Engineering the Next Generation of Multi-Agent Systems: A Community Roadmap from EMAS 2025

Sebastian Rodriguez[1][0000−0002−0514−9221], Akhila Bairy[2][0000−0002−8796−1474], Matteo Baldoni[3][0000−0002−9294−0408], Patrick Benjamin[4][0009−0009−2652−4957], Constantin Blessing[5][0009−0005−8516−0269], Nicolas Brandstetter[6][0009−0007−9462−4736], Amit K. Chopra[7][0000−0003−4629−7594], Thomas Clemen[8][0000−0002−8200−5141], Louise A. Dennis[9][0000−0003−1426−1896], Ahmad Esmaeili[10][0000−0003−0612−2351], Lu Feng[11][0000−0002−4651−8441], Angelo Ferrando[12][0000−0002−8711−4670], Zahra Ghorrati[13][0000−0003−1396−7248], Victor Guillet[14][0009−0006−6758−3425], Önder Gürcan[15][0000−0001−6982−5658], Soham Hans[16][1111−1111−1111−1111], James Herber[17][0009−0003−6040−5608], Viviana Mascardi[18][0000−0002−2261−9926], Marcel Mauri[19][0000−0002−4135−1945], Jörg P. Müller[20][0000−0001−7533−3852], John Thangarajah[1][0000−0002−7699−6444], Rafał Tyl[21][0009−0007−0528−2181], and Yi Yang[22][0000−0001−9565−1559]

[1] RMIT Universtiy {sebastian.rodriguez,john.thangarajah}@rmit.edu.au
[2] Karlsruhe Institute of Technology akhila.bairy@kit.edu
[3] University of Torino matteo.baldoni@unito.it
[4] University of Oxford patrick.benjamin@cs.ox.ac.uk
[5] Esslingen University constantin.blessing@hs-esslingen.de
[6] Universidad de Chile nicolasbrandstetter@ug.uchile.cl
[7] Lancaster University amit.chopra@lancaster.ac.uk
[8] Hamburg University of Applied Sciences thomas.clemen@haw-hamburg.de
[9] University of Manchester louise.dennis@manchester.ac.uk
[10] Wichita State University ahmad.esmaeili@wichita.edu
[11] University of Virginia lu.feng@virginia.edu
[12] University of Modena and Reggio Emilia angelo.ferrando@unimore.it
[13] Purdue University zghorrat@purdue.edu
[14] ONERA (French aerospace research labs) victor.guillet@onera.fr
[15] NORCE Research AS ongu@norceresearch.no
[16] USC Institute for Creative Technologies sohamhan@usc.edu
[17] Independent Researcher jimherber@gmail.com
[18] University of Genova viviana.mascardi@unige.it
[19] Goethe University Frankfurt mauri@cs.uni-frankfurt.de
[20] Technische Universität Clausthal jmue@tu-clausthal.de
[21] Grail, Gentle Viking rafftyl@wp.pl
[22] KU Leuven yi.yang@kuleuven.be

Abstract. This paper presents the outcomes of an open-floor session held at the 13th International Workshop on Engineering Multi-Agent Systems (EMAS 2025), aimed at co-developing a research roadmap for the EMAS community. Participants collaboratively identified and prioritised challenges in engineering large-scale, adaptive multiagent systems, particularly considering the need to engineer systems that can seamlessly integrate *learning* and *reasoning*. Through structured group discussions, four key challenges emerged: explainability in heterogeneous environments, environment modeling, handling dynamic contexts, and communication standardisation. For each of the challenges, participants proposed and ranked potential solutions based on impact and effort. The resulting roadmap highlights concrete research directions toward engineering intelligent, explainable, and interoperable multiagent systems that effectively integrate reasoning and learning in dynamic environments.

1 Introduction

The design, development and deployment of multi-agent systems (MAS) have evolved significantly over the past two decades, driven by advances in artificial intelligence, distributed computing, and autonomous decision-making. As MAS technologies are increasingly embedded in dynamic and uncertain environments, the need to engineer systems that can seamlessly integrate *learning* and *reasoning* has become a central challenge. This integration raises fundamental questions about adaptability, explainability, coordination, and assurance—key concerns for the EMAS community.

In recent years, the rise of large language models (LLMs) [?] and agentic AI systems [?] has further transformed the landscape of intelligent systems engineering. LLM-based agents – or "AI agents", or "Agentic AI agents" – increasingly demonstrate the ability to make decisions, exchange information, and coordinate tasks. Google defines LLM-based agents as *software systems that use AI to pursue goals and complete tasks on behalf of users. They show reasoning, planning, and memory and have a level of autonomy to make decisions, learn, and adapt.* Their capabilities are *made possible in large part by the multimodal capacity of generative AI and AI foundation models*[23].

Whether LLM-based agents show planning and reasoning capabilities is still open to discussion, with strong claims from some LLMs developers and vendors mitigated by skepticism by some scientists [?,?,?].

Whatever the individual point of view, the EMAS community is discussing about the relationships between agentic AI systems and more traditional approaches in many venues[24]. There is general agreement in recognizing that LLM-based agents may play a role in addressing tasks that are properly coped with by

[23] Google definition of AI agent, `https://cloud.google.com/discover/what-are-ai-agents`, accessed on March 25, 2026.

[24] For example, the *Agents vs. LLMs* panel at the 26th Workshop From Objects to Agents in July 2025, `https://sites.google.com/view/woa2025`; the *Agent Toolkits 2025 Community Session* at the 22nd European Conference on Multi-Agent Systems in early September 2025, `https://interactions.ics.unisg.ch/agent-toolkits-2025/` [?]; the *Rethinking Multi-Agent Systems in the Era of LLMs* Workshop in Oxford in mid September 2025, `https://sites.google.com/view/rethinking-mas/home`.

neither traditional symbolic agents nor data-driven learning architectures, and that a synthesis is needed [?,?,?].

