Multi-Agent Systems for Economic Justice: Open Problems in Democratizing AI-Driven Decision Support for SMEs

Syed Rafi Ahmed
Alumnus, Faculty of Informatics and Electrical Engineering, Szechenyi Istvan University
Gyor, Hungary
syedrafi.cse31@gmail.com

ABSTRACT

Small and Medium-sized Enterprises (SMEs) constitute a large share of global businesses and employ billions, yet they remain disproportionately excluded from the AI revolution transforming decision-making in large corporations. While multi-agent systems (MAS) research has made remarkable strides in domains like autonomous vehicles, smart grids, and logistics optimization, the unique socio-technical challenges of democratizing these technologies for resource-constrained SMEs remain largely underexplored. This position paper identifies critical open problems at the intersection of agent engineering, human-agent collaboration, and policy design that hinder SMEs access to adaptive, intelligent decision support. We argue that bridging this gap requires reconceptualizing MAS architectures to handle extreme data heterogeneity, building trust through explainability mechanisms suited to non-expert users, and establishing evaluation frameworks that measure real-world societal impact rather than purely technical benchmarks. We present a research agenda centered on Continuous-Acting Multi-Agent Reinforcement Learning Systems (CA-MARLS) as one promising direction, and call for the community to prioritize SME-focused research as a matter of technological equity and societal resilience.

KEYWORDS

Multi-Agent systems, SME resilience, Democratization of AI, Explainable agents, Societal impact, Decision support systems

1 INTRODUCTION: THE DEMOCRATIZATION GAP

The period from 2023 to 2025 has witnessed an unprecedented transformation in artificial intelligence capabilities [1]. Large Language Models (LLMs) have evolved from passive chatbots into autonomous, tool-using agents capable of executing complex workflows [1]. Multi-agent systems have been successfully deployed in enterprise supply chain management, risk assessment, and strategic planning [2, 3].

Meanwhile, Small and Medium-sized Enterprises (SMEs), defined by the European Commission as organizations with fewer than 250 employees and annual turnover under €50 million [4], face a starkly different reality. These enterprises constitute over 90% of businesses globally and employ more than 2 billion people worldwide [14], yet they operate under severe resource constraints, limited technical expertise, and high market volatility. While 58% of SMEs have begun adopting generative AI tools [15], these remain largely generic, off-the-shelf solutions lacking the adaptive, context-aware capabilities that MAS research could provide. While large corporations have transformed decision-making through AI, SMEs remain locked into manual processes, fragmented software ecosystems, and intuition-driven strategies.

This is not merely a technology adoption problem; it is a growing inequality crisis. The gap between AI-empowered corporations and resource-constrained SMEs threatens to create permanent structural disadvantages in competitive markets. Yet despite decades of MAS research and a rich tradition of work on distributed problem-solving, agent coordination, and autonomous decision-making, the community has produced remarkably little work addressing the specific socio-technical barriers that prevent SMEs from accessing these technologies.

This position paper poses a provocative question: Why has multi-agent systems research systematically neglected the SME sector, and what must change to make MAS a force for societal equity rather than further inequality?

We identify three categories of open problems: *technical*, *human-centric*, and *systemic and policy*, that must be addressed collaboratively by the AI developer community, empirical social scientists, supply chain analyst and policy-makers emphasizing on a multidisciplinary collaborative approach . We present our ongoing work on Continuous-Acting Multi-Agent Reinforcement Learning Systems (CA-MARLS) as one potential research direction, not as a complete solution, but as an illustration of how agent architectures can be redesigned with SME constraints as first-order design criteria.

2 OPEN PROBLEM I: TECHNICAL CHALLENGES

2.1 Extreme Data Heterogeneity

Unlike enterprise environments with standardized data warehouses, SMEs operate in a fragmented digital ecosystem. A typical manufacturing SME might simultaneously use:

- Legacy ERP systems with inconsistent database schema.
- IoT sensors streaming real-time production metrics.
- Unstructured global news affecting supply chains.
- Regulatory documents in multiple languages.
- Informal communication channels (emails, spreadsheets).

