AI Infrastructure: From Gigastructure to Edge Intelligence with Multi-Agent Systems

Andrea Omicini[1(✉)], Alessandro Ricci[1], and Viviana Mascardi[2]

[1] Department of Computer Science and Engineering (DISI), Alma Mater Studiorum–Università di Bologna, Cesena, Italy
{andrea.omicini,a.ricci}@unibo.it

[2] Department of Computer Science, Bioengineering, Robotics and Systems Engineering (DIBRIS), Università degli Studi di Genova, Genova, Italy
viviana.mascardi@unige.it

Abstract. The rise of generative AI and LLMs is reshaping the global landscape of computational infrastructures. Massive investments in hardware and software are required, raising pressing questions about technological monopolies, digital divides, and the role of public institutions. Recalling the historical evolution of ICT infrastructures, we argue in this paper for a different long-term perspective on AI development—one where agents and MASs serve as the conceptual and technical foundation of scalable, sustainable, and open AI frameworks. Such an approach can help address the challenges of sustaining AI research within public contexts and promote democratic control over AI technologies and applications in the public interest.

Keywords: AI Infrastructure · Multi-Agent Systems · Open Science · Neurosymbolic AI · Internet of Intelligent Things · Digital Divide · Public Interest · AI & Democracy

1 AI Infrastructure Today: Towards Gigastructure

The popularisation of artificial intelligence (AI) is still quite a recent phenomenon, which has basically taken everyone's life by storm. Every contemporary human capable of using a personal device has nowadays access to an unprecedented amount of knowledge typically delivered by LLM-based applications; and generative AI apps have extended the impact of AI technologies even further—to music, images, videos, animation, ...

However, the world-wide accessibility of generative AI applications does not come cheap, and has instead pushed the demand for computational and network resources beyond any previously-imaginable limit. In fact, the training, testing, and deployment of large language models (LLMs) not only require massive datasets, but also ask for the availability of a specialised technological infrastructure developed on an enormous scale. The unprecedented level of the currently-available computational and networking infrastructure – which we refer to as

V. Mascardi and A. Omicini (Eds.): *The Agents Journey*, LNCS 16395, pp. 366–384, 2026.
https://doi.org/10.1007/978-3-032-22940-3_14

AI *gigastructure*, here and henceforth – requires a corresponding and equally-unprecedented level of financial investments. However, few are the organisations that can actually afford that level of investments. Mostly, they are private technology giants, which are rapidly overcoming public organisations—governments and academia: while in 2019 the public sector was controlling more than 60% of the AI supercomputers, in 2025 companies in the private sector own almost 80% of the AI supercomputers—see Fig. 1.

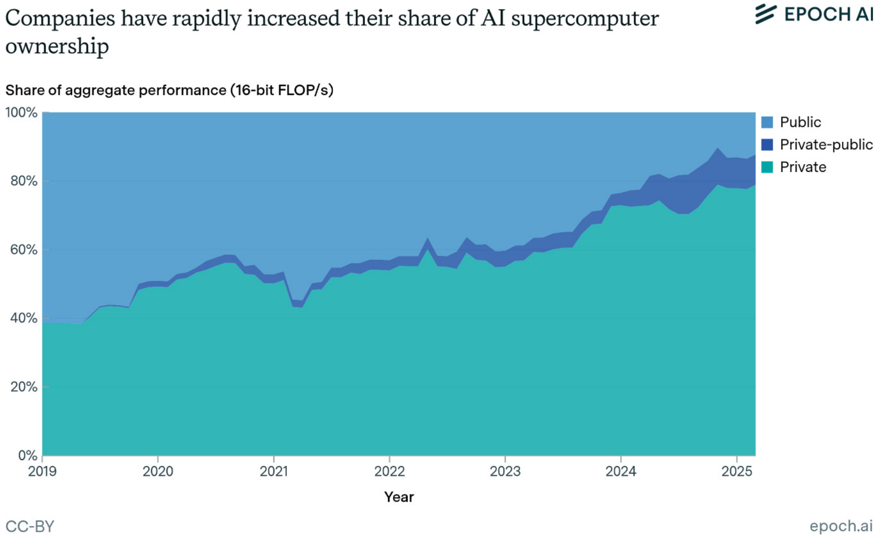

Fig. 1. In the last years, private companies have steadily increased their share of AI supercomputer ownership over public sector [89].

As we are quite accustomed to hear and read bold claims about research from the private sector, whereas the academia keeps on producing most of the actual scientific results, the imbalance in the availability of supercomputing resources for public organisations threatens public research on AI wherever in the world universities are mostly public. In the words of AI scientists Fei-Fei Li, "Public-sector investment in AI is so abysmal. Not a single university today can train a ChatGPT model... academia cannot fully develop its own versions so that it can be used for more open scientific research. That is a problem." [63]

So, the first critical questions becomes: Can public universities keep the pace with private organisation research on AI? Or, instead, should academic researchers just give up, or, leave academia for the big giants in the private sector to pursue their research goals in the field of intelligent systems? Should scientific and technological departments within public universities just leave AI research to rich private organisations, and focus instead on less money-intensive research themes? Should, in the end, the public sector lose control of the AI landscape completely, and leave it all to technological monopolies?

Meanwhile, the centralisation of the AI infrastructure in the hands of a few (private) bodies capable of affording the required financial investments also has a detrimental impact on the geopolitical side. Fact is, private organisations owning the AI gigastructure are primarily based in the United States (75%) and China (15%), with Europe significantly lagging behind (less than 5%), other countries like Japan owning small amounts, and even continents like South America and Africa practically out of the global picture—see Fig. 2.

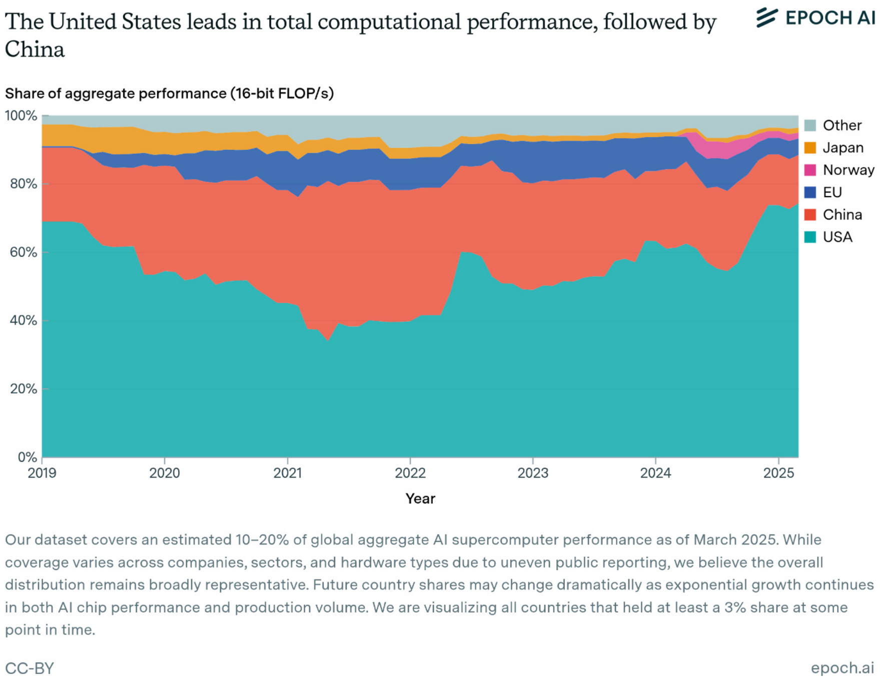

Fig. 2. Geopolitical share of AI supercomputer ownership [89].