This shift opens unprecedented opportunities for building open, adaptive, and human-aligned MAS—but it also introduces new challenges related to controllability, accountability, transparency, and safety. These developments call for a renewed examination of the principles, methods, and tools that underpin the engineering of multi-agent systems.

The 13^{th} International Workshop on Engineering Multi-Agent Systems (EMAS 2025) provided an ideal forum to reflect on these emerging challenges and to chart a shared vision for the next decade of MAS research. In addition to the traditional paper presentations and discussions, the workshop featured an open-floor, hands-on session aimed at collaboratively developing a research roadmap for the EMAS community. This initiative sought to capture the collective insight of researchers and practitioners and to identify priorities that can guide future work in the field.

The roadmap session adopted a participatory and exploratory format. Participants worked together to articulate key challenges, opportunities, and research questions that arise when engineering MAS capable of combining symbolic reasoning, data-driven learning, and emergent capabilities of LLM-powered agents. To ground the discussion, a search and rescue scenario was used as a concrete example, providing a shared context that illustrated the complex interaction between autonomous agents, humans, and uncertain environments. Through structured dialogue and synthesis, participants proposed conceptual frameworks and technical pathways toward addressing these challenges.

This paper summarizes the outcomes of that collaborative exercise. It consolidates the insights generated during the session into a coherent research roadmap that reflects the current consensus and diversity of perspectives within the EMAS community. The roadmap identifies short-, medium-, and long-term priorities for advancing the top four challenges identified by participants in the engineering of intelligent, adaptive, and trustworthy multi-agent systems.

Beyond documenting the outcomes, this paper also reflects on the process of collective roadmap creation as a community-building and knowledge-integration activity. By fostering open dialogue across methodological traditions and application domains, the workshop demonstrated how participatory approaches can support the co-evolution of the EMAS research agenda and practice that we hope will influence future editions of EMAS.

The remainder of the paper is structured as follows. Section 2 situates this work in the broader landscape of research roadmap initiatives in agent-oriented software engineering and related fields. Section 3 describes the design and facilitation of the workshop, including participant composition and methodological approach. Section 4 presents the main findings derived from the session discussions. Section 5 introduces the resulting research roadmap, outlining the envisioned directions and milestones. Section 6 offers a discussion of cross-cutting insights and reflections, and Section 7 concludes with next steps for sustaining the roadmap as a living artifact of the EMAS community.

2 Background and Related Work

This section summarises the trajectory of previous MAS roadmap and frameworks and defines the key conceptual framing surrounding the challenge.

2.1 Foundational challenges in MAS

The EMAS community has been active in defining shared views and identifying key challenges to overcome when engineering multi-agent systems. A significant prior endeavor was reported following the 6th International Workshop on Engineering Multi-Agent Systems (EMAS 2018)[20]. That initiative centered on identifying the challenges in three areas: Cognitive Agent Architectures, Agent programming and Machine Learning & MAS. Many of these core challenges are still present today as this work shows.

A crucial implication derived from the persistence of these foundational concerns is that the mere introduction of systems with more powerful capabilities, such as those enabled by Large Language Models (LLMs), does not inherently resolve the underlying engineering deficits. Instead, integrating advanced learning and reasoning must be accomplished on top of robust, verifiable engineering frameworks. Without this foundational rigor, new cognitive capabilities risk inheriting and potentially amplifying existing systemic problems, particularly in areas like complexity, unpredictability, and maintenance overhead.

As noted by Dix et al. [14] over a decade ago, several open challenges persist in the EMAS field. While foundational research remains vital, there is an urgent need to strengthen the engineering dimension, particularly in developing and maintaining large-scale multi-agent systems with clear standards of quality and practice. Technological barriers further limit deployment, underscoring the need for novel engineering techniques and the sharing of empirical lessons from real-world applications using authentic agent technologies rather than general-purpose programming languages [14].

Research of the EMAS community has been tackling these foundational challenges for almost 30 years [35].

2.2 Challenges in related areas

The engineering challenges are mirrored and amplified across other key MAS research communities. Challenges that the EMAS community should seek to embrace and provided appropriate methodological support.

The Coordination, Organization, Institutions, Norms, and Ethics in Agent Systems (COINE) community focuses on the scientific and technological aspects of social coordination, organizational theory, normative MAS, and ethics. The close interests of the EMAS and COINE communities were evident in this year's joint panel and discussions[25].

[25] https://emas.in.tu-clausthal.de/2025/programme

The challenge is the governance of open systems, where agents, their interactions, or the system's purpose may dynamically change over time. Although extensive work has been done in this area on topics such as norms [1,12] still a number of open challenges remain [10].

Furthermore, ensuring the system is *designed responsibly* is crucial to trusting its behaviour. Developing a Responsible AI requires more than some "add-on" features [13], that poses its own set of challenges [3]. The role of the EMAS community to provide the appropriate frameworks is critical in this domain.

The scalability imperative was recognized as a major requirement for MAS applications. Roadmaps, such as the E4MAS community's work [33], focused heavily on techniques for developing Large-Scale Multi-Agent Systems (LSMAS). This research established that achieving efficiency at scale relies on decentralized loci of control, moving beyond monolithic architectures. A crucial architectural realization was that the agent environment—the structure defining how agents perceive, act, and interact—is the primary mechanism for solving scaling issues and managing coordination complexity. That concern is still present in the community today (see section 4).