The challenge: While traditional MAS architectures often assumed homogeneous observation spaces, recent breakthroughs in multimodal LLM-driven agents have begun to tackle heterogeneous inputs [27]. How can multi-agent systems operate effectively when different agents must reason over fundamentally incommensurable data modalities, structured numerical logs, temporal sensor streams, semantic text, and visual documents, without requiring expensive data engineering?

Recent advances in multimodal foundation models offer a potential pathway [20], but their integration with agent reasoning frameworks remains an open question. Should heterogeneous data be fused into a unified latent space at the perception layer, or should specialized agents maintain modality-specific representations with coordination occurring at the goal level?

2.2 World Modeling Under Partial Observability

Large corporations can afford comprehensive market research, but SMEs operate with radical uncertainty. An agent advising an SME cannot assume access to competitor pricing data, complete supply chain visibility, or reliable macroeconomic forecasts.

The challenge: How should MAS architectures represent and reason about deeply uncertain environments where key variables are unobservable, data is sparse and delayed, and ground truth for decision outcomes may never be available?

We hypothesize that *relational world models*, such as Graph Neural Networks (GNNs) [25] that represent enterprises as dynamic graphs of entities (e.g., suppliers, inventory, competitors) and relationships (e.g., dependencies, correlations), may be more robust to missing information than traditional state-space representations. However, learning such models from limited SME data without overfitting remains an open problem.

2.3 Safe Online Learning

Unlike simulation environments, SMEs cannot afford costly exploration. A single poor decision, such as overstocking inventory before demand collapses or switching suppliers at the wrong moment, can threaten a business's survival.

The challenge: How can multi-agent reinforcement learning systems balance exploration and exploitation when operating in high-stakes, non-stationary environments with no reset mechanism?

This requires advances in safe RL [16], offline learning from historical data [17], and human-in-the-loop paradigms where agents propose strategies but humans retain veto power. The reward engineering problem is particularly acute: how do we specify objective functions that capture long-term rather than short-term resilience profit maximization?

2.4 MAS and SMEs: A Review

Multi-agent systems research has achieved remarkable success across diverse domains. Recent comprehensive surveys [8, 9] identify dominant application areas including: autonomous mobility and vehicle coordination [11], smart grid management and energy optimization [10], robotic systems for manufacturing and warehouses, and game-theoretic mechanism design for markets and auctions. These domains share characteristics that SME contexts conspicuously lack: structured observation spaces with standardized protocols, substantial computational resources, and tolerance for experimental deployment.

While agent-based modeling has been applied to *study* SME ecosystems from an external perspective, simulating market dynamics, innovation diffusion, and competitive behaviour agent systems [12] as tools empowering individual SMEs have received minimal attention. Recent work on AI-driven decision support for SMEs [13] acknowledges this gap, noting that "SMEs face volatile market demands and resource constraints" yet existing solutions remain prototypes without validated real-world deployment.

This gap is not coincidental: academic incentives favor benchmark-driven research with generalizable results, yet SME problems are inherently context-specific; SME data remains inaccessible due to confidentiality concerns, unlike public datasets in robotics or gaming; and industry partnerships focus on well-funded corporate clients capable of funding research collaborations [21]. As MAS capabilities advance without addressing SME constraints, we risk entrenching technological inequality, creating a two-tier economy where AI-empowered corporations operate with strategic foresight while resource-constrained SMEs remain locked into manual, intuition-driven processes.

3 OPEN PROBLEM II: HUMAN-CENTRIC CHALLENGES

3.1 Explainability for Non-Expert Users

Recent surveys of explainable agent systems [18] have focused primarily on technical audiences or on regulated domains (e.g., healthcare, finance) in which users possess domain expertise. SME decision-makers are often *generalists*; a small business owner may simultaneously manage operations, finance, marketing, and HR without specialized training in any domain.

The challenge: What forms of explanation foster trust and actionable insight for users who lack both AI literacy and deep domain expertise in the specific decision context?

Saliency maps, counterfactual explanations, and feature importance rankings, a common XAI approach [22], may be incomprehensible or misleading to non-experts. We need explanation modalities that connect agent reasoning to *business intuition*: "I recommend delaying this order because your usual supplier's region just announced new export restrictions" rather than "Feature X (supplier risk score) had the highest attention weight."