Even though European countries are among the richest on the planet, only by considering it as a whole Europe possesses the economic capacity to compete—and then, what about, say, African or South American countries? The whole development of the AI gigastructure is apparently exacerbating the geopolitical imbalance in AI capabilities, which seriously threatens to negatively impact on the potential of development of most countries in the world.

So, the next questions becomes: are we witnessing the emergence of a new kind of digital divide – a huge, geopolitical one –, based no longer on essential information and communication technology (ICT) features, such as Internet access, but rather on access to AI tools and infrastructure? A digital divide where only few superpowers on Earth are able to secure their control over a significant portion of the overall AI gigastructure, quickly gaining industrial and military predominance over the rest of the world?

So, in the remainder of this chapter we try and answer those two fundamental questions about our future as it is going to be predictably shaped by the development of the AI infrastructure: *(i)* Should universities abandon AI research and leave it to the technology monopolists? *(ii)* Is the AI gigastructure going to drive an unbridgeable technology wedge between few superpowers and the rest of the planet?

In spite of the magnitude of the above questions, and – perhaps more critically –, of their economic, social, and geopolitical nature, in this chapter we mostly adopt a scientific and technical viewpoint, showing how agents and multi-agent system (MAS), working as the fundamental reference paradigm for the engineering of intelligent systems [101], can provide us with a less catastrophist perspective on the foreseeable future of AI technology and infrastructure—and of the world to come for all of us.

2 AI Infrastructure Tomorrow: A Perspective

Actually, both critical questions make (more) sense when their common premise is acknowledged: which is, that the future of the AI relies on an overwhelmingly-costly computational & networked infrastructure—which is basically the stage where we found ourselves now, at the end of the first quarter of the 21st century. To put it simply, this means that the future we envision is one where increasingly large datasets are fed to ever-larger models to build massive centralised AI architectures running over a vast distributed computational infrastructure—the AI gigastructure.

However, before we simply accept such a premise, some further considerations need to be made. First of all, we need to address the fundamental issue of technical obsolescence, whose pace in AI is seemingly accelerating: new model architectures, new hardware accelerators, new training techniques emerge frequently, rendering existing infrastructure outdated within a few years. Unlike traditional infrastructure investments – like transportation and urban commodities – AI infrastructure is going to require continual reinvestment, rather than mostly maintenance expenses. On the one hand, this apparently suggests that even fewer governments and private monopolies will be actually able to sustain the costs of an always up-to-date AI infrastructure. On the other hand, early investors may miss dynamics such as a steep reduction in hardware costs accompanied by increasing software efficiency. Conversely, investors who enter at the "right time" could in principle benefit from lower costs, possibly without suffering the drawbacks of a late market entry.

Yet, a deeper question still needs to be asked: Will the need for such an intensive infrastructure persist indefinitely? Or, instead, is it possible that future innovations would make AI development more efficient, more distributable, more sustainable, obliterating the need for a gigastructure for AI applications? Is there the chance that breakthroughs in algorithmic efficiency or new paradigms could significantly alter the forthcoming AI infrastructure landscape?

In order to get the chance to have a glimpse about the future of the AI infrastructure, it may be instructive to take a quick look back at the historical evolution of ICT. First of all, the initial phase of ICT was characterised by centralised mainframes accessible only to a few people and organisations, and with the widespread idea that a handful of exceptionally powerful computers would have provided enough computational power to the whole planet. Even thought Thomas Watson is typically misquoted for technology myth-busting ("I think there's a world market for maybe five computers"), he was still actually referring to a new IBM supercomputer in terms of numbers that would nowadays make us smile.

Instead, as we now know, that stage was followed by the miniaturisation and democratisation of computing through microcomputers, first, and personal devices, then—with industry leaders often unable to see what was coming—such as "There is no reason for any individual to have a computer in his home" (Ken Olsen, CEO of Digital Equipment Corporation, 1977), or Microsoft's Steve Ballmer "There's no chance that the iPhone is going to get any significant market share".

The advent of the Internet, first, and cloud computing, then, continued the trend, decentralising access and functionality across distributed systems. More recently, the rise of edge and fog computing has further pushed intelligence to the periphery, allowing small, cheap, networked devices to perform sophisticated computations.

Overall, the historical pattern is clear: initial centralisation, with a few super-costly, super-powerful computational units doing all the work, is followed by a progressive decentralisation process, with huge numbers of distributed low-power and low-cost computational units. The question now is whether AI infrastructure will follow the very same trend, or, instead, it will keep on growing towards gigastructure stimulated by exponentially growing AI application requests. Could the future of AI be marked by distributed intelligence across edge devices [28], rather than centralised supercomputing hubs? And, in case, which computational paradigms would drive us in that direction, harnessing the complexity of AI applications, and preventing the need for the AI gigastructure?

3 Agents and MAS as the AI Paradigm

3.1 Agentic AI

The very idea of agentic AI exploded in the AI landscape in the last few years [64], based on the fundamental idea that *intelligent systems should be built around "intelligent, autonomous AI agents that can reason, plan, and adapt"* [27]. Two essential things about the concept of agentic AI should be made crystal clear from the very start, before we proceed: *(i)* that idea is not new at all, heading back to the original notions of agents and MAS in the scientific literature [73], *(ii)* what is new there is the predominant notion of agent as an LLM/generative AI container.

So, the novelty of the agentic AI approach is mostly coming from the widespread diffusion and spectacular power of generative AI technologies and frameworks, such as OpenAI's Operator[1] or Stanford's Self-Taught Reasoner (STaR) [102]. On the other hand, the notion that agents and MAS provide intelligent systems engineers with the most expressive and coherent foundation for the design, development, and deployment of AI systems is a long-standing one, in academia, and can be dated back at least to the last decade of the 20th century. At the time, the three original and most important agent conferences – AGENTS (International Conference on Autonomous Agents), ICMAS (International Conference on Multi-Agent Systems), and ATAL (International Workshop on Agent Theories, Architectures, and Languages) – were born, which later merged into the annual International Conference on Autonomous Agents and Multiagent Systems (AAMAS) [59]. More or less at the same time, Workshop on Objects and Agents (WOA) was initiated [41], where ideas like "the deliberative capability of an agent is the natural place for whichever sort of intelligence is needed, in whichever form." [78] were immediately shared by the community.