3 Workshop Design and Methodology

The roadmap session took place during the EMAS workshop, co-located with AAMAS in Detroit, USA. Participants included EMAS authors and general attendees from the broader AAMAS community. They represented a wide range of backgrounds across agent-oriented software engineering, artificial intelligence, distributed systems, and human–agent interaction, united by a shared interest in advancing the engineering principles and methodologies underpinning multi-agent systems. The diversity of perspectives provided a rich foundation for identifying both enduring and emerging research priorities within the field.

3.1 Overview

The hands-on roadmap session was designed as an interactive, collaborative exercise. This format emphasises structured, time-boxed activities that move from problem identification to collective prioritisation and solution design. The workshop spanned an entire afternoon and combined individual reflection with group discussion to maximise participation and synthesis.

Participants were introduced to the session goals and briefly reviewed framing materials that summarised current challenges in engineering multi-agent systems, with a particular emphasis on integrating reasoning and learning. In small groups, participants shared their initial thoughts, identified key challenges, and highlighted promising research directions. This ensured a shared understanding of the context and stimulated creative thinking.

In order to co-create a research roadmap for Engineering Multi-Agent Systems, participants then engaged in a sequence of structured collaborative activities. Participants engaged in a silent brainstorming phase to identify key opportunities and problems related to engineering multi-agent systems. Each idea was captured on sticky notes, ensuring that all voices and perspectives were represented. Groups then moved to voting and prioritisation, using dot stickers or stars to highlight the most pressing and impactful challenges. These prioritised items were subsequently reframed into standardized research challenges, enabling a common structure for later analysis.

Once the core challenges had been articulated, participants proposed potential solutions and research pathways, again using sticky notes to encourage breadth and diversity of ideas. After group discussion and voting, the most promising solutions were retained and mapped within an impact–effort matrix. This visual exercise helped distinguish initiatives that could yield immediate results from those requiring more substantial long-term investment.

The session concluded with a plenary sharing and discussion, where each group presented their roadmap and reflected on common patterns, differences, and cross-cutting themes. This final exchange fostered convergence around shared priorities and revealed complementary approaches across the participating groups.

Throughout the workshop, participants used post-it notes, stickers, and handwritten notes to capture challenges, ideas, and proposed solutions. These physical artefacts were collected by the session chairs at the end of the workshop. The materials were later digitised, clustered, and analysed to identify recurrent themes, underlying challenges, and relationships among proposed research directions. The synthesis process informed the thematic structure and roadmap presented in this paper.

3.2 Methodological Reflection and Limitations

The open-floor, participatory workshop format proved effective in stimulating discussion and capturing diverse perspectives from the EMAS community. By framing challenges as research questions and ranking potential solutions by impact and effort, participants were encouraged to move beyond abstract debate toward actionable priorities. This interactive structure fostered inclusivity and creativity, enabling a collective articulation of the field's research roadmap.

However, some limitations must be acknowledged. The outcomes reflect the views of a self-selected group of workshop participants, who—while representative of active researchers in engineering multi-agent systems—may not encompass the full diversity of perspectives within the broader AI and systems engineering communities. Time constraints restricted deeper exploration of interdependencies among challenges and cross-cutting themes such as ethics, scalability, or sustainability.

Future iterations of this roadmap process could benefit from complementary methods to validate and refine the identified priorities.

Then participants were provided with a case study description to ground the discussions. The domain of *Coordinated Drone Swarm & AI-Supported Command for Disaster* was chosen for this purpose. Appendix A contains the guiding material and case study text.

4 Key Challenges and Opportunities

4.1 Opportunities and Advancements

The first phase of the workshop focused on identifying positive aspects and opportunities within the EMAS research community, particularly regarding the integration of learning and reasoning in the design and implementation of autonomous multi-agent systems. Participants highlighted a number of well-established strengths that position EMAS as a mature yet continuously evolving research domain.

Established Engineering Foundations Participants emphasized that the EMAS community has developed a solid foundation of engineering principles and abstractions that continue to enable the systematic design of complex autonomous systems [9,29,35]. Core multi-agent concepts—such as centralized and decentralised architectures, coordination mechanisms, and communication protocols—provide proven building blocks for scalable agent-based solutions. Simple yet effective interaction models remain relevant due to their robustness and modest computational requirements.

The community has also achieved significant progress in human–autonomy teaming, multi-level and hierarchical architectures, and normative systems. These efforts have yielded frameworks capable of managing goals, actions, and roles at varying levels of abstraction, exemplified by holonic and hierarchical multi-agent systems [8,22,26]. Extensive work on multi-agent planning and coordination has produced models that support autonomy while maintaining coherence at the system level [32]. Participants also highlighted EMAS's enduring strength in formal methods and safety verification, providing tools for rigorous design assurance even in dynamic and uncertain domains [2,11].

Furthermore, the development of dedicated programming languages and development environments for multi-agent systems continues to distinguish EMAS from broader AI research communities [4,5,17,24]. These languages encapsulate decades of accumulated knowledge about agent-oriented abstractions, decision processes, and interaction patterns—assets that can inform emerging paradigms in agentic AI.

Integration with Learning and Data-Driven Approaches While traditionally grounded in symbolic and model-based reasoning, the EMAS community has increasingly explored data-driven and hybrid approaches. Participants observed that deductive models remain essential where data availability is limited, such as in time-critical or safety-sensitive environments (e.g., disaster response). At the same time, data-driven decision-making mechanisms have matured to support adaptive command-and-control capabilities and to bridge the gap between autonomous systems and human operators.