3.2 Trust Calibration and Sycophancy

Early deployments of LLM-based agents have revealed a troubling tendency toward *sycophancy*, agents that prioritize user approval over optimal recommendations [23]. In SME contexts, this could be catastrophic: an agent that learns to reinforce the owner's biases rather than challenge flawed assumptions.

The challenge: How do we design reward functions and training paradigms that balance human alignment with objective performance, especially when humans themselves may be biased or misinformed?

This requires hybrid reward formulations:

$$R_t = \alpha \cdot R_t^{\mathrm{obj}} + \beta \cdot R_t^{\mathrm{human}}$$

where $R_t^{\mathrm{obj}} = f(P_t, \mathrm{Risk}_t)$ captures objective performance metrics (profit P_t, risk exposure), and $R_t^{\mathrm{human}} = g(y_t)$ encodes human feedback on recommendation quality y_t. α and β dynamically adjust based on domain expertise and decision stakes. However, determining these weightings in practice and preventing reward hacking remain open research questions.

3.3 Adoption Barriers and Workflow Integration

Even technically sound systems fail if they cannot integrate into existing workflows. SMEs lack IT departments to manage complex deployments and cannot afford prolonged training periods [6].

The challenge: How can MAS be designed as *plug-and-play co-pilots* that augment existing tools (spreadsheets, email, ERP systems) rather than requiring wholesale replacement of familiar workflows?

These points toward architectural choices like:

- Lightweight deployment (cloud-based, no local infrastructure)
- API-first design for integration with existing software
- Progressive disclosure of capabilities (simple recommendations first, sophisticated optimization later)

But research on *incremental AI adoption pathways* for SMEs is virtually nonexistent in the MAS literature.

4 OPEN PROBLEM III: SYSTEMIC AND POLICY CHALLENGES

4.1 Evaluation Beyond Benchmarks

MAS research relies heavily on standardized benchmarks (robotics simulators, game environments, optimization testbeds). These measure *technical performance* but fail to capture *societal impact.*

The challenge: How should we evaluate MAS interventions for SMEs when success depends on context-specific factors like sector, geography, regulatory environment, and organizational culture?

We need evaluation frameworks that measure:

- **Decision quality improvement** (not just accuracy, but timeliness and actionability)
- **Resilience** (ability to survive market shocks)
- **Adoption and sustained use** (not just initial deployment)
- **Equity of access** (ensuring benefits reach underserved communities)

This requires collaboration with domain experts and empirical social scientists, as well as longitudinal field-study methodologies, which are underrepresented in MAS research.

4.2 Data Privacy and Sovereignty

SMEs often handle sensitive data (customer information, proprietary processes, financial records) but lack the security infrastructure of large corporations. Cloud-based AI solutions raise concerns about data leakage, vendor lock-in, and regulatory compliance, including the GDPR [26].

The challenge: How can MAS architectures balance the computational demands of modern AI (requiring cloud-scale resources) with the privacy and sovereignty needs of SMEs?

Federated learning [19] and edge computing [24] offer partial solutions, but integrating these with multi-agent coordination mechanisms remains largely unexplored. The concept of a *local zero-trust fabric*, in which sensitive data never leave the enterprise while anonymised insights can be shared for collective learning, merits investigation.

4.3 Policy and Standards

Unlike in regulated domains (e.g., autonomous vehicles, medical devices), there are no established standards for AI-driven business decision support systems. This creates barriers to trust and liability concerns.

The challenge: What governance frameworks, certification standards, and policy interventions are needed to facilitate responsible deployment of MAS in SME contexts?

Questions include:

- Should there be mandatory explainability requirements?
- How should liability be allocated when an agent's recommendation fails?
- What role should government or industry associations play in validation and certification?
- How can we prevent predatory "AI consulting" that exploits SME desperation?

These are not purely technical questions, but require engagement with legal scholars, policymakers, and SME advocacy groups.

