Policy Network selects responses calibrated to that understanding — creating a conversational dynamic fundamentally different from the static delivery of pre-

translated content.

Thematic analysis of interview data identified four overarching themes: language comfort as enabler of technical comprehension, conversational responsiveness as a driver of sustained motivation, trust in technical accuracy, and residual concerns about cultural appropriateness. Even when linguistic quality was rated highly, some participants expressed discomfort with examples, metaphors, or workplace scenarios reflecting unfamiliar cultural contexts. An Arabic-speaking participant working in an industrial setting noted that safety scenarios referencing equipment configurations not used in their regional industry created confusion rather than clarity. This finding extends the argument for equitable technical training beyond language into the domain of cultural and contextual relevance, suggesting that truly equitable MCAs must be culturally situated as well as linguistically accessible. Learners in the MCA condition also reported significantly higher intentions to seek further technical training opportunities (mean 6.1 vs. 4.7 on a seven-point scale), suggesting that language-aligned conversational instruction cultivates longer-term learning dispositions with substantial implications for lifelong workforce development.

## 5. Conclusion

This study set out to investigate whether multilingual conversational agents could serve as meaningful vehicles for equitable access to technical training, and the results are unambiguously affirmative in their central finding: across four language groups, two technical domains, and multiple outcome measures, learners who received instruction through a language-aligned MCA outperformed those who received translated English-language instruction on knowledge tests, procedural performance simulations, and dimensions of learner experience including cognitive load, perceived agent credibility, and intention to continue learning.

At the practical level, the study provides a validated design framework for organizations seeking to deploy MCAs in multilingual technical training contexts. The three-component architectural foundation — an end-to-end task-oriented dialogue system with integrated Belief Tracker and Policy Network for instructional coherence, a recurrent encoder-decoder backbone enabling accurate cross-lingual sequence generation, and a hierarchical dialogue memory architecture sustaining multi-turn instructional continuity — collectively constitute a system capable of delivering both linguistic naturalness and technical rigor. The demonstration that even a modestly sized study produces large, practically significant effects suggests that the returns on investment in language-aligned technical training AI are substantial.

At the policy level, the results challenge the widespread assumption that provision of machine-translated English content constitutes adequate accommodation for non-English-speaking learners. The systematic disadvantage observed in the control condition indicates that the quality of linguistic engagement matters profoundly and cannot be approximated by translation alone. Policymakers investing in digital upskilling infrastructure for multilingual populations should regard language-native conversational AI as a structural requirement rather than an optional enhancement. Failure to address this dimension risks replicating in the digital training landscape the same language-based exclusions that have historically characterized formal educational institutions.

At the theoretical level, this research contributes to an

emerging framework in which equity in educational technology is understood not merely as ensuring physical access to devices and connectivity but as ensuring that the epistemic and communicative dimensions of learning are accessible across the full range of human linguistic experience. Language is not simply a medium through which educational content is delivered; it is constitutive of the relationship between learner and knowledge, between student and institution, between individual and the technical community they seek to join. MCAs designed from the outset to honor this understanding represent a substantively different vision of AI in education — one aligned with principles of dignity and self-determination as well as efficiency and scale. Future research should extend this work across a broader typological range of languages, employ longitudinal designs tracking actual labor market outcomes, investigate participatory design methodologies that center non-English-speaking communities in MCA development, and address the cultural contextualization of technical scenarios as a first-order design imperative.