Participants noted the opportunity to leverage lightweight cognitive architectures, such as BDI-based agents, in contexts where deploying large language models or deep neural architectures is impractical. The ability to convert raw data into human-readable formats also emerged as a strength, highlighting the community's commitment to interpretability and human-centered design.

Applied Domains and Case Studies With regards to the provided scenario, participants highlighted the community's ongoing engagement with applied multi-agent domains, such as swarm robotics and collective coordination, and was identified as another area of advancement. Significant progress in swarm MAS algorithms and robotic testbeds has provided both empirical validation and reusable experimental infrastructure. These advances offer a concrete foundation for testing hybrid reasoning and learning approaches in realistic, high-stakes environments such as search and rescue—the scenario anchoring this workshop.

Community Culture and Emerging Opportunities Beyond technical contributions, participants recognized several cultural and strategic strengths of the EMAS community. There is a shared understanding that meaningful progress requires interdisciplinary collaboration, integrating insights from AI, cognitive science, robotics, human–computer interaction, and software engineering. The community also demonstrates a consistent openness to emerging technologies and paradigms, rapidly engaging with new trends such as generative AI and LLM-based agents to explore their potential for agent programming and coordination.

Looking ahead, participants identified promising opportunities to create shared testbeds and evaluation infrastructures that could foster collaboration and cumulative progress across institutions. Integrating sensing systems and infrastructure into these testbeds, and experimenting with generative AI for agent programming, were noted as particularly fruitful directions for collective exploration.

4.2 Open Challenges

The community identified a wide range of challenges that were grouped into major thematic areas, each reflecting ongoing technical, methodological, and social barriers in the engineering of multi-agent systems.

The most prominent themes were modeling the environment (8 votes), explainability (7 votes), and communication support (6 votes). These highlight the need for more accurate and adaptable representations of complex, changing environments; mechanisms to make agent behaviors interpretable and transparent to humans; and robust communication frameworks that ensure coordination and understanding among heterogeneous agents and between agents and humans.

Other significant areas included dynamic environments (5 votes), emphasising the need for agents capable of adapting to rapidly changing conditions; acceptance of agents (5 votes), pointing to the social and psychological challenges of integrating autonomous systems into human-centered contexts; and natural language to formal representation (3 votes), addressing the gap between intuitive human inputs and formal computational models.

Additional concerns, though receiving fewer votes, remain critical for system robustness and long-term adoption: human–agent autonomy balance, testing and verification (including real-time validation), standards, interoperability, and data fusion. These touch on the reliability, safety, and governance of multi-agent systems—essential factors for scaling and deploying them beyond controlled environments.

Several specialised topics—such as stakeholder engagement, safety and risk management, sensor conflict integration, model representation, scalability, programming abstractions, and hardware restrictions—were also recognized. Although not perceived as the most urgent, they underline the interdisciplinary and practical dimensions of agent-based system design.

Finally, the discussion acknowledged the potential for collaboration with other AAMAS research domains, particularly in areas like multi-agent based simulation to complement real-world testing, multi-agent simultaneous localization and mapping (SLAM) in robotics, and commercialization support. These intersections suggest pathways for advancing both foundational research and applied innovation in multi-agent systems engineering.

5 Roadmap: Solutions for the Top Four Challenges

Following the collective identification of key research challenges, participants were divided into groups to address the four most urgent issues for the EMAS community: Explainability, Modeling the Environment, Dynamic Environments, and Interoperable Communication Platforms. Each challenge was reframed as a research question to sharpen its focus and align it with the community's vision. For each question, participants proposed candidate solutions, then voted on the most promising ones based on two dimensions:

Impact the potential transformative effect on the field if achieved.
Effort the estimated complexity and resources required.

This created a four-quadrant prioritisation (High Impact / Low Effort; High Impact / High Effort; Low Impact / Low Effort; Low Impact / High Effort) that helped identify "quick wins" as well as longer-term strategic objectives. The roadmap below integrates these insights into coherent research pathways for each of the four key challenges.

5.1 Explainability

Research question: How can we create explainable systems in heterogeneous environments?

Participants agreed that explainability must support heterogeneous ecosystems that include symbolic agents, learning-based components, and human collaborators. They envisioned explainability mechanisms that are not just post-hoc, but integrated into communication and reasoning processes.

Most promising solutions (in order of community votes):

Text-based communication that can be summarized: Encourage agents to exchange information in structured language formats that can later be summarized for different stakeholders. This approach offers immediate benefits for transparency and debugging.

Formalize reasoning for any action: Develop frameworks that allow agents to produce explicit justifications for their actions—using logical or causal reasoning traces—to enable rigorous inspection of decision processes.

Transparency and trust mechanisms: Embed explainability and transparency into design by providing traceability of decisions and provenance of data, forming a basis for measurable trust.

Stakeholder-tailored explanations: Generate explanations adapted to the audience—whether human operators, other agents, or machines—using context-aware templates or modality switches.

Multimodal explanations: Combine text, visualization, and symbolic traces for richer human–agent interaction, particularly in control and mission settings.

Ranked Solutions:

Solution	**Votes**	**Impact**	**Effort**
Text-based communication that can be summarized	5	High	Low
Formalize reasoning for any action	4	High	High
Transparency and trust mechanisms	3	High	High
Stakeholder-tailored explanations	2	High	Low
Multimodal explanations	3	High	High

Roadmap direction: Short-term work should prototype text-based, summarisable communication within existing MAS frameworks. Mid-term research will integrate formal reasoning traces with stakeholder-aware explanation generation. Long-term, the community should pursue explainability-by-design standards, enabling explainable behavior to be an intrinsic feature of agent architectures rather than an add-on.

Relevant work: [13,16,**?**,**?**,25,27,28,36]

5.2 Modeling the Environment

Research question: How might we model the environment?

