From Cognition to Simulation: A Systematic Mapping of Visualisation and Interaction in Multi-Agent System

Guilherme Monteiro[1], **Gleifer V. Alves**[1], **André P. Borges**[1], **Rafael C. Cardoso**[2]

Postgraduate Program in Computer Science (PPGCC-PG)
Federal University of Technology – Paraná (UTFPR)
Ponta Grossa – PR – Brazil

Department of Computing Science
University of Aberdeen
Aberdeen – UK

`luipad@alunos.utfpr.edu.br`

`{gleifer, apborges}@utfpr.edu.br`

`rafael.cardoso@abdn.ac.uk`

Abstract. *Multi-Agent Systems have achieved high cognitive complexity, but their experimentation is hindered by architectural limitations and the rigid coupling between cognition and visualisation in current simulation platforms. To investigate this scenario, we present a systematic literature mapping based on an adaptation of PRISMA and Kitchenham's methodologies. The analysis revealed three critical gaps: a lack of standardised integration, a conflict between 3D realism and flexibility, and low usability due to hardcoded programming. These results demonstrate the need to develop decoupled architectures. The delegation of the logical deliberation cycle to cognitive frameworks and volumetric rendering to game engines, interconnecting them via standardised asynchronous communication bridges, emerges as the structural pathway to provide an efficient and scalable visual infrastructure for research in the field.*

1. Introduction

Artificial Intelligence (AI) has evolved into distributed ecosystems, establishing Multi-Agent Systems (MAS) [Wooldridge 2009] as foundations for modelling autonomous entities. To validate the emergent behaviour of these agents, computational simulation acts as an essential empirical laboratory [Abar et al. 2017]. However, historical simulations based on discrete two-dimensional matrices have become insufficient when confronting algorithms that utilise sensors and actuators in continuous environments [Ghallab et al. 2025], demanding platforms with greater physical and spatial realism [Juliani et al. 2020]. Consequently, the scientific community has converged towards the use of Game Engines, which natively provide stable three-dimensional environments and complex interaction logics [Juliani et al. 2020].

The Python ecosystem stands out in the development of AI models and in supporting deliberative architectures integrated with machine learning, such as the *Multi-Agent System for Python* (MASPY) framework [Mellado et al. 2025, Mellado et al. 2026]. However, the language is poorly equipped for three-dimensional rendering [Beeching et al. 2021], creating a bottleneck that forces the delegation of spatial simulation to external platforms. The core of the scientific challenge, therefore, lies in efficiently coupling the agent's abstract mind to its simulated body without introducing high communication latency or massive processing constraints [Beeching et al. 2021, Al Shukairi and Cardoso 2023, Lazarin et al. 2026].

This paper presents a systematic mapping study to diagnose gaps and trends in the literature. The remainder of this work is organised as follows: Section 2 presents the systematic mapping methodology; Section 3 details the results of the review and identifies the main trends and gaps in the field; finally, Section 4 presents the conclusion.

2. Systematic Mapping Methodology

To systematically investigate the state of the art of MAS simulation platforms and virtual environments while ensuring methodological rigour and research reproducibility, this study follows the guidelines of the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA)[1]. The execution of this review was structured into three essential phases suggested by [Kitchenham and Charters 2007] planning, conducting, and reporting to organise the workflow and maintain a definitive focus during the study screening process. To centralise and manage this entire lifecycle, the Parsifal platform[2] was employed, supporting the collaborative management of the protocol from initial search planning to final data extraction.

2.1. Research Questions

To guide this mapping study, six Research Questions (RQs) were formulated and are presented below:

- RQ1 Which agent modelling paradigms are most widely used in simulators?
- RQ2 Which game engines or dedicated simulation platforms are predominantly employed in recent literature?
- RQ3 What are the most common approaches for integrating the agents' reasoning engines with spatial rendering environments?
- RQ4 What level of visual and spatial fidelity is offered by current tools?
- RQ5 Which application domains or practical problems are most frequently addressed in recent literature?
- RQ6 Which interaction paradigms for agent configuration occur most frequently in recent literature?