3.2 MASs for Intelligent Systems

So, multi-agent systems (MASs) have long been a staple in classical AI research, in particular since the emergence of the DAI (distributed artificial intelligence) research area [77]. There, an agent is typically understood as an autonomous software entity capable of perceiving its environment, reasoning, taking actions to achieve goals—yet, a social entity [82], capable of non-trivial interaction within a multi-agent system. In turn, a MAS consists of multiple interacting agents, coordinating their behaviour in some way [31], possibly exhibiting some form of collective intelligence—e.g., [37].

MASs offer a robust and flexible framework for the engineering of intelligent systems [101], since they provide a uniform and coherent way to encapsulate all the many different AI techniques available, from rule-based reasoning to machine learning, within modular and interacting components [87]. Agent-oriented software engineering (AOSE) [39] further extends these ideas by offering methodologies and tools for building large-scale [85], adaptive systems out of agent-based abstractions [46]. Safety requirements of critical applications are a fundamental research topic inside the AOSE community, where formal and semi-formal approaches for testing and verification of agents and MASs have been and are still being developed [8,11,13,14,38,56,60,68,96,100].

What are the basic ideas here, in short? *(i)* The software engineering principle of *encapsulation* applies to intelligence, too, and agents are the fittest containers for intelligence—so, for any AI technique and method. *(ii)* Agents of whatever specific sort can be assigned specific *goals* or tasks that require a specific sort of intelligence in order to be achieved or completed. *(i)* Complex intelligent systems can be built around the *divide et impera* principle, based on multiple intelligent agents each one pursuing its own goal/task with its limited intelligence, knowledge, and capabilities, *coordinating* [31] with other autonomous agents within a MAS so as to achieve the global (intelligent) system goal.

[1] Previously at https://operator.chatgpt.com/, since 17/7/2025 part of ChatGPT as a ChatGPT agent—see https://openai.com/index/introducing-operator/.

So, bottom line, agentic AI is basically just the standard MAS view over intelligent systems engineering, but enhanced with the almighty tools and whistles of contemporary generative AI—which spells distributed artificial intelligence (DAI), in the end [77]: as nowadays any non-trivial computational system is mostly nothing but a distributed one, almost any intelligent system is a distributed intelligent system since its design, and not just for the distributed infrastructural support. Which is why, in the end, AOSE [46] has been synonymous with intelligent systems engineering since its very inception.

What are the potential consequences of that, in terms of the AI infrastructure? Some clues could come from exploring the issues of *edge intelligence* [28]—in the remainder of this section.

3.3 MASs for the Internet of Intelligent Things (IoIT)

The same trajectory of physical distribution and miniaturisation of computational devices that lead us from centralised mainframes to pervasive computing is nowadays observable in the area of AI —where the notion of *pervasive intelligence* [51] is going to play a central role in the next decades. There, in particular, the growing computational power in embedded systems and the expansion of IoT devices are laying the foundation for the Intelligent Things (IoIT). Beyond traditional Internet of Things (IoT), which focuses first and foremost on connectivity and data acquisition, IoIT emphasises autonomous decision-making, exploitation of local knowledge, and situated intelligence.

As discussed elsewhere in this book [28], MASs are ideally suited to work as the sources of abstractions, technologies, and methodologies for the IoIT. Individual agents can be exploited as to handle situated interaction [34,70,98]; to handle local knowledge [47,53,86]; to rationally reason and suitably (re)act in physical and augmented environments [6]. MASs can work as distributed aggregators of knowledge [29] and intelligence, providing for agent coordination [31,71] even on the large scale [85].

Notions such as *micro-intelligence* [34] can then flourish at the device level, reducing the need for constant communication with centralised servers, in the same way as techniques for *localised learning* [36] align well with the DAI interpretation of intelligent systems intrinsically driven by the MAS paradigm. In fact, MAS provide intelligent systems engineers with a structured approach for coordinating device-level intelligence across networks, as "the incorporation of intelligent agents has become widespread to provide smartness within distributed system as well as to design and manage complex scenarios, such as the device-edge- cloud continuum and IoT ecosystem" [65].

So, rather than the AI gigastructure, a likely projection for the future of the AI infrastructure is the one of a *pulverised infrastructure* [54] where smart agents of any sort, suitably coordinated within large-scale MAS, will provide the conceptual, technical, and methodological foundations for the design, development, and deployment of intelligent systems.

Yet: which kind of agents are going to work as the basic components for the forthcoming AI systems?

3.4 MASs for Neurosymbolic Systems

As witnessed by the first years of WOA, in their first decades agent models and technologies mostly developed along two parallel lines: *(i)* rational agents based on a logically-sound architecture [7] – *á la* belief-desire-intention (BDI) [6,44,72,75,76,91] –, mostly focussing on classic AI issues, such as automated reasoning and planning, using symbolic approaches [99], and *(ii)* middleware agents – *á la* JADE [15,16,19,22,23,26,88] –, mostly focussing on the language paradigm shift (from objects to agents), as well as on computational autonomy for harnessing the many issues of distributed systems [25]. The subsequent, overwhelming success of deep learning (DL) led to an intense yet heterogeneous flow of research activities, where DL techniques are integrated with agent-oriented models and technologies in many different ways – e.g., [62,67,96] –, obviously including agent-based modelling & simulation (ABMS)—e.g., [4,66,97].

In the last years, however, the impact of agentic AI techniques has pushed many researchers and practitioners in the area of intelligent system to focus on the almost trivial notion of agents as software containers for DL and generative AI technologies. There, the integration of subsymbolic techniques within the agent frameworks mostly occurs at the syntactic level [1]—so that the expressive power of the agent and MAS abstractions is mostly wasted. Yet, the widespread adoption of subsymbolic techniques has brought into the MAS scenario the same AI issues – transparency, interpretability, trustability [95], accountability, ... – that have led to the recent emergence of explainable artificial intelligence (XAI) as one of the hottest research areas [61]. There, one of the most promising research lines involves the integration of symbolic and subsymbolic techniques into neurosymbolic ones [35], a hybrid approach to intelligent systems that combines the raw power and effective capabilities of subsymbolic approaches with the interpretability and reasoning ability of symbolic systems.

Agents and MAS serve as the ideal architectural layer for this integration. Within an agent, subsymbolic components – e.g., neural networks (NNs) – can handle perception and pattern matching, while symbolic components can manage logic, reasoning, and explanations [2,20,69,90]. Neurosymbolic agents become then the containers where heterogeneous AI paradigms coexist and cooperate fruitfully, making the overall MAS more powerful and effective, and at the same time more understandable, explainable, and trustworthy.

However, if neurosymbolic agents – and MAS whose components are neurosymbolic agents – are the conceptual, technical, and methodological foundation of forthcoming AI systems, one might argue that each of them would require the same huge amounts of computational & networking resources as current generative AI applications do. Yet, first of all, neurosymbolic agents can be designed and implemented so as to exploit subsymbolic techniques – such as generative AI ones – only when actually required by the task at hand, while relying on symbolic techniques – which are typically less computationally intensive – for their general operation.