Participants discussed what should be modeled, how consistent centralised or distributed representations should be, and how evolving models can be validated. The consensus was that environment modeling is both foundational and multilayered, requiring trade-offs between abstraction, accuracy, and computational cost.

Most promising solutions (in order of community votes):

Define what and how to model, and at which level of abstraction: Propose partitioned spatial and conceptual scaffolds with variable granularity, linked by annotation ontologies for attaching new data, local strategies, and situational knowledge.

Maintain consistency in centralized or consolidated models: Construct multi-level MAS structures where agents operate at different abstraction levels with specific roles and authorities, ensuring coherence across layers.

Validation of evolving models: Employ simulation-based validation, runtime monitoring, and agile experimentation loops to continuously evaluate environmental representations.

Integrate agents into the model itself: *(digital twin paradigm)* Treat agents as both participants and components of the model, enabling co-evolution between the MAS and its digital twin.

Ranked Solutions:

Solution	**Votes**	**Impact**	**Effort**
Define what and how to model, and level of abstraction	7	High	High
Maintain consistency in centralized or consolidated models	7	High	Low
Validation of evolving models	6	High	Medium
Integrate agents into the model itself	3	High	Low

Roadmap direction: Early stages should define hybrid modeling templates and corresponding ontologies. Mid-term actions include developing simulation-based validation pipelines and model-fusion methods. Long-term efforts should produce standardised hybrid environment representations and open testbeds for benchmarking. This efforts are closely related to Digital Twins required in multiple domains.

Relevant work: [15,23,33,34]

5.3 Dynamic Environments

Research question: How might we handle dynamic environments?

This group explored strategies to manage rapid change, uncertainty, and adaptation under resource and time constraints. The core idea was to embed adaptivity within the environment itself, moving from passive to active environments.

Most promising solutions (in order of community votes):

Create an "active environment": Develop environments capable of morphing in response to their own internal processes (e.g., earthquakes, floods). The approach requires a library of pre-built models and situations that can dynamically match sensor input and reconfigure accordingly.

Integrate geotagged data from multiple scouts: Use dynamic, layered maps that evolve over time based on spatially distributed sensor input, enabling situational awareness and faster adaptation.

Combine accountability, human annotation, and task reporting: Encourage hybrid systems where agents and humans cooperatively maintain environmental awareness through periodic status reports and manual updates.

Ranked Solutions:

Solution	Votes	Impact	Effort
Create an "active environment"	9	High	High
Use geotagged sensor data from multiple scouts	8	High	Medium
Combine accountability, human annotations, and task reporting	7	Medium	Medium

Roadmap direction: The short-term goal is to prototype layered environment maps integrating real-time sensor data. Mid-term research should focus on constructing and testing active environment frameworks that dynamically fuse situational models. Long-term efforts should aim for self-adaptive digital twins capable of autonomously reshaping mission models and task allocations.

Relevant work: [21,31]

5.4 Interoperable Communication Platforms

Research question: How might we standardize interoperable communication platforms?

Interoperability was recognized as a critical enabler for collaboration between heterogeneous agents, LLM-powered modules, and human users. Participants highlighted the need for standardization across multiple communication layers.

Most promising solutions (in order of community votes):

Standardized communication protocols: Develop and publish standardized protocol languages and wire encodings to ensure compatibility across platforms and domains, along with extensions for message forwarding and adaptation.

Select and consolidate a shared protocol standard: Reach community consensus on one or a few widely adopted standards, complemented by reference implementations and middleware adapters.

Identify communication layers and semantics: Establish a clear separation of communication layers—(1) Network, (2) Communication, (3) Interaction, (4) Organization, and (5) Coordination—and define their associated semantics.

Ranked Solutions:

Solution	Votes	Impact	Effort
Standardize communication protocols	11	High	Medium
Select and consolidate a shared protocol standard	7	High	High
Identify communication layers and semantics	6	High	High

Roadmap direction: Initial steps involve defining canonical message schemas and minimal interoperable protocols. Medium-term actions will establish translation adapters between MAS frameworks. Long-term, EMAS can lead a community-wide standardization initiative defining interoperable stacks for communication, coordination, and organization layers.

Relevant work: [?,?,?,6,7,?,18,19,?,30,?]

5.5 Cross-Cutting Takeaways

The workshop's participatory roadmap development demonstrated that high-impact, low-effort opportunities—such as text-based explainability and ontology-driven environment models—can deliver early wins for EMAS, while ambitious long-term objectives—such as active environments and protocol standardization across communication layers—offer transformative potential.

By structuring research pathways around clear research questions, measurable milestones, and community-shared evaluation artifacts, EMAS is well-positioned to bridge reasoning and learning paradigms, support agentic AI systems, and maintain leadership in the design and engineering of next-generation multi-agent systems.

6 Discussion

The outcomes of this workshop reveal a vibrant and forward-looking EMAS community that is actively redefining the boundaries of multi-agent systems engineering in the era of agentic AI and large language models. The discussions and roadmap activities underscored both the maturity of EMAS as an engineering discipline and the need for renewed integration with recent advances in machine learning, natural language interaction, and human–autonomy teaming.

6.1 Converging Foundations and Emerging Directions

Participants agreed that the EMAS community has built a strong foundation in formal methods, coordination mechanisms, normative reasoning, and architectural design. These provide stable and reusable engineering abstractions—distinct from the more ad hoc nature of current LLM-based agent systems. However, the emergence of LLMs opens an opportunity to revisit long-standing MAS principles through a new lens. The roadmap discussions reflected a shared understanding that the next phase of EMAS research will involve reconciling symbolic, model-driven reasoning with data-driven adaptivity, enabling agents that can both learn and reason in dynamic, human-centered contexts.