By answering these six questions, this systematic mapping aims to establish a comprehensive diagnosis of the current technological landscape.

2.2. Search Strategy

The primary search for this study was conducted in the Scopus and Web of Science electronic databases, covering the period from 2020 to 2025. The main search string combined terms related to MAS and Simulation Platforms. The final structured search string applied to the databases was:

[1] `https://www.prisma-statement.org/` (Accessed: 19 June 2026)
[2] `https://parsif.al` (Accessed: 19 June 2026)

```
(Multi-Agent System OR Intelligent Agents OR MAS)
AND (Simulation Platform OR Agent-Based Modelling
OR Simulation Framework OR Simulator OR Testbed
OR Virtual Environment)
```

In addition to this structured search, the Backward Snowballing technique was applied as a complementary search strategy. This stage consisted of "mining" the references of secondary and tertiary studies (such as systematic reviews and surveys) identified in the initial search.

2.3. Study Selection Criteria

To select the most relevant articles and ensure that the mapping focused strictly on simulation tools, the retrieved studies underwent a two-phase screening process in accordance with predefined Inclusion Criteria (IC) and Exclusion Criteria (EC). The adopted ICs were:

- IC1 Studies that propose or evaluate platforms, simulation frameworks, or virtual environments for a single agent or for MAS.
- IC2 Studies that detail the software architecture used for communication and integration between the agents and the simulation environment.

Conversely, the EC applied to eliminate out-of-scope studies were:

- EC1 Studies focusing exclusively on mathematical or theoretical formulations of agents, without the presentation or practical use of an executable simulator.
- EC2 Studies in which the simulation is conducted in a physical medium, i.e., without the use of simulation software.
- EC3 Secondary or tertiary studies, which were set aside to be used exclusively in the snowballing phase.

By applying these criteria, the initial pool of retrieved literature was systematically filtered. This process ensured that only primary studies containing substantial architectural and empirical evidence regarding the coupling of virtual environments and agent reasoning engines advanced to the quality assessment and data extraction phases.

2.4. Quality Assessment

Accepted primary articles underwent a Quality Assessment using five equally weighted Quality Questions (QQs) to evaluate their methodological rigour. Each question was scored using a strict rubric, 1.0 (Full), 0.5 (Partial), or 0.0 (Insufficient), resulting in a maximum total score of 5.0 per article. The formulated QQs were:

- QQ1 Is the platform/tool clearly identified and documented (name, version, link)?
- QQ2 Does the article clearly describe the application domain?
- QQ3 Does the article clearly describe how the integration and communication architecture occur between the agent's logical engine (backend) and the spatial rendering environment (frontend)?
- QQ4 Is the interaction paradigm for configuring agents described (e.g., via code, GUI, XML files)?

QQ5 Is the implementation reproducible (providing a repository link or detailed pseudocode)?

A cut-off threshold of ≥ 3.0 was applied to select articles with high architectural rigour for the final corpus. Papers scoring below this limit were excluded due to insufficient documentation and a lack of explicit evidence. This measure prevented subjective inferences, ensuring a reliable quantitative aggregation to answer the research questions.

3. Mapping Results

The application of the research methodology yielded a final set of studies that were systematically analysed to address the six RQs. This section presents the study selection process flow that led to the final corpus of articles, the extracted data, and the quantitative results for each RQ.

3.1. Study Selection Flow

The study selection process followed a structured flow to ensure the methodological transparency of this screening, as detailed in Figure 1.