Also, at the core of AOSE is the notion of agent *society* [78,81], enabling (de)composition techniques like *zooming* [74], which promote fine-grained goal and task (dis)aggregation and stepwise refinement—so that the computational load of every agent can in principle be adjusted to its computational power. For instance, one already quite visible trend in knowledge-intensive intelligent systems is partitioning application knowledge into very specific application subdomains, train some individual agents over each subdomain – which typically do not require huge computational resources, e.g. [40] – and make them cooperate within a MAS covering altogether all the application knowledge. Also, real application domains often do not require the full power of general-purpose AI approaches, so that simplification and reduction techniques can be used to exploit AI techniques on constrained devices–even hyper-constrained devices, as in the case of [3].

In the end, it seems safe to conclude that neurosymbolic agents, too, can be expected to suitably fit the "device-edge-cloud continuum and IoT ecosystem" [65], providing for edge intelligence, specifically, and for the general needs of intelligent systems, as well, without necessarily require the support of the AI gigastructure.

4 The Future of AI Infrastructure: Some Answers

In Sect. 1 we summarise the two fundamental questions to be faced when reasoning about the future of the AI infrastructure:

(i) should universities abandon AI research and leave it to the technology monopolists?, and *(ii)* is the AI gigastructure going to drive an unbridgeable technology wedge between few superpowers and the rest of the planet?

which we can finally try and answer—the latter in Sect. 4.1, the former in Sect. 4.2.

4.1 The Future of AI Infrastructure is Agent-Based

Definitely, the future of AI does not belong exclusively to massive, superpowerful, centralised infrastructures. The most likely trajectory for intelligent systems will presumably include all the possible spectrum of (networked) computational devices, from a number of large supercomputers down to a myriad hyper-constrained embedded devices—balancing the drive in opposite directions, centralisation vs. decentralisation.

Far beyond the current wave of agentic AI, agents and MASs are going to provide the conceptual and technical means to build intelligent systems of any sort. AOSE methods and processes will shape the forthcoming tide of well-designed intelligent systems, allowing engineers to adapt them to the available computational infrastructure—up to a point, as obvious. In particular, MASs can bridge the gap between large-scale infrastructures and decentralised networks, since they can be either deployed on high-performance servers or embedded within edge devices. Neurosymbolic agents, such as *self-explaining agents* [79], demonstrate how a MAS can exploit effective AI subsymbolic techniques while maintaining interpretability, transparency, and trustability—without necessarily depending on a large and costly infrastructure.

Since agents and MASs date far back than the notion of agentic AI, a large, interconnected, worldwide scientific community (which obviously includes the Italian one, grown around twenty-five years of WOA) is already in place, mostly from public institutions, sharing common values such as open science, and prone to listen to the needs of the global citizens of the world—so that the benefits of AI can be shared by everyone on the planet, and the dangers of AI can be limited for everyone at the best of human knowledge. So, for instance, the MAS community has already developed a suitable set of open standards through the FIPA[2] standardisation body [43], which are also actively sustained by the research work—e.g., [30,80]. Also, agent-based notions like *electronic institutions* [12] and *social norms* [92] seemingly provide MAS engineers with a straightforward path towards intelligent systems that are not only powerful and promoting sustainable infrastructure principles, but also more easily aligned with democratic values and ethical norms.

4.2 AI and Public Interest: The Issue of Research

The current dominance of private corporations in AI technology has sparked intense debate, with someone (mostly from private organisations themselves) claiming that universities and public institutions should step back from AI research due to the high costs involved, arguing that the resource requirements has already placed meaningful AI research beyond the reach of academia. That would be a grave mistake. AI technologies already impact every facet of human life, from healthcare to education, from governance to employment: leaving their development solely in the hands of private entities would threaten transparency, accountability, and ethical oversight. Moreover, public norms and regulations, while essential, typically lag behind technological change, and are often insufficient on their own to safeguard public interests—and private ownerships are not historically the best remedy for that problem.

Luckily, public institutions have already reacted constructively and proactively to the risk of losing control of a relevant asset such as AI research and technology: from US[3] to China[4] and Singapore,[5] from EU[6] to UK[7] and Canada,[8] from India[9] to Japan,[10] all the richest countries in the world have implemented

[2] http://www.fipa.org.

[3] https://www.nsf.gov/news/nsf-announces-100-million-investment-national-artificial.

[4] https://www.chinadaily.com.cn/a/202504/18/WS6802358ea3104d9fd38204b5.html.

[5] https://nsstc.niar.org.tw/en/highlights/analysis/2024-07-29.

[6] https://digital-strategy.ec.europa.eu/en/policies/european-approach-artificial-intelligence.

[7] https://www.ukri.org/opportunity/responsible-ai-uk-keystone-projects.

[8] https://ised-isde.canada.ca/site/ai-strategy/.

[9] https://www.pib.gov.in/PressReleasePage.aspx?PRID=2108810.

[10] https://www.meti.go.jp/english/policy/0704_001.pdf.

big public AI funding programs; and many other countries all over the world have done/are doing the same – e.g., in Africa[11] –, at the best of their economic power and sociopolitical context.

However, significantly increasing public funding dedicated to AI is just the first step towards the main goal of preserving democratic oversight and ensuring socially beneficial outcomes. Also, efforts to foster competitive AI research within (public) universities and public research institutions and to promote connections with private sector will be essential, in particular where small and medium-sized enterprises (SME)[12] – whose size typically makes access to AI research problematic – are numerous and relevant. Along this line, Open Science presents a promising path forward, since it emphasises transparency, reproducibility, and collaboration. Yet, it is already evident – even in the academic environment – that any effort toward openness of science should in some way navigate the complexities of intellectual property, data privacy, and (perhaps mostly, today) security. Academic freedom is vital to this effort, offering a foundation for independent exploration and innovation—and maybe for that very reason often under attack.[13]

Nonetheless, academic freedom alone may not suffice in the medium to long term. Alternative paradigms should be pursued, looking for approaches that could reduce reliance on massive computational infrastructure – what we called the AI gigastructure –, and lead us instead towards more flexible, scalable, and inclusive AI architectures. As we discussed in Sect. 3, that is precisely the role that agents and MASs are expected to play in the next decades, working as the main paradigm for the design, development, and deployment of the intelligent systems to come.

5 Conclusion

In conclusion, the dominance of the AI gigastructure in the hands of few private technology giants is not an inevitable future. Drawing from historical trends in ICT as well as from the expressive power of agents and MAS – as well as from more than three decades of academic research –, we argue first of all that the future of AI infrastructure will likely include every size of computational devices, from large supercomputers down to constrained embedded devices.

AOSE offers a paradigm enabling intelligent systems engineers to scale across infrastructural levels—from centralised data centres to distributed edge networks. In the long term, agents and MAS are likely to become the cornerstone of AI infrastructure, offering a balanced response to the challenges of digital divide, public oversight, and technological monopolisation. In that, they represent not only a technical solution but a socio-technical framework for responsible AI.