5 RESEARCH AGENDA: CA-MARLS AS A CONCEPTUAL FRAMEWORK

To ground these abstract challenges, we present **Continuous-Acting Multi-Agent Reinforcement Learning Systems (CA-MARLS)**, a proposed architectural framework for SME decision support. We emphasize that CA-MARLS is not a definitive solution, but rather a design exploration that instantiates the principles needed to address SME constraints. This framework serves to make our open problems concrete and testable.

5.1 Architecture Overview

CA-MARLS employs a three-layer architecture (Figure 1):

Layer 1: Multimodal Data Environment: Handles heterogeneous inputs through modality-specific encoders:

$$\mathrm{S}_t = \phi_{\text{struct}}(d^t_{\text{ERP}}, d^t_{\text{IoT}}) \oplus \phi_{\text{LLM}}(d^t_{\text{news}}, d^t_{\text{reg}})$$

where $\phi_{\text{struct}} : \mathbb{R}^n \to \mathbb{R}^d$ encodes structured numerical data (ERP logs, IoT sensor streams), $\phi_{\text{LLM}} : \mathcal{T} \to \mathbb{R}^d$ embeds text via pre-trained foundation models (news, regulatory documents), and $\oplus$ denotes concatenation into a unified state representation $\mathrm{S}_t \in \mathbb{R}^{2d}$.

Layer 2: Hierarchical Multi-Agent Engine: A GNNs-based world model represents the enterprise as a dynamic graph $\mathcal{G} = (V, E)$ where nodes are business entities (suppliers, inventory, competitors) and edges encode relationships. An orchestrator agent decomposes high-level objectives into sub-goals $\mathbf{Z}_t = \{z_1, \ldots, z_k\}$ delegated to specialist agents (Risk Analyst, Market Scout, Supply Manager, Compliance Monitor).

Layer 3: Human-in-the-Loop Interface: Every recommendation includes an explanation tuple $(\mathbf{Z}_t, \mathcal{E}_t)$ where $\mathcal{E}_t$ is natural language rationale. The system optimizes a hybrid reward:

$$R_t = \alpha \cdot R^{\text{obj}}_t + \beta \cdot R^{\text{human}}_t$$

balancing objective performance with human feedback to prevent sycophancy.