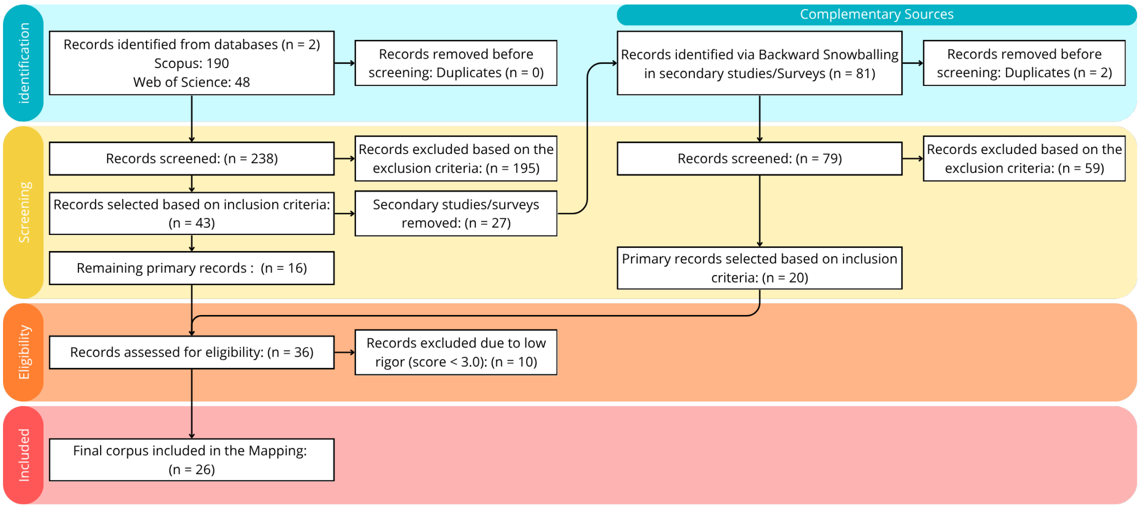

Figure 1. PRISMA-based study selection flow detailing the identification, screening, and eligibility phases.

The initial bibliographic identification retrieved a total of 238 records (190 from Scopus and 48 from Web of Science). Following the application of the predefined inclusion (IC) and exclusion criteria (EC) during the first screening phase, 195 studies were excluded, resulting in 43 selected articles. To isolate the primary architectural evidence, these 43 articles were divided into two distinct groups: 16 primary studies that proceeded directly to the eligibility phase and 27 secondary or tertiary studies (e.g., literature reviews and surveys). These 27 reviews were removed from the primary analysis, but served as a foundation for a backward snowballing procedure.

The snowballing phase identified 81 additional records from the references of the 27 reviews. After the removal of 2 duplicates, 79 records underwent the rigorous screening process. Applying the same ICs and ECs, 59 papers were excluded, resulting in 20 new primary records. A total of 36 primary articles (16 from the initial search and 20 from the snowballing) advanced to the Quality Assessment phase.

Table 1. Grouped Quality Assessment Scores of Primary Articles.

Score	References
5.0	[Chen et al. 2024], [Wang et al. 2023], [Al Shukairi and Cardoso 2023]
4.5	[Du et al. 2024], [Eichelbeck et al. 2026], [Schrittwieser et al. 2020]
4.0	[de Mello et al. 2024], [Rajvanshi et al. 2024], [Angelotti and Díaz-Rodríguez 2023], [Liu et al. 2023], [Silver et al. 2021], [Zhang et al. 2025], [Goel et al. 2024], [Schuck et al. 2025], [Rakhman et al. 2021], [Grosvenor et al. 2025], [Moon 2021]
3.5	[Zellers et al. 2021], [Sovrano et al. 2022], [Zhang et al. 2024], [Sohn et al. 2020], [Yaw et al. 2020], [Cai et al. 2025]
3.0	[Labiosa et al. 2025], [Lorang et al. 2024], [Rehm et al. 2024]
2.5	[Subramanian et al. 2024], [Bahamid et al. 2024], [Dang-Nhu 2020]
2.0	[Deane and Ray 2025], [Xu et al. 2020], [Wang et al. 2022]
1.5	[Thilak and Chandrasekar 2025], [Uddin et al. 2025], [Sharif et al. 2020]
0.5	[Núñez-Molina 2022]

These 36 articles were then subjected to a full quality assessment; Table 1 presents the evaluated articles and their respective scores. Each study was evaluated, ensuring that only publications with a final score of 3.0 or higher advanced. Following the application of this qualitative filter, the final corpus was consolidated into 26 articles. This set constitutes the validated baseline that supports the data extraction and gap identification of this research.