[11] https://blog.startuplist.africa/articles/ai-revolution-in-africa-2025.
[12] https://eur-lex.europa.eu/eli/reco/2003/361/oj/.
[13] https://www.nytimes.com/2025/08/19/us/politics/trump-universities-financial-penalties.html.

That path, however, strongly depends on the support from public institutions: goals, fundings, and proper regulations from governments; research, open science and technology from public research bodies, such as universities. Whereas many governments worldwide are already pushing in that direction, twenty-five years of research results at WOA – most of which from Italian public universities – suggest that the academic community, too, is ready to accept the challenge.

Acknowledgments. This chapter was partially supported by *(i)* “ENGINES – ENGineering INtElligent Systems around intelligent agent technologies” project, funded by the European Union – Next Generation EU within the framework of the National Recovery and Resilience Plan NRRP – Mission 4 “Education and Research” – Component 2 – Investment 1.1 “National Research Program and Projects of Significant National Interest Fund (PRIN)” – Call PRIN 2022 – D.D. n. 104 of 02/02/2022, under grant number 20229ZXBZM, *(ii)* the “AEQUITAS” project funded by European Union’s Horizon Europe research and innovation programme under grant number 101070363, *(iii)* PNRR âĂŞ M4C2 âĂŞ Investimento 1.3, Partenariato Esteso PE00000013 âĂŞ “FAIR—Future Artificial Intelligence Research” âĂŞ Spoke 8 “Pervasive AI”, funded by the European Commission under the NextGenerationEU programme.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Agiollo, A., Calegari, R., Ciatto, G., Magnini, M., Omicini, A., Sabbatini, F.: Intelligent agents from symbolic to neurosymbolic systems: The quest for integration. In: Mascardi and Omicini [73]
2. Agiollo, A., Ciatto, G., Omicini, A.: Graph neural networks as the copula mundi between logic and machine learning: a roadmap. In: Calegari et al. [33], pp. 98–115. http://ceur-ws.org/Vol-2963/paper18.pdf
3. Agiollo, A., Omicini, A.: Load classification: a case study for applying neural networks in hyper-constrained embedded devices. Appl. Sci. **11**(24) (2021). https://doi.org/10.3390/app112411957. special issue “Artificial Intelligence and Data Engineering in Engineering Applications”
4. Albericci, T., Cecconello, T., Gibertini, A., Vizzari, G.: A curriculum-based reinforcement learning approach to pedestrian simulation. In: Calegari et al. [33], pp. 224–240. http://ceur-ws.org/Vol-2963/paper11.pdf
5. Alderighi, M., Baldoni, M., Baroglio, C., Micalizio, R., Tedeschi, S.: WOA 2024 – 25th Workshop “From Objects to Agents”, CEUR Workshop Proceedings, vol. 3735. Bard, AO, Italy (2024). http://ceur-ws.org/Vol-3735/
6. Amato, A., Di Martino, B., Venticinque, S.: BDI intelligent agents for augmented exploitation of pervasive environments. In: Fortino et al. [58], pp. 81–88. http://ceur-ws.org/Vol-741/ID16_AmatoDiMartinoVenticinque.pdf
7. Ancona, D., Briola, D., Ferrando, A., Martelli, M., Mascardi, V.: 25 years of declarative agent technologies in Italy. In: Mascardi and Omicini [73]
8. Ancona, D., Briola, D., Ferrando, A., Mascardi, V.: MAS-DRiVe: a practical approach to decentralized runtime verification of agent interaction protocols. In: Santoro et al. [94], pp. 35–43. http://ceur-ws.org/Vol-1664/w7.pdf

9. Armano, G., Paoli, F.D., Omicini, A., Vargiu, E.: WOA 2003 – 4th Workshop "From Objects to Agents". Pitagora Editrice Bologna, Villasimius, CA, Italy (2003). http://giuseppevizzari.github.io/WOA-proceedings-archive/woa-2003.html
10. Baldoni, M., Baroglio, C., Bergenti, F., Garro, A.: WOA 2013 – 14th Workshop "From Objects to Agents", CEUR Workshop Proceedings, vol. 1099. Turin, Italy (2013). http://ceur-ws.org/Vol-1099/
11. Baldoni, M., Baroglio, C., Berio, G., Marengo, E.: Declarative representation of curricula models: an LTL-and UML-based approach. In: Baldoni, M., Boccalatte, A., De Paoli, F., Martelli, M., Mascardi, V. (eds.) WOA 2007 – 8th Workshop "From Objects to Agents", pp. 34–41. Seneca Edizioni Torino, Genova, Italy (2007). http://woa07.dibris.unige.it/papers/BaldoniDeclarative.pdf
12. Baldoni, M., Baroglio, C., Capuzzimati, F., Micalizio, R.: Endowing business artifacts with a normative coordination layer. In: De Meo et al. [48], pp. 71–77. http://ceur-ws.org/Vol-1867/w13.pdf
13. Baldoni, M., Baroglio, C., Martelli, A., Patti, V.: Conformance and interoperability in open enviroments. In: De Paoli et al. [49], pp. 151–157. http://ceur-ws.org/Vol-204/P11.pdf
14. Baldoni, M., et al.: Personalization, verification and conformance for logic-based communicating agents. In: Corradini et al. [42], pp. 177–183. http://lia.deis.unibo.it/books/woa2005/papers/25.pdf
15. Baldoni, M., Boella, G., Dorni, M., Mugnaini, A., Grenna, R.: powerJADE: Organizations and roles as primitives in the JADE framework. In: Baldoni et al. [17], pp. 84–92. http://www.pa.icar.cnr.it/woa08/materiali/paper/paper_1.pdf
16. Baldoni, M., Boella, G., Grenna, R.: Modeling organizations and roles using a middleware Jade-based. In: Bergenti [21], pp. 100–107. http://www.ailab.unipr.it/woa09/papers/Baldoni.pdf
17. Baldoni, M., Cossentino, M., De Paoli, F., Seidita, V.: WOA 2008 – 9th Workshop "From Objects to Agents". Seneca Edizioni Torino, Palermo, Italy (2008). http://www.pa.icar.cnr.it/woa08/materiali/Proceedings.pdf
18. Baldoni, M., De Paoli, F., Martelli, A., Omicini, A.: WOA 2004 – 5th Workshop "From Objects to Agents". Pitagora Editrice Bologna, Torino, Italy (2004). http://lia.deis.unibo.it/books/woa2004/atti.pdf
19. Bellifemine, F., Poggi, A., Rimassa, G., Turci, P.: An object-oriented framework to realize agent systems. In: Corradi et al. [41], pp. 52–57. http://giuseppevizzari.github.io/WOA-proceedings-archive/pdfs/woa2000/WOA11.pdf
20. Beretta, D., Monica, S., Bergenti, F.: Capturing a recursive pattern in neural-symbolic reinforcement learning. In: Falcone et al. [55], pp. 17–31. http://ceur-ws.org/Vol-3579/paper2.pdf
21. Bergenti, F.: WOA 2009 – 10th Workshop "From Objects to Agents". Seneca Edizioni Torino, Parma, Italy (2009). http://www.ailab.unipr.it/woa09/papers/
22. Bergenti, F., Caire, G., Gotta, D.: Agents on the move: JADE for android devices. In: Santoro and Bergenti [93], pp. 44–47. http://ceur-ws.org/Vol-1260/paper9.pdf
23. Bergenti, F., Iotti, E., Poggi, A.: Outline of a formalization of JADE multi-agent systems. In: Di Napoli et al. [52], pp. 123–128. http://ceur-ws.org/Vol-1382/paper19.pdf
24. Bergenti, F., Monica, S.: WOA 2019 – 20th Workshop "From Objects to Agents", CEUR Workshop Proceedings, vol. 2404. Parma, Italy (2019). http://ceur-ws.org/Vol-2404/
25. Bergenti, F., Monica, S.: Rethinking software agents as building blocks of software systems. In: Mascardi and Omicini [73]