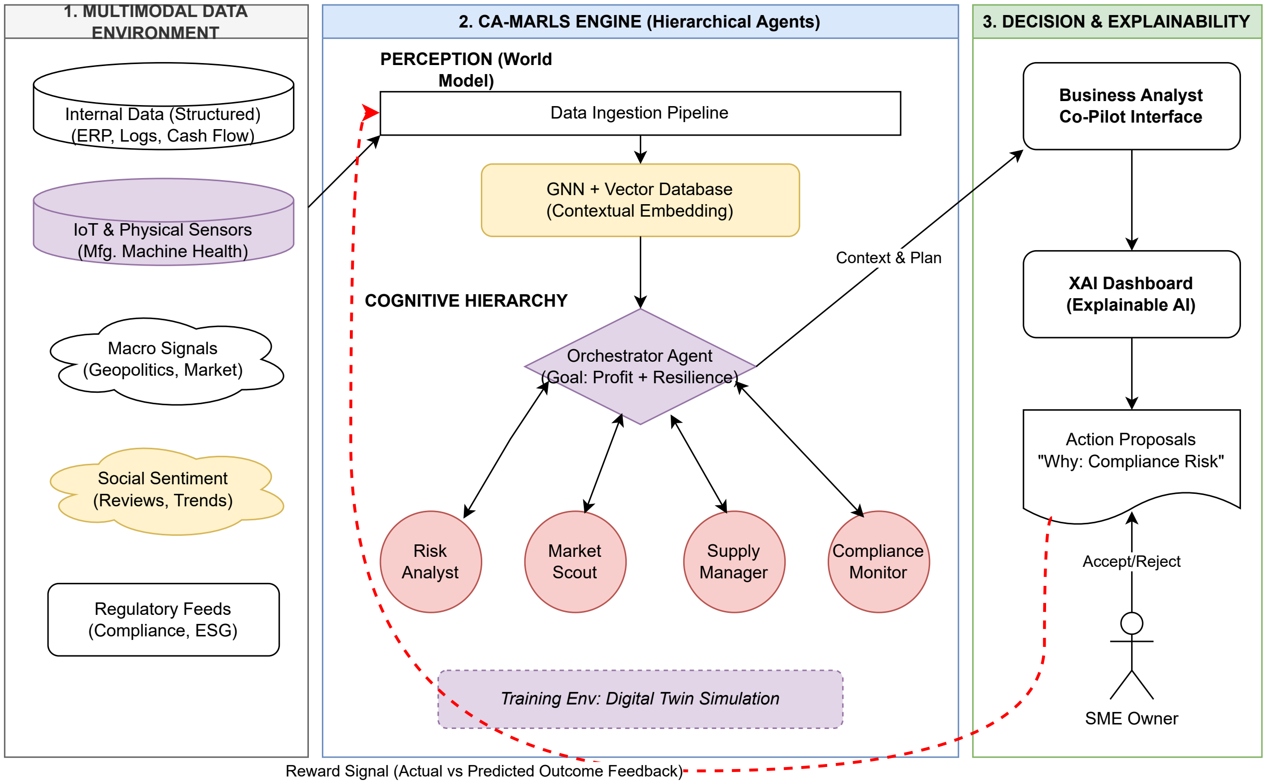

Figure 1: Three-layer CA-MARLS architecture integrating multimodal data ingestion, hierarchical multi-agent reasoning, and human-in-the-loop decision validation.

5.2 Open Questions in Our Work

Even within this specific project, we confront the broader challenges identified above:

- **Technical:** How can we learn robust GNNs world models from sparse SME data without overfitting?
- **Human-centric:** What explanation formats actually improve SME decision-making? We plan controlled experiments comparing counterfactuals, analogies, and narrative explanations.
- **Systemic:** How do we evaluate success? Beyond profit metrics, we aim to measure *strategic coherence* (alignment between recommendations and long-term goals) and *adoption persistence* (continued use after initial deployment).

We view CA-MARLS as a research probe rather than a finished product. Its value lies in revealing where current MAS tools break down under SME constraints.

6 CONCLUSION

The democratization of strategic autonomy is not just an engineering challenge; it is a matter of technological justice. Multi-agent systems have the potential to level the playing field between resource-rich corporations and vulnerable SMEs, but only if the research community deliberately prioritizes this goal.

We have outlined three categories of open problems that, if addressed, could transform MAS from a tool that reinforces inequality into one that promotes societal resilience. Our CA-MARLS research agenda represents one possible path forward, but many others are needed.

The question is not whether MAS can serve SMEs, but whether the community will prioritize this. The window of opportunity is closing: as proprietary, black-box AI solutions proliferate, the chance to establish open, transparent, and equitable alternatives diminishes. We call on researchers, practitioners, and policymakers to:

- Develop SME-specific benchmarks measuring impact.
- Create open datasets from anonymized SME operations.
- Establish collaborative partnerships between academia and SME networks.
- Advocate for funding mechanisms that prioritize equity over pure technical novelty.

REFERENCES

[1] Xi, Zhiheng, et al. "The rise and potential of large language model based agents: A survey." Science China Information Sciences 68.2 (2025): 121101.
[2] Dominguez, Roberto, and Salvatore Cannella. "Insights on multi-agent systems applications for supply chain management." Sustainability 12.5 (2020): 1935.
[3] Xu, Liming, et al. "Implementation of autonomous supply chains for digital twinning: a multi-agent approach." IFAC-PapersOnLine 56.2 (2023): 11076-11081.