It can be noted that the final corpus exhibits a diverse spectrum of application domains, this heterogeneity is a direct consequence of the non-domain-restricted search string employed. By not artificially restricting the search to a single domain or to more specific domains, this mapping reflects the current cross-disciplinary reality of the literature: agent cognitive paradigms are being coupled in different technological contexts, highlighting the lack of generalised integration frameworks.

3.2. Data Extraction

Following the consolidation of the corpus consisting of 26 primary articles, this subsection details the data extraction process. The primary objective of this stage was to collect structured evidence to address the six RQs.

RQ1 sought to identify the cognitive architectures in simulated MAS. Learning-based agents (predominantly utilising Reinforcement Learning (RL) and Neural Networks) appear in 19 of the mapped studies. The Symbolic/Deliberative approach accounts for 15 occurrences. Furthermore, models based on Large Language Models (LLMs) appear 10 times. Reactive or rule-based agents accounted for 6 occurrences, while utility-based agents recorded 4 occurrences. Notably, the sum of these frequencies exceeds the total number of 26 analysed articles, as several works propose compositional systems that integrate multiple paradigms within the same agent.

RQ2 investigated the software ecosystem used to host, process, and render MAS experiments. Table 2 synthesises the identified platform categories. The “Frequency” metric indicates the number of articles citing the use of at least one tool from that category. The extracted data demonstrate the use of logical, mathematical, or RL tools in 19 occurrences. Custom simulators built from scratch appear in 13 studies, while physical simulators and embodied AI platforms appear in 16 studies. Finally, commercial game engines were documented in only 2 of the analysed works.

Table 2. Categorisation of simulation tools and platforms (RQ2).

Platform Category	Frequency
Logical, Mathematical, or RL Frameworks	19
Custom Simulator or Framework	13
Physical and Robotics Simulators	10
Embodied AI Platforms	6
Commercial Game Engines	2

RQ3 focused on mapping the network topology and software architectural patterns used for communication between the reasoning engine and the spatial rendering environment. The coupling architecture is omitted or unspecified in 15 out of the 26 evaluated articles. Among the minority of studies that document this integration, 6 occurrences utilise decoupled architectures via APIs and message passing, adopting a client-server model mediated by sockets with the JSON protocol. Other mapped occurrences include shared memory (2 occurrences), event-driven integrations (2 occurrences), and static file exchange (1 occurrence).

RQ4 aimed to map the spatial abstraction level and the graphical rendering fidelity of the virtual environments. Table 3 synthesises these categories; the total exceeds 26 articles because some studies adopt hybrid visual approaches. Data extraction evidences the use of 3D environments in 16 articles. 2D abstractions (Grid-worlds) appear in 10 articles. Finally, 7 articles operate without visual fidelity, providing no graphical representation of the sociotechnical space and relying on semantic, algebraic, or textual worlds.

Table 3. Categorisation of visual fidelity and spatial representation (RQ4).

Visual / Spatial Fidelity Level	Frequency	Characteristics and Mapped Examples
3D Volumetric and Virtual Reality (VR)	16	Realistic 3D rendering; photorealistic 3D for Embodied AI; Voxel-based 3D worlds; high continuous visual fidelity; VR
2D (Grid, Cellular, Isometric, or Continuous)	10	Top-down grid/cellular abstraction (Grid-world); continuous and analytical 2D; pixelated matrices for classic games; 2.5D isometric
Null / Purely Textual, Algebraic, or Logical	7	Total absence of spatial representation; purely informational and logical-semantic environment; purely textual simulation based on natural language; strictly algebraic processing

RQ5 sought to map the practical contexts and application domains in which the proposed MAS are instantiated. Table 4 categorises the extracted domains.

Table 4. Categorisation of application domains and research contexts (RQ5).