26. Bergenti, F., Petrosino, G.: Overview of a scripting language for JADE-based multi-agent systems. In: Cossentino et al. [45], pp. 57–62. http://ceur-ws.org/Vol-2215/paper_10.pdf
27. Biswas, A., Talukdar, W.: Building Agentic AI Systems: Create intelligent, autonomous AI agents that can reason, plan, and adapt. Packt Publishing Ltd (2025). https://www.packtpub.com/en-it/product/building-agentic-ai-systems-9781801079273
28. Bouleanu, D.C., Loaiza Carrillo, M.A., Savaglio, C., Bădică, C., Gravina, R., Fortino, G.: From objects to agents, and back to smart objects: software agents for intelligent Internet of Things (IoT) systems. In: Mascardi and Omicini [73]
29. Bouquet, P., Serafini, L., Zanobini, S.: Semantic coordination in systems of autonomous agents: The approach and an implementation. In: Armano et al. [9], pp. 179–186. http://giuseppevizzari.github.io/WOA-proceedings-archive/pdfs/woa2003/17.pdf
30. Briola, D., Locoro, A., Mascardi, V.: Ontology agents in FIPA-compliant platforms: a survey and a new proposal. In: Baldoni et al. [17], pp. 68–75. http://www.pa.icar.cnr.it/woa08/materiali/paper/paper_7.pdf
31. Cabri, G., Leonardi, L., Mariani, S., Zambonelli, F.: Coordination of software agents: models and languages. In: Mascardi and Omicini [73]
32. Calegari, R., Ciatto, G., Denti, E., Omicini, A., Sartor, G.: WOA 2020 – 21st Workshop "From Objects to Agents", CEUR Workshop Proceedings, vol. 2706. Bologna, Italy (2020). http://ceur-ws.org/Vol-2706/
33. Calegari, R., Ciatto, G., Denti, E., Omicini, A., Sartor, G.: WOA 2021 – 22nd Workshop "From Objects to Agents", CEUR Workshop Proceedings, vol. 2963. Bologna, Italy (2021). http://ceur-ws.org/Vol-2963/
34. Calegari, R., Ciatto, G., Mariani, S., Denti, E., Omicini, A.: Logic programming in space-time: The case of situatedness in LPaaS. In: Cossentino et al. [45], pp. 63–68. http://ceur-ws.org/Vol-2215/paper_11.pdf
35. Calegari, R., Ciatto, G., Omicini, A.: On the integration of symbolic and sub-symbolic techniques for XAI: a survey. Intelligenza Artificiale **14**(1), 7–32 (2020). https://doi.org/10.3233/IA-190036. special issue for the Twentieth Edition of the Workshop 'From Objects to Agents'
36. Calegari, R., Denti, E., Mariani, S., Omicini, A.: Towards logic programming as a service: Experiments in tuProlog. In: Santoro et al. [94], pp. 79–84. http://ceur-ws.org/Vol-1664/w14.pdf
37. Casadei, M., Gardelli, L., Viroli, M.: Collective sorting tuple spaces. In: De Paoli et al. [49], pp. 173–180. http://ceur-ws.org/Vol-204/P01.pdf
38. Casadei, M., Viroli, M.: A framework to specify and verify computational fields for pervasive computing systems. In: De Paoli, F., Vizzari, G. (eds.) WOA 2012 – 13th Workshop "From Objects to Agents". CEUR Workshop Proceedings, vol. 892, pp. 72–81. Milano, Italy (2012). http://ceur-ws.org/Vol-892/paper2.pdf
39. Ciancarini, P., Wooldridge, M.J.: Agent-Oriented Software Engineering, Lecture Notes in Computer Science, vol. 1957. Springer, Cham (2001). https://doi.org/10.1007/3-540-44564-1. 1st International Workshop (AOSE 2000), Limerick, Ireland, 10 Jun. 2000. Revised Papers
40. Ciatto, G., Agiollo, A., Magnini, M., Omicini, A.: Large language models as oracles for instantiating ontologies with domain-specific knowledge. Knowl. Based Syst. **310**, 112940:1–22 (2025). https://doi.org/10.1016/j.knosys.2024.112940
41. Corradi, A., Omicini, A., Poggi, A.: WOA 2000 – 1st Workshop "From Objects to Agents", Atti di Congressi, vol. 1195. Pitagora Editrice Bologna, Parma,