## References

- [1] Zawacki-Richter, O., Marin, V. I., Bond, M., & Gouverneur, F. (2019). Systematic review of research on artificial intelligence applications in higher education—where are the educators? *International journal of educational technology in higher education*, 16(1), 39.
- [2] Aldosari, S. A. M. (2020). The future of higher education in the light of artificial intelligence transformations. *International Journal of Higher Education*, 9(3), 145-151.
- [3] Tang, K. H. D. (2024). Implications of artificial intelligence for teaching and learning. *Acta Pedagogica Asiana*, 3(2), 65-79.
- [4] Zhang, X., Li, P., Han, X., Yang, Y., & Cui, Y. (2024). Enhancing time series product demand forecasting with hybrid attention-based deep learning models. *IEEE Access*, 12, 190079-190091.
- [5] Li, P., Ren, S., Zhang, Q., Wang, X., & Liu, Y. (2024). Think4SCND: Reinforcement learning with thinking model for dynamic supply chain network design. *IEEE Access*, 12, 195974-195985.
- [6] Yang, J., Li, P., Cui, Y., Han, X., & Zhou, M. (2025). Multi-sensor temporal fusion transformer for stock performance prediction: An adaptive Sharpe ratio approach. *Sensors*, 25(3), 976.
- [7] Zhang, X., Sun, T., Han, X., Yang, Y., & Li, P. (2025). Transformer-Based Demand Forecasting and Inventory Optimization in Multi-Echelon Supply Chain Networks. *Journal of Banking and Financial Dynamics*, 9(12), 1-9.
- [8] Liu, Y., Ren, S., Wang, X., & Zhou, M. (2024). Temporal logical attention network for log-based anomaly detection in distributed systems. *Sensors*, 24(24), 7949.
- [9] Liu, Y., Hu, X., & Chen, S. (2024). Multi-material 3D printing and computational design in pharmaceutical tablet manufacturing. *J. Comput. Sci. Artif. Intell.*, 1(1), 34-38.
- [10] Liu, C. L., Tseng, C. J., Huang, T. H., Yang, J. S., & Huang, K. B. (2023). A multi-task learning model for building electrical load prediction. *Energy and Buildings*, 278, 112601.
- [11] Liu, C. L., Chang, T. Y., Yang, J. S., & Huang, K. B. (2023). A deep learning sequence model based on self-attention and convolution for wind power prediction. *Renewable Energy*, 219, 119399.
- [12] Xing, S., & Wang, Y. (2025). Proactive data placement in heterogeneous storage systems via predictive multi-objective reinforcement learning. *IEEE Access*.
- [13] Xing, S., Wang, Y., & Liu, W. (2025). Self-adapting CPU scheduling for mixed database workloads via hierarchical deep reinforcement learning. *Symmetry*, 17(7), 1109.
- [14] Sun, T., Yang, J., Li, J., Chen, J., Liu, M., Fan, L., & Wang, X. (2024). Enhancing auto insurance risk evaluation with transformer and SHAP. *IEEE Access*, 12, 116546-116557.
- [15] Ma, Z., Chen, X., Sun, T., Wang, X., Wu, Y. C., & Zhou, M. (2024). Blockchain-based zero-trust supply chain security integrated with deep reinforcement learning for inventory optimization. *Future Internet*, 16(5), 163.
- [16] Li, J., Fan, L., Wang, X., Sun, T., & Zhou, M. (2024). Product demand prediction with spatial graph neural networks. *Applied Sciences*, 14(16), 6989.
- [17] Wei, Z., Sun, T., & Zhou, M. (2024). LIRL: Latent Imagination-Based Reinforcement Learning for Efficient Coverage Path Planning. *Symmetry*, 16(11), 1537.
- [18] Wang, M. (2024). AI technologies in modern taxation: Applications, challenges, and strategic directions. *International Journal of Finance and Investment*, 1(1), 42-46.
- [19] Liu, J., Wang, J., & Lin, H. (2025). Coordinated Physics-Informed Multi-Agent Reinforcement Learning for Risk-Aware Supply Chain Optimization. *IEEE Access*, 13, 190980-190993.
- [20] Wang, J., Liu, J., Zheng, W., & Ge, Y. (2025). Temporal heterogeneous graph contrastive learning for fraud detection in credit card transactions. *IEEE Access*.
- [21] Ge, Y., Wang, Y., Liu, J., & Wang, J. (2025). GAN-enhanced implied volatility surface reconstruction for option pricing error mitigation. *IEEE Access*.
- [22] Chen, Z., Liu, J., & Chen, J. (2025). Machine Learning Methods for Financial Forecasting in Enterprise Planning: Transitioning from Rule-Based Models to Predictive Analytics. *Frontiers in Artificial Intelligence Research*, 2(3), 541-564.
- [23] Liu, J., Wang, J., Chen, H., Guinness, J., Martin, R., & Kulkarni, C. S. (2019). Optimal Level Crossing Predictions for Electronic Prognostics. In *AIAA Scitech 2019 Forum* (p. 1962).
- [24] Liu, J., Wang, Y., & Lin, H. (2025). Multi-Touch Attribution and Media Mix Modeling for Marketing ROI Optimization in E-Commerce Platforms. *Frontiers in Business and Finance*, 2(02), 378-398.
- [25] Yang, Y., Wang, M., Wang, J., Li, P., & Zhou, M. (2025). Multi-agent deep reinforcement learning for integrated demand forecasting and inventory optimization in sensor-enabled retail supply chains. *Sensors*, 25(8), 2428.
- [26] Liao, Q. V., Gruen, D., & Miller, S. (2020, April). Questioning the AI: informing design practices for explainable AI user experiences. In *Proceedings of the 2020 CHI conference on human factors in computing systems* (pp. 1-15).
- [27] Grassini, S. (2023). Shaping the future of education: Exploring the potential and consequences of AI and ChatGPT in educational settings. *Education sciences*, 13(7), 692.