Domain Category	Frequency	Examples of Identified Real-World Problems
Robotics and Motion Planning	10	Service robotics and task and motion planning; drone and swarm choreography; mobile robot navigation; robot soccer; continuous manipulation; search missions with UAVs; tactical surface surveillance
Embodied AI and Open Worlds	5	Multi-object autonomous visual navigation; embodied AI in open survival worlds; semantic grounding in domestic environments; adaptation and crafting in partially observable worlds; simulated home automation
Autonomous Vehicles, Traffic, and Urban Simulation	4	Urban simulation and autonomous traffic; urban autonomous driving; driving in regulated environments; driver education and training
Energy, Industry, and Advanced Physical Simulation	4	Smart grid and microgrid management; building energy modelling; power generation systems and industrial optimisation; computational fluid dynamics and spatial domain awareness
Classic Games, Strategy, and Benchmarks	3	Tactical combat in RPGs; board games and retro video games; real-time strategy games
Logical and Negotiation	1	Automated negotiation and bilateral conflict resolution.

The quantitative data reveal that the robotics and motion planning domain appears in 10 articles, followed by Embodied AI (5 articles) and autonomous vehicles/traffic (4 articles). Classic games, strategy, and benchmarks account for 3 articles, while logical and negotiation systems each account for 1 occurrence.

RQ6 aimed to map the interaction paradigms required to parameterise the simulation environment and to instantiate the agents' initial cognitive state.

The predominant paradigm is the direct manipulation of code and scripts (19 occurrences). The use of external configuration files and Domain-Specific Languages (DSLs) appears in 14 studies, whilst natural language parameterisation via prompt engineering appears in 8 studies. Finally, Graphical User Interfaces (GUIs) were identified in only 3 occurrences. Knowledge graphs and databases accounted for 2 occurrences, whilst 1 occurrence was unspecified.

3.3. Trends and Gaps

The synthesis of the extracted data enables delineation of the boundaries of the state of the art in MAS simulation. The analytical intersection of the RQs highlights a transitional scenario marked by accelerated advances in agent cognitive engineering in stark contrast to the stagnation of software engineering in the virtual environments that host them. From this analysis, it was possible to map the main methodological trends in the field. The main trend concerns the *Cognitive and Spatial Frontier*: the research field is rapidly shifting toward "foundational cognition," characterised by the abandonment of purely reactive models in favour of hybrid and LLM-driven agents.

Despite the trend toward highly complex entities, the available technological toolkit has not kept pace with this evolution, revealing three systemic gaps. *Gap 1 (Coupling Black Box)*: there are no standardised middlewares to integrate agent brains into their respective environments, a scarcity that forces scientists to dedicate valuable research time to basic engineering for building ad hoc simulators, undermining the scalability of scientific findings. Furthermore, *Gap 2 (Realism versus Flexibility)*: researchers find themselves balancing realism and flexibility when choosing their tools, forced to choose between the photorealism of rigid, complex-to-integrate physical simulators and the flexibility of primitive 2D representations (Grid Worlds).

Finally, *Gap 3 (Usability Deficit)*: the interaction paradigm with simulators acts as an operational bottleneck in the literature, where the near-total absence of GUIs forces researchers to act primarily as rigid-syntax programmers. The reliance on direct manipulation of hardcoded scripts slows experimentation and increases susceptibility to syntactic errors in test environments. Ultimately, the evidence of this structural mismatch proves that, for MAS to continue evolving, the development of new visual bridges and decoupled architectures that resolve such bottlenecks is methodologically imperative.

4. Conclusion

The analysis of the corpus revealed that, although agent cognitive engineering has reached new heights driven by hybrid approaches, continuous learning, and foundational models, the software infrastructures hosting and rendering these intelligences are stagnant.

The three identified systemic gaps, the coupling black box of architectural integration, the methodological conflict between realism and flexibility, and the usability deficit, expose weaknesses in contemporary experimentation. The absence of standardised middleware and the historical reliance on low-level, rigid (hardcoded) parameterisations force researchers to build ad hoc simulators. This isolated and repetitive practice consumes valuable research time and actively compromises the reproducibility and scalability of scientific findings.