Italy (2000). http://giuseppevizzari.github.io/WOA-proceedings-archive/woa-2000.html
42. Corradini, F., De Paoli, F., Merelli, E., Omicini, A.: WOA 2005 – 6th Workshop "From Objects to Agents". Pitagora Editrice Bologna, Camerino, MC, Italy (2005). http://lia.deis.unibo.it/books/woa2005/atti.pdf
43. Cossentino, M., Hopmans, G., Odell, J.: FIPA standardization activities in the software engineering area. In: Armano et al. [9], pp. 71–77. http://giuseppevizzari.github.io/WOA-proceedings-archive/pdfs/woa2003/23.pdf
44. Cossentino, M., et al.: GIMT: a tool for ontology and goal modeling in BDI multi-agent design. In: Santoro and Bergenti [93], pp. 81–88. http://ceur-ws.org/Vol-1260/paper10.pdf
45. Cossentino, M., Sabatucci, L., Seidita, V.: WOA 2018 – 19th Workshop "From Objects to Agents", CEUR Workshop Proceedings, vol. 2215. Palermo, Italy (2018). http://ceur-ws.org/Vol-2215/
46. Cossentino, M., Sabatucci, L., Seidita, V.: Designing agent-oriented systems. In: Mascardi and Omicini [73]
47. Costantini, S., Tocchio, A.: Learning by knowledge exchange in logical agents. In: Corradini et al. [42], pp. 1–8. http://lia.deis.unibo.it/books/woa2005/papers/1.pdf
48. De Meo, P., Postorino, M.N., Rosaci, D., Sarnè, G.M.L.: WOA 2017 – 18th Workshop "From Objects to Agents", CEUR Workshop Proceedings, vol. 1867. Scilla, RC, Italy (2017). http://ceur-ws.org/Vol-1867/
49. De Paoli, F., Di Stefano, A., Omicini, A., Santoro, C.: WOA 2006 – 7th Workshop "From Objects to Agents", CEUR Workshop Proceedings, vol. 204. Catania, Italy (2006). http://ceur-ws.org/Vol-204/
50. De Paoli, F., Manzoni, S., Poggi, A.: WOA 2002 – 3rd Workshop "From Objects to Agents". Pitagora Editrice Bologna, Milano, Italy (2002). http://giuseppevizzari.github.io/WOA-proceedings-archive/woa-2002.html
51. Delicato, F.C., Al-Anbuky, A., Wang, K.I.K.: Editorial: smart cyber-physical systems: toward pervasive intelligence systems. Futur. Gener. Comput. Syst. **107**, 1134–1139 (2020). https://doi.org/10.1016/j.future.2019.06.031
52. Di Napoli, C., Rossi, S., Staffa, M.: WOA 2015 – 16th Workshop "From Objects to Agents", CEUR Workshop Proceedings, vol. 1382. Naples, Italy (2015). http://ceur-ws.org/Vol-1382/
53. Di Nitto, E., Pianciamore, M., Selvini, P.: The role of agents in knowledge management. In: De Paoli et al. [50], pp. 29–34. http://giuseppevizzari.github.io/WOA-proceedings-archive/pdfs/woa2002/22.pdf
54. Domini, D., Farabegoli, N., Aguzzi, G., Viroli, M.: Towards intelligent pulverized systems: a modern approach for edge-cloud services. In: Alderighi et al. [5], pp. 252–270. http://ceur-ws.org/Vol-3735/paper_19.pdf
55. Falcone, R., Castelfranchi, C., Sapienza, A., Cantucci, F.: WOA 2023 – 24th Workshop "From Objects to Agents", CEUR Workshop Proceedings, vol. 3579. Roma, Italy (2023). http://ceur-ws.org/Vol-3579/
56. Ferrando, A., Malvone, V.: Hands-on VITAMIN: a compositional tool for model checking of multi-agent systems. In: Alderighi et al. [5], pp. 157–169. http://ceur-ws.org/Vol-3735/paper_12.pdf
57. Ferrando, A., Mascardi, V.: WOA 2022 – 23rd Workshop "From Objects to Agents", CEUR Workshop Proceedings, vol. 3261. Genova, Italy (2022). http://ceur-ws.org/Vol-3261/

58. Fortino, G., Garro, A., Palopoli, L., Russo, W., Spezzano, G.: WOA 2011 – 12th Workshop "From Objects to Agents", CEUR Workshop Proceedings, vol. 741. Rende, Italy (2011). http://ceur-ws.org/Vol-741/
59. Gini, M., Ishida, T., Castelfranchi, C., Johnson, W.L.: AAMAS 2002: Proceedings of the 1st International Joint Conference on Autonomous Agents and Multiagent Systems: Part 1. Association for Computing Machinery, New York, NY, USA (2002). https://doi.org/10.1145/544862
60. Giordano, L., Martelli, A., Terenziani, P., Bottrighi, A., Montani, S.: A temporal approach to the specification and verification of interaction protocols. In: Corradini et al. [42], pp. 171–176. http://lia.deis.unibo.it/books/woa2005/papers/24.pdf
61. Gunning, D.: Explainable artificial intelligence (XAI). Funding Program DARPA-BAA-16-53, DARPA (2016)
62. Habbash, N., Bottoni, F., Vizzari, G.: Reinforcement learning for autonomous agents exploring environments: an experimental framework and preliminary results. In: Calegari et al. [32], pp. 84–100. http://ceur-ws.org/Vol-2706/paper5.pdf
63. Hammond, G.: AI scientist Fei-Fei Li: "Maths is pretty clean. Humans are messy". Financial Times (2025). https://www.ft.com/content/d5f91c27-3be8-454a-bea5-bb8ff2a85488
64. Huang, K.: Agentic AI: Theories and Practices. Springer, Cham (2025). https://doi.org/10.1007/978-3-031-90026-6
65. Loaiza, M., Savaglio, C., Arijo, N.H., Aloi, G., Fortino, G., Gravina, R.: Agents in software development architectures. In: Falcone et al. [55], pp. 66–77. http://ceur-ws.org/Vol-3579/paper5.pdf
66. Lombardo, G., Pellegrino, M., Poggi, A.: Unsupervised continual learning from synthetic data generated with agent-based modeling and simulation: A preliminary experimentation. In: Ferrando and Mascardi [57], pp. 116–126. http://ceur-ws.org/Vol-3261/paper9.pdf
67. Losapio, G., Minutoli, F., Mascardi, V., Ferrando, A.: Smart balancing of e-scooter sharing systems via deep reinforcement learning. In: Calegari et al. [33], pp. 83–97. http://ceur-ws.org/Vol-2963/paper16.pdf
68. Maggi, P., Sisto, R.: Experiments on formal verification of mobile agent data integrity properties. In: De Paoli et al. [50], pp. 131–136. http://giuseppevizzari.github.io/WOA-proceedings-archive/pdfs/woa2002/15.pdf
69. Magnini, M., Ciatto, G., Omicini, A.: A view to a KILL: knowledge injection via lambda layer. In: Ferrando and Mascardi [57], pp. 61–76. http://ceur-ws.org/Vol-3261/paper5.pdf
70. Mariani, S., Omicini, A.: TuCSoN coordination for MAS situatedness: towards a methodology. In: Santoro and Bergenti [93], pp. 48–57. http://ceur-ws.org/Vol-1260/paper11.pdf
71. Mariani, S., Omicini, A.: Multi-paradigm coordination for MAS: integrating heterogeneous coordination approaches in MAS technologies. In: Santoro et al. [94], pp. 91–99. http://ceur-ws.org/Vol-1664/w16.pdf
72. Mascardi, V., Demergasso, D., Ancona, D.: Languages for programming BDI-style agents: an overview. In: Corradini et al. [42], pp. 9–15. http://lia.deis.unibo.it/books/woa2005/papers/2.pdf
73. Mascardi, V., Omicini, A.: The Agents Journey: Twenty-five Years of Multi-agent Systems at WOA. Lecture Notes in Computer Science – State-of-the-Art Surveys, Springer (2026)