6.2 Key Research Priorities

Across the four prioritized challenges—Explainability, Modeling the Environment, Dynamic Environments, and Interoperability—participants consistently emphasized the importance of trust, transparency, and adaptability. These themes are deeply interconnected: explainable reasoning depends on meaningful models of the environment; adaptability in dynamic situations relies on communication and coordination standards; and interoperability serves as the foundation for integrating diverse agent architectures, learning systems, and human interfaces.

The high-impact, low-effort solutions identified—such as ontology-based modeling—represent practical starting points that can yield early community benefits. Conversely, high-impact, high-effort directions, including the development of active environments and standardized communication layers, point to longer-term collective goals requiring shared infrastructures and collaboration across institutions.

6.3 The Role of EMAS in the LLM and Agentic AI Era

One of the most notable outcomes of the workshop is the recognition that EMAS research provides essential engineering principles for the growing landscape of agentic AI systems powered by LLMs. While LLMs offer powerful capabilities, they lack the explicit organizational, normative, and verification structures that EMAS research has refined over decades. The community therefore has a pivotal role in bringing rigor, safety, and accountability to the emerging generation of generative and autonomous agents.

Several groups highlighted that future EMAS research could focus on LLM-augmented agent architectures, where learning-based modules are embedded within well-defined multi-agent frameworks. Such systems would combine adaptive intelligence with verifiable coordination, providing a clear pathway toward explainable, trustworthy, and human-aligned multi-agent ecosystems.

6.4 Toward a Collaborative Research Agenda

The participatory approach used in this workshop, i.e. combining structured brainstorming, collective prioritization, and impact–effort mapping, proved effective in aligning diverse perspectives from across the EMAS and AAMAS community. The resulting roadmap offers both short-term research actions (e.g., communication templates for explainability, hybrid simulation environments) and strategic long-term goals (e.g., standardization efforts, community-wide testbeds).

Beyond the specific challenges discussed, the workshop reinforced the identity of EMAS as a collaborative, engineering-oriented community—one capable of addressing the complexity of modern autonomous systems through rigorous design, interdisciplinary integration, and open research infrastructure.

7 Conclusion and Next Steps

This paper has presented the outcomes of the EMAS 2025 Roadmapping Workshop, which brought together researchers and practitioners from across the AA-

MAS community to collaboratively identify key challenges and promising directions for the engineering of multiagent systems in the era of LLMs and agentic AI. Through a structured and participatory process inspired by the Lightning Decision Jam methodology, participants jointly explored, categorized, and prioritized issues ranging from explainability and environmental modeling to communication support and system adaptability.

The resulting roadmap highlights a collective vision for the future of multiagent systems engineering—one that emphasizes integrating data-driven and symbolic reasoning, enhancing human–agent collaboration through explainable and trustworthy interfaces, and developing robust methodologies for deployment in dynamic and uncertain environments. The identified challenges and corresponding research directions not only reflect long-standing concerns within the EMAS community but also illustrate its evolving engagement with emerging paradigms in AI.

Beyond the immediate outcomes, the workshop demonstrates the value of community-driven approaches for defining shared research agendas. The roadmap is intended as a living document that can inform future EMAS editions and inspire collaborative efforts across disciplines such as robotics, simulation, and human–AI interaction. Continuing this dialogue will be essential to ensure that the engineering of multiagent systems remains rigorous, transparent, and responsive to the rapidly changing landscape of agentic AI.

We hope that the outcomes of this workshop will serve as the foundation for a living roadmap, to be refined and extended through subsequent EMAS editions and related community initiatives. Continued dialogue among researchers, practitioners, and stakeholders will be essential to ensure that the roadmap remains dynamic, inclusive, and responsive to the rapid evolution of agent and AI technologies.

References

1. Alechina, N., Dastani, M., Logan, B.: Programming norm-aware agents. In: Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 2. pp. 1057–1064. AAMAS '12, International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC (Jun 2012)
2. Bakar, N.A., Selamat, A.: Agent systems verification : Systematic literature review and mapping. Applied Intelligence **48**(5), 1251–1274 (May 2018). `https://doi.org/10.1007/s10489-017-1112-z`
3. Barredo Arrieta, A., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., Garcia, S., Gil-Lopez, S., Molina, D., Benjamins, R., Chatila, R., Herrera, F.: Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. Information Fusion **58**, 82–115 (Jun 2020). `https://doi.org/10.1016/j.inffus.2019.12.012`
4. Bordini, R.H., El Fallah Seghrouchni, A., Hindriks, K., Logan, B., Ricci, A.: Agent programming in the cognitive era. Autonomous Agents and Multi-Agent Systems **34**(2), 37 (May 2020). `https://doi.org/10.1007/s10458-020-09453-y`
5. Bordini, R.H., Hübner, J.F., Wooldridge, M.: Programming Multi-agent Systems in AgentSpeak Using Jason. Wiley-Blackwell, Chichester, 1 edition edn. (Oct 2007)