These results demonstrate the need to develop decoupled architectures to mitigate these gaps. The delegation of the logical deliberation cycle to cognitive frameworks and volumetric rendering to game engines, interconnecting them via asynchronous communication bridges, emerges as the pathway to eliminate integration bottlenecks.

As future work, this research will focus on proposing standardised integration protocols and expanding the implementations of a bridge to link cognitive frameworks and game engines. The objective is to improve experiment parameterisation, allowing researchers to adjust environmental variables and cognitive states at runtime without the need to modify the source code. Additionally, it is planned to conduct studies to investigate the performance of new decoupled architectures in comparison with tools already established in the literature.

Acknowledgments

This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brasil (CAPES) - Finance Code 001, through the CAPES Social Demand scholarship (CAPES-DS). This work is also funded by project 444568/2024-7, CNPq/MCTI/FNDCT No. 22/2024, Conhecimento Brasil Programme – Support for Network Projects with Brazilian Researchers Abroad.

References

Abar, S., Theodoropoulos, G. K., Lemarinier, P., and O'Hare, G. M. P. (2017). Agent based modelling and simulation tools: A review of the state-of-art software. *Computer Science Review*, 24:13–33.

Al Shukairi, H. and Cardoso, R. C. (2023). ML-MAS: A hybrid AI framework for self-driving vehicles. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*, page 1191–1199. International Foundation for Autonomous Agents and Multiagent Systems.

Angelotti, G. and Díaz-Rodríguez, N. (2023). Towards a more efficient computation of individual attribute and policy contribution for post-hoc explanation of cooperative multi-agent systems using myerson values. *Knowledge-Based Systems*, 260:110189.

Bahamid, A., Mohd Ibrahim, A., and Shafie, A. A. (2024). Crowd evacuation with human-level intelligence via neuro-symbolic approach. *Advanced Engineering Informatics*, 60:102356.

Beeching, E., Debangoye, J., Simonin, O., and Wolf, C. (2021). Godot reinforcement learning agents. arXiv preprint. Version v1, 7 Dec 2021.

Cai, Z., Cardenas, C. R., Leo, K., Zhang, C., Backman, K., Li, H., Li, B., Ghorbanali, M., Datta, S., Qu, L., Gutierrez, J., Ignatiev, A., Li, Y.-F., Vered, M., Stuckey, P. J., de la Banda, M. G., and Rezatofighi, H. (2025). Neusis: A compositional neuro-symbolic framework for autonomous perception, reasoning, and planning in complex uav search missions. *IEEE Robotics and Automation Letters*, 10(9):9502–9509.

Chen, Y., Arkin, J., Zhang, Y., Roy, N., and Fan, C. (2024). Scalable multi-robot collaboration with large language models: Centralized or decentralized systems? In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, page 4311–4317. IEEE.

Dang-Nhu, R. (2020). Plans: neuro-symbolic program learning from videos. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, NIPS '20, Red Hook, NY, USA. Curran Associates Inc.

de Mello, R. R. P., de Santiago, R., Silveira, R. A., and Gelaim, T. A. (2024). Neural-symbolic bdi-agent as a multi-context system: A case study with negotiating agent. *Expert Systems with Applications*, 238:121656.

Deane, O. and Ray, O. (2025). Symplex: Learning social norm hierarchies by combining autonomous exploration and expert imitation. In *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '25, page 2493–2495, Richland, SC. International Foundation for Autonomous Agents and Multiagent Systems.

Du, Q., Gray, A., Li, B., Li, Z., Luo, J., Ravikumar, P., Scherer, S., Si, X., Stepputtis, S., Sycara, K., Wang, C., Wang, W., and Xie, Y. (2024). Logicity: Advancing neuro-symbolic ai with abstract urban simulation. In *Advances in Neural Information Processing Systems 37*, NeurIPS 2024, page 69840–69864. Neural Information Processing Systems Foundation, Inc. (NeurIPS).

Eichelbeck, M., Markgraf, H., and Althoff, M. (2026). Commonpower: A framework for safe data-driven smart grid control. *IEEE Transactions on Smart Grid*, 17(1):71–82.