74. Molesini, A., Omicini, A., Ricci, A., Denti, E.: Zooming multi-agent systems. In: Müller, J.P., Zambonelli, F. (eds.) Agent-Oriented Software Engineering VI, Lecture Notes in Computer Science, vol. 3950, pp. 81–93. Springer, Heidelberg (2006). https://doi.org/10.1007/11752660_7. 6th International Workshop (AOSE 2005), Utrecht, The Netherlands, 25–26 Jul. 2005. Revised and Invited Papers
75. Morreale, V., Bonura, S., Francaviglia, G., Centineo, F., Cossentino, M., Gaglio, S.: Reasoning about goals in BDI agents: The PRACTIONIST framework. In: De Paoli et al. [49], pp. 187–194. http://ceur-ws.org/Vol-204/P20.pdf
76. Morreale, V., et al.: PRACTIONIST: a framework for developing BDI agent systems. In: De Paoli et al. [49], pp. 4–5. http://ceur-ws.org/Vol-204/D02.pdf
77. O'Hare, G.M., Jennings, N.R.: Foundations of Distributed Artificial Intelligence. Sixth-Generation Computer Technology, Wiley, Hoboken. hardcover edn. (1996). https://www.wiley.com/en-gb/Foundations+of+Distributed+Artificial+Intelligence-p-9780471006756
78. Omicini, A.: From objects to agent societies: abstractions and methodologies for the engineering of open distributed systems. In: Corradi et al. [41], pp. 29–34. http://giuseppevizzari.github.io/WOA-proceedings-archive/pdfs/woa2000/WOA07.pdf
79. Omicini, A.: Not just for humans: explanation for agent-to-agent communication. In: Vizzari, G., Palmonari, M., Orlandini, A. (eds.) AIxIA 2020 DP —- AIxIA 2020 Discussion Papers Workshop. AI*IA Series, vol. 2776, pp. 1–11. Sun SITE Central Europe, RWTH Aachen University, Aachen, Germany (2020). http://ceur-ws.org/Vol-2776/paper-1.pdf
80. Omicini, A., Ricci, A., Rimassa, G., Viroli, M.: Integrating objective & subjective coordination in FIPA: A roadmap to TuCSoN. In: Armano et al. [9], pp. 85–91. http://giuseppevizzari.github.io/WOA-proceedings-archive/pdfs/woa2003/07.pdf
81. Omicini, A., Ricci, A., Viroli, M.: Artifacts in the A&A meta-model for multi-agent systems. Auton. Agents Multi-Agent Syst. **17**(3), 432–456 (2008). https://doi.org/10.1007/s10458-008-9053-x. special issue on Foundations, Advanced Topics and Industrial Perspectives of Multi-Agent Systems
82. Omicini, A., Tolksdorf, R., Zambonelli, F.: Engineering Societies in the Agents World, Lecture Notes in Artificial Intelligence, vol. 1972. Springer Berlin Heidelberg (2000). https://doi.org/10.1007/3-540-44539-0. 1st International Workshop (ESAW'00), Berlin, Germany, 21 Aug. 2000. Revised Papers
83. Omicini, A., Viroli, M.: WOA 2001 – 2nd Workshop "From Objects to Agents". Pitagora Editrice Bologna, Modena, Italy (2001). http://giuseppevizzari.github.io/WOA-proceedings-archive/woa-2001.html
84. Omicini, A., Viroli, M.: WOA 2010 – 11th Workshop "From Objects to Agents", CEUR Workshop Proceedings, vol. 621. Rimini, Italy (2010). http://ceur-ws.org/Vol-621/
85. Omicini, A., Zambonelli, F.: Coordination of large-scale socio-technical systems: challenges and research directions. In: Di Napoli et al. [52], pp. 76–79. http://ceur-ws.org/Vol-1382/paper11.pdf
86. Passadore, A., Grosso, A., Boccalatte, A.: Indexing enterprise knowledge bases with AgentSeeker. In: Bergenti [21], pp. 69–75. http://www.ailab.unipr.it/woa09/papers/Passadore.pdf
87. Perini, A., Bresciani, P., Giorgini, P., Giunchiglia, F., Mylopoulos, J.: Towards an agent oriented approach to software engineering. In: Omicini and Viroli [83], pp. 72–77. http://giuseppevizzari.github.io/WOA-proceedings-archive/pdfs/woa2001/pdf/19.pdf

88. Petrosino, G., Iotti, E., Monica, S., Bergenti, F.: Prototypes of productivity tools for the Jadescript programming language. In: Calegari et al. [33], pp. 14–28. http://ceur-ws.org/Vol-2963/paper4.pdf
89. Pilz, K., Rahman, R., Sanders, J., Heim, L.: Trends in AI supercomputers (2025). https://epoch.ai/blog/trends-in-ai-supercomputers
90. Pisano, G., Ciatto, G., Calegari, R., Omicini, A.: Neuro-symbolic computation for XAI: towards a unified model. In: Calegari et al. [32], pp. 101–117. http://ceur-ws.org/Vol-2706/paper18.pdf
91. Ribino, P., Cossentino, M., Lodato, C., Lopes, S., Sabatucci, L., Seidita, V.: Ontology and goal model in designing BDI multi-agent systems. In: Baldoni et al. [32], pp. 66–72. http://ceur-ws.org/Vol-1099/paper12.pdf
92. Robol, M., Giorgini, P., Busetta, P.: Applying social norms to implicit negotiation among non-player characters in serious games. In: Santoro et al. [10], pp. 23–28. http://ceur-ws.org/Vol-1664/w5.pdf
93. Santoro, C., Bergenti, F.: WOA 2014 – 15th Workshop "From Objects to Agents", CEUR Workshop Proceedings, vol. 1260. Catania, Italy (2014). http://ceur-ws.org/Vol-1260/
94. Santoro, C., Messina, F., De Benedetti, M.: WOA 2016 – 17th Workshop "From Objects to Agents", CEUR Workshop Proceedings, vol. 1664. Catania, Italy (2016). http://ceur-ws.org/Vol-1664/
95. Sapienza, A., Cantucci, F., Castelfranchi, C., Falcone, R.: Trust evolution in agent and multi-agent systems: a computational modeling perspective. In: Mascardi and Omicini [73]
96. Unniyankal, H., Belardinelli, F., Ferrando, A., Malvone, V.: RMLGym: a formal reward machine framework for reinforcement learning. In: Falcone et al. [55], pp. 1–16. http://ceur-ws.org/Vol-3579/paper1.pdf
97. Vidali, A., Crociani, L., Vizzari, G., Bandini, S.: A deep reinforcement learning approach to adaptive traffic lights management. In: Bergenti and Monica [24], pp. 42–50. http://ceur-ws.org/Vol-2404/paper07.pdf
98. Vizzari, G., Bandini, S.: Coordinated change of state for situated agents. In: Baldoni et al. [18], pp. 69–76. http://giuseppevizzari.github.io/WOA-proceedings-archive/pdfs/woa2004/11.pdf
99. Vozna, A., et al.: Evolution of programming languages in agent systems and the role of computational logic. In: Mascardi and Omicini [73]
100. Yan, E., Burattini, S., Hübner, J.F., Ricci, A.: Towards a multi-level explainability framework for engineering and understanding BDI agent systems. In: Falcone et al. [55], pp. 216–231. http://ceur-ws.org/Vol-3579/paper17.pdf
101