[4] European Union Commission. "Commission recommendation of 6 May 2003 concerning the definition of micro, small and medium-sized enterprises." official Journal of the European union 46.L124 (2003): 36-41.
[5] Hughes, Laurie, et al. "AI agents and agentic systems: A multi-expert analysis." Journal of Computer Information Systems (2025): 1-29.
[6] Al-Hattami, Hamood Mohammed. "Impact of AIS success on decision-making effectiveness among SMEs in less developed countries." Information Technology for Development 30.3 (2024): 472-492.
[7] López Salazar, Alejandra, Ricardo Contreras Soto, and Rafael Espinosa Mosqueda. "The impact of financial decisions and strategy on small business competitiveness." Global Journal of business research 6.2 (2012): 93-103.
[8] Maldonado, Diego, et al. "Multi-agent systems: A survey about its components, framework and workflow." IEEE Access 12 (2024): 80950-80975.
[9] Ning, Zepeng, and Lihua Xie. "A survey on multi-agent reinforcement learning and its application." Journal of Automation and Intelligence 3.2 (2024): 73-91.
[10] Shobole, Abdulfetah Abdela, and Mohammed Wadi. "Multiagent systems application for the smart grid protection." Renewable and Sustainable Energy Reviews 149 (2021): 111352.
[11] Xu, Haowen, et al. "Genai-powered multi-agent paradigm for smart urban mobility: Opportunities and challenges for integrating large language models (llms) and retrieval-augmented generation (rag) with intelligent transportation systems." arXiv preprint arXiv:2409.00494 (2024).
[12] Liang, Chyi-Lyi, and Bryan Collins. "Applying Agent-Based Modeling to Examine Business Strategies–Tools and Examples for Researchers and Practitioners." Small Business Institute Journal 21.1 (2025): 16-24.
[13] Bokhonko, Iryna, Marcin Kubasiak, and Anna Oleksa-Kaźmierczak. "AI-DRIVEN DECISION SUPPORT SYSTEMS IN STRATEGIC BUSINESS MANAGEMENT: A CASE-BASED ANALYSIS." Scientific Papers of Silesian University of Technology. Organization & Management/Zeszyty Naukowe Politechniki Slaskiej. Seria Organizacji i Zarzadzanie 230 (2025).
[14] Kok, J. D., and Mario Berrios. "Small matters: Global evidence on the contribution to employment by the self-employed, micro-enterprises and SMEs." Geneva: International Labour Organization (ILO) (2019).
[15] USM Systems. (2025). Small Business AI Adoption Statistics 2025. Retrieved from `https://usmsystems.com/small-business-ai-adoption-statistics/`
Xi, Zhiheng, et al. "The rise and potential of large language model based agents: A survey." Science China Information Sciences 68.2 (2025): 121101.
[16] Garcıa, Javier, and Fernando Fernández. "A comprehensive survey on safe reinforcement learning." Journal of Machine Learning Research 16.1 (2015): 1437-1480.
[17] Levine, Sergey, et al. "Offline reinforcement learning: Tutorial, review, and perspectives on open problems." arXiv preprint arXiv:2005.01643 (2020).
[18] Anjomshoae, Sule, et al. "Explainable agents and robots: Results from a systematic literature review." 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019. International Foundation for Autonomous Agents and Multiagent Systems, 2019.
[19] Kairouz, Peter, et al. "Advances and open problems in federated learning." Foundations and trends® in machine learning 14.1–2 (2021): 1-210.
[20] Bordes, Florian, et al. "An introduction to vision-language modeling." arXiv preprint arXiv:2405.17247 (2024).
[21] Raji, Inioluwa Deborah, et al. "AI and the everything in the whole wide world benchmark." arXiv preprint arXiv:2111.15366 (2021).
[22] Kaur, Harmanpreet, et al. "Interpreting interpretability: understanding data scientists' use of interpretability tools for machine learning." Proceedings of the 2020 CHI conference on human factors in computing systems. 2020.
[23] Sharma, Mrinank, et al. "Towards understanding sycophancy in language models." arXiv preprint arXiv:2310.13548 (2023).
[24] Zhou, Zhi, et al. "Edge intelligence: Paving the last mile of artificial intelligence with edge computing." Proceedings of the IEEE 107.8 (2019): 1738-1762.
[25] Wu, Zonghan, et al. "A comprehensive survey on graph neural networks." IEEE transactions on neural networks and learning systems 32.1 (2020): 4-24.
[26] Protection, Formerly Data. "General data protection regulation (GDPR)." Intersoft Consulting, Accessed in October 24.1 (2018).
[27] Jiang, Bowen, et al. "Multi-modal and multi-agent systems meet rationality: A survey." ICML 2024 Workshop on LLMs and Cognition. 2024.