Ghallab, M., Nau, D., and Traverso, P. (2025). *Acting, Planning, and Learning*. Cambridge University Press.

Goel, S., Lymperopoulos, P., Thielstrom, R., Krause, E., Feeney, P., Lorang, P., Schneider, S., Wei, Y., Kildebeck, E., Goss, S., Hughes, M. C., Liu, L., Sinapov, J., and Scheutz, M. (2024). A neurosymbolic cognitive architecture framework for handling novelties in open worlds. *Artificial Intelligence*, 331:104111.

Grosvenor, A., Zemlyansky, A., Wahab, A., Bohachov, K., Dogan, A., and Deighan, D. (2025). Hybrid intelligence systems for reliable automation: advancing knowledge work and autonomous operations with scalable ai architectures. *Frontiers in Robotics and AI*, 12.

Juliani, A., Berges, V.-P., Teng, E., Cohen, A., Harper, J., Elion, C., Goy, C., Gao, Y., Henry, H., Mattar, M., Lange, D., et al. (2020). Unity: A general platform for intelligent agents. *arXiv preprint arXiv:1809.02627*. arXiv:1809.02627v2, 6 May 2020.

Kitchenham, B. and Charters, S. (2007). Guidelines for performing systematic literature reviews in software engineering. EBSE Technical Report EBSE-2007-01, Software Engineering Group, School of Computer Science and Mathematics, Keele University and Department of Computer Science, University of Durham. Version 2.3. Dated 9 July 2007.

Labiosa, A., Wang, Z., Agarwal, S., Cong, W., Hemkumar, G., Harish, A. N., Hong, B., Kelle, J., Li, C., Li, Y., Shao, Z., Stone, P., and Hanna, J. P. (2025). Reinforcement learning within the classical robotics stack: A case study in robot soccer. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, page 14999–15006. IEEE.

Lazarin, N., Pantoja, C., and Viterbo, J. (2026). My body, my perceptions: A shift from computationalism to embodied cognition in bdi-agent-based embedded systems. In *Proceedings of the 25th International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '26, page 800–809, Richland, SC. International Foundation for Autonomous Agents and Multiagent Systems.

Liu, B., Jiang, Y., Zhang, X., Liu, Q., Zhang, S., Biswas, J., and Stone, P. (2023). Llm+p: Empowering large language models with optimal planning proficiency.

Lorang, P., Goel, S., Shukla, Y., Zips, P., and Scheutz, M. (2024). A framework for neurosymbolic goal-conditioned continual learning in open world environments. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, page 12070–12077. IEEE.

Mellado, A. L. L., Neres, G. G., Borges, A. P., Cardoso, R. C., and Alves, G. V. (2026). *MASPY: A Python Framework for Developing BDI Agents with Reinforcement Learning*, page 4074–4076. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC.

Mellado, A. L. L., Pinz Borges, A., and Alves, G. V. (2025). Maspy: A python-based framework for developing bdi multi-agent systems. In *Advances in Practical Applications of Agents, Multi-Agent Systems, and Computational Social Science: The PAAMS Collection*, pages 216–227, Cham. Springer Nature Switzerland.

Moon, J. (2021). Plugin framework-based neuro-symbolic grounded task planning for multi-agent system. *Sensors*, 21(23):7896.

Núñez-Molina, C. (2022). Application of neurosymbolic ai to sequential decision making. In Raedt, L. D., editor, *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, pages 5863–5864. International Joint Conferences on Artificial Intelligence Organization. Doctoral Consortium.

Rajvanshi, A., Sikka, K., Lin, X., Lee, B., Chiu, H.-P., and Velasquez, A. (2024). Saynav: Grounding large language models for dynamic planning to navigation in new environments. *Proceedings of the International Conference on Automated Planning and Scheduling*, 34:464–474.

Rakhman, U., Ahn, J., and Nam, C. (2021). Fully automatic data collection for neuro-symbolic task planning for mobile robot navigation. In *2021 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, page 450–455. IEEE.