6. Chopra, A.K., Christie V, S.H.: Communication Meaning: Foundations and Directions for Systems Research. In: Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems. pp. 1786–1791. AAMAS '23, International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC (May 2023)
7. Christie, S.H., Singh, M.P., Chopra, A.K.: Kiko: Programming Agents to Enact Interaction Models. In: Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems. pp. 1154–1163. AAMAS '23, International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC (May 2023)
8. Correa e Silva Fernandes, K.C.: Systémes Multi-Agents Hybrides: Une Approche Pour La Conception de Systémes Complexes. Ph.D. thesis, Université Joseph Fourier- Grenoble 1 (2001)
9. Cossentino, M., Hilaire, V., Molesini, A., Seidita, V. (eds.): Handbook on Agent-Oriented Design Processes. Springer-Verlag, Berlin Heidelberg (2014)
10. Criado, N., Argente, E., Botti, V.: Open issues for normative multi-agent systems. AI Communications **24**(3), 233–264 (Aug 2011). https://doi.org/10.3233/AIC-2011-0502
11. Dennis, L.A., Fisher, M.: Verifiable Autonomous Systems: Using Rational Agents to Provide Assurance about Decisions Made by Machines. Cambridge University Press, Cambridge, United Kingdom ; New York, NY, USA (2023)
12. Dignum, F.: Autonomous agents with norms. Artificial Intelligence and Law **7**(1), 69–79 (Mar 1999). https://doi.org/10.1023/A:1008315530323
13. Dignum, V.: Responsible Artificial Intelligence: How to Develop and Use AI in a Responsible Way (2019)
14. Dix, J., Hindriks, K.V., Logan, B., Wobcke, W.: Engineering Multi-Agent Systems (Dagstuhl Seminar 12342). Dagstuhl Reports **2**(8), 74–98 (2012). https://doi.org/10.4230/DagRep.2.8.74
15. Galland, S., Balbo, F., Gaud, N., Rodriguez, S., Picard, G., Boissier, O.: A multidimensional environment implementation for enhancing agent interaction. In: Bordini, R., Elkind, E. (eds.) 14th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS15). pp. 1801–1802. ACM In-Cooperation, Istanbul, Turkey (May 2015)
16. Gatti, A., Mascardi, V., Ferrando, A.: ChatBDI: Think BDI, Talk LLM. In: Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems. pp. 2541–2543. AAMAS '25, International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC (Jun 2025)
17. Hindriks, K.V., Dix, J.: GOAL: A Multi-agent Programming Language Applied to an Exploration Game. In: Agent-Oriented Software Engineering, pp. 235–258. Springer, Berlin, Heidelberg (2014). https://doi.org/10.1007/978-3-642-54432-3_12
18. Huget, M.P.: Agent Communication. In: Agent-Oriented Software Engineering, pp. 101–133. Springer, Berlin, Heidelberg (2014). https://doi.org/10.1007/978-3-642-54432-3_6
19. Lewis, J., Matson, E.T., Wei, S., Min, B.C.: Implementing HARMS-based indistinguishability in ubiquitous robot organizations. Robotics and Autonomous Systems **61**(11), 1186–1192 (Nov 2013). https://doi.org/10.1016/j.robot.2013.04.001
20. Mascardi, V., Weyns, D., Ricci, A., Earle, C.B., Casals, A., Challenger, M., Chopra, A., Ciortea, A., Dennis, L.A., Díaz, Á.F., El Fallah-Seghrouchni, A., Ferrando, A., Fredlund, L.Å., Giunchiglia, E., Guessoum, Z., Günay, A., Hindriks, K., Iglesias,

C.A., Logan, B., Kampik, T., Kardas, G., Koeman, V.J., Larsen, J.B., Mayer, S., Méndez, T., Nieves, J.C., Seidita, V., Teze, B.T., Varga, L.Z., Winikoff, M.: Engineering Multi-Agent Systems: State of Affairs and the Road Ahead. SIGSOFT Softw. Eng. Notes **44**(1), 18–28 (Oct 2020). https://doi.org/10.1145/3310013.3322175
21. Moradi, H., Iskandar, R., Rodriguez, S., Singh, D., Dugdale, J., Tzempelikos, D., Sfetsos, A., Bakogianni, E., Pavlidi, E., Díaz, J., Ribas, M., Moragues, A., Estrany, J.: Improving evacuation policies through agent-based modeling and stakeholder engagement in hazard-prone areas. International Journal of Disaster Risk Reduction **119**, 105280 (Mar 2025). https://doi.org/10.1016/j.ijdrr.2025.105280
22. Odell, J., Nodine, M., Levy, R.: A Metamodel for Agents, Roles, and Groups. In: Odell, J., Giorgini, P., Müller, J.P. (eds.) Agent-Oriented Software Engineering V, pp. 78–92. No. 3382 in Lecture Notes in Computer Science, Springer Berlin Heidelberg (Jan 2005)
23. Ricci, A., Piunti, M., Viroli, M.: Environment programming in multi-agent systems: An artifact-based perspective. Autonomous Agents and Multi-Agent Systems **23**(2), 158–192 (Sep 2011). https://doi.org/10.1007/s10458-010-9140-7
24. Rodriguez, S., Gaud, N., Galland, S.: SARL: A general-purpose agent-oriented programming language. In: The 2014 IEEE/WIC/ACM International Conference on Intelligent Agent Technology. vol. 3, pp. 103–110. IEEE Computer Society Press, Warsaw, Poland (2014). https://doi.org/10.1109/WI-IAT.2014.156
25. Rodriguez, S., Hilaire, V.: A Methodological Approach for the analysis and design of Human-Swarm interactions based upon Feedback Loops. Expert Systems With Applications **217**, 119482 (2023). https://doi.org/10.1016/j.eswa.2022.119482
26. Rodriguez, S., Hilaire, V., Gaud, N., Galland, S., Koukam, A.: Holonic Multi-Agent Systems. In: Di Marzo Serugendo, G., Gleizes, M.P., Karageorgos, A. (eds.) Self-Organizing Software: From Natural to Artificial Adaptation, pp. 238–263. Self-Organising Software From Natural to Artificial Adaptation - Natural Computing Series, Springer, 1 edn. (2011)
27. Rodriguez, S., Thangarajah, J.: Explainable Agents (XAg) by Design. In: Proceedings of the 2024 International Conference on Autonomous Agents and Multiagent Systems (Blue Sky). pp. 2712–2716. AAMAS '24, Auckland, New Zeland (2024)
28. Rodriguez, S., Thangarajah, J., Davey, A.: Design Patterns for Explainable Agents (XAg). In: Proceedings of the 2024 International Conference on Autonomous Agents and Multiagent Systems. pp. 1621–1629. AAMAS '24, Auckland, New Zeland (2024)
29. Shehory, O., Sturm, A. (eds.): Agent-Oriented Software Engineering. Springer, Berlin, Heidelberg (2014). https://doi.org/10.1007/978-3-642-54432-3
30. Singh, M.P.: Information-driven interaction-oriented programming: BSPL, the blindingly simple protocol language. In: The 10th International Conference on Autonomous Agents and Multiagent Systems - Volume 2. pp. 491–498. AAMAS '11, International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC (May 2011)
31. Swarup, S., Mortveit, H.S.: Live simulations. In: Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems. pp. 1721–1725. AAMAS '20, International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC (2020)
32. Torreño, A., Onaindia, E., Komenda, A., Štolba, M.: Cooperative Multi-Agent Planning: A Survey. ACM Comput. Surv. **50**(6), 84:1–84:32 (Nov 2017). https://doi.org/10.1145/3128584