Rehm, J., Reshodko, I., Børresen, S. Z., and Gundersen, O. E. (2024). The virtual driving instructor: Multi-agent system collaborating via knowledge graph for scalable driver education. *AI Magazine*, 45(4):514–525.

Schrittwieser, J., Antonoglou, I., Hubert, T., Simonyan, K., Sifre, L., Schmitt, S., Guez, A., Lockhart, E., Hassabis, D., Graepel, T., Lillicrap, T., and Silver, D. (2020). Mastering atari, go, chess and shogi by planning with a learned model. *Nature*, 588(7839):604–609.

Schuck, M., Dahanaggamaarachchi, D. O., Sprenger, B., Vyas, V., Zhou, S., and Schoellig, A. P. (2025). Swarmgpt: Combining large language models with safe motion planning for drone swarm choreography. *IEEE Robotics and Automation Letters*, 10(11):12237–12244.

Sharif, M., Yasmin, M., Hussain, S., and Shah, J. H. (2020). A review of reinforcement learning based medical image analysis. *Multimedia Tools and Applications*, 79(31):22621–22652.

Silver, T., Chitnis, R., Curtis, A., Tenenbaum, J. B., Lozano-Pérez, T., and Kaelbling, L. P. (2021). Planning with learned object importance in large problem instances using graph neural networks. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(13):11962–11971.

Sohn, S., Woo, H., Choi, J., and Lee, H. (2020). Meta reinforcement learning with autonomous inference of subtask dependencies.

Sovrano, F., Raymond, A., and Prorok, A. (2022). Explanation-aware experience replay in rule-dense environments. *IEEE Robotics and Automation Letters*, 7(2):898–905.

Subramanian, C., Liu, M., Khan, N., Lenchner, J., Amarnath, A., Swaminathan, S., Riegel, R., and Gray, A. (2024). A neuro-symbolic approach to multi-agent rl for interpretability and probabilistic decision making.

Thilak, K. R. and Chandrasekar, P. (2025). Modeling and simulation of anaerobic digestion-gasification integration using madrl-fahp. *Biofuels*, 0(0):1–24.

Uddin, M. S., Ahmed, A., Aktarujjaman, M., Monirujjaman, M., Ahmed, M., Mridha, M. F., and Hossen, M. J. (2025). A hybrid reinforcement learning and knowledge graph framework for financial risk optimization in healthcare systems. *Scientific Reports*, 15(1):29057.

Wang, G., Wei, F., Jiang, Y., Zhao, M., Wang, K., and Qi, H. (2022). A multi-auv maritime target search method for moving and invisible objects based on multi-agent deep reinforcement learning. *Sensors*, 22(21).

Wang, G., Xie, Y., Jiang, Y., Mandlekar, A., Xiao, C., Zhu, Y., Fan, L., and Anandkumar, A. (2023). Voyager: An open-ended embodied agent with large language models.

Wooldridge, M. (2009). *An Introduction to MultiAgent Systems*. John Wiley & Sons, 2 edition.

Xu, J., Wang, H., Niu, Z., Wu, H., and Che, W. (2020). Knowledge graph grounded goal planning for open-domain conversation generation. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):9338–9345.

Yaw, C. T., Yap, K. S., Wong, S. Y., Yap, H. J., and Paw, J. K. S. (2020). Enhancement of neural network based multi agent system for classification and regression in energy system. *IEEE Access*, 8:163026–163043.

Zellers, R., Holtzman, A., Peters, M., Mottaghi, R., Kembhavi, A., Farhadi, A., and Choi, Y. (2021). Piglet: Language grounding through neuro-symbolic interaction in a 3d world. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, page 2040–2050. Association for Computational Linguistics.

Zhang, H., Du, W., Shan, J., Zhou, Q., Du, Y., Tenenbaum, J. B., Shu, T., and Gan, C. (2024). Building cooperative embodied agents modularly with large language models.

Zhang, L., Ford, V., Chen, Z., and Chen, J. (2025). Automatic building energy model development and debugging using large language models agentic workflow. *Energy and Buildings*, 327:115116.