33. Weyns, D., Michel, F.: Agent Environments for Multi-agent Systems – A Research Roadmap. In: Weyns, D., Michel, F. (eds.) Agent Environments for Multi-Agent Systems IV. pp. 3–21. Springer International Publishing, Cham (2015). https://doi.org/10.1007/978-3-319-23850-0_1
34. Weyns, D., Omicini, A., Odell, J.: Environment as a first class abstraction in multiagent systems. Autonomous Agents and Multi-Agent Systems **14**(1), 5–30 (Feb 2007). https://doi.org/10.1007/s10458-006-0012-0
35. Winikoff, M.: 30 Years of Engineering Multi-Agent Systems: What and Why? | Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems (2024)
36. Yan, E., Burattini, S., Hübner, J.F., Ricci, A.: A multi-level explainability framework for engineering and understanding BDI agents. Autonomous Agents and Multi-Agent Systems **39**(1), 9 (Jan 2025). https://doi.org/10.1007/s10458-025-09689-6

A Case Study Notes

Coordinated Drone Swarm & AI-Supported Command for Disaster

A powerful earthquake has devastated a densely populated urban area. Many streets are blocked by rubble or fire, and communications infrastructure is degraded. Emergency services must act fast to locate survivors, assess structural damage, and coordinate rescue.

To support this, the response team deploys:

- **Robot swarms** to search for survivors, map hazards, and relay information. Semi-autonomous robots with perception (e.g., thermal cameras), capable of local decision-making, path planning, and ad-hoc communication with swarm peers.
- An **AI-enhanced Command & Control (C2) system** to help human operators analyse data, prioritize responses, and coordinate ground units and air assets in real time. A centralized or distributed system that:
  - Aggregates drone and responder data,
  - Identifies high-priority zones,
  - Recommends mission plans to human operators,
  - Monitors team status and suggests role reassignments.
- **First Responders** to assist survivors and contain hazards. Ground units (e.g., firefighter bots or human-agent teams) that receive guidance from the C2 system and may themselves be semi-autonomous.

Consider the following points:

- How to engineer the coordination logic for a robot swarm under uncertainty.
- How to design interaction protocols between drones, the C2 AI system, and human operators and first responders.
- How to integrate machine learning for perception and data summarization, and symbolic reasoning for mission planning and role allocation.
- How to ensure trust, transparency, ethics and accountability in a human-AI team operating in a high-stakes environment.
- How to define the engineering process to develop and deploy the system.

Engineering Questions:

- What kind of agent architecture suits the drones, the C2 system, and human interfaces.
- How to handle distributed task allocation and resilience to partial observability.
- How can LLMs or data-driven models assist in reporting or coordination, without compromising latency or reliability.
- How can we validate decisions and ensure safety in the face of ambiguous data.

Main Components to consider

1. System Architecture.
   - Define a modular architecture showing how drone agents, the AI C2 system, and first responder agents interact.
   - Consider modules for Perception; Planning (deliberative agent logic); Communication; Human-AI Interface; Coordination & Data Fusion.
2. Key Specifications for Each Module.
   - Specify whether each module is a single agent or multi-agent subsystem.
   - Clarify which modules will use (non-exhaustive list):
     - ML / Data-driven approaches.
     - BDI (Belief-Desire/Intentions) or Goal oriented; Automated Planning.
     - Rule-based logic for constraints and fallback behaviour.
   - Identify ethics and accountability constraints:.
     - Explainable decision.
     - Human oversight on actions and fail-safes.
3. Interaction Support.
   - Define communication protocols and technologies (e.g., FIPA ACL, BSPL, MQTT, ROS2).
   - Use or design a shared ontology for status, location, health, mission type, etc.
   - Consider bandwidth constraints and dynamic network topologies (ad hoc mesh).
4. Engineering Process.
   - Outline how the system will be:
     - Tested (e.g., unit testing, BDD) and validated (e.g., in simulation, emulated disaster zones).
     - Monitored during operation for faults and accountability.
     - Development and Deployment, including formal methods, agent programming languages, edge-agent deployment etc.
   - Discuss deployment constraints:
     - Energy and hardware limits on drones.
     - Offline fallback behaviour for disconnected agents.
     - Real-time guarantees and failover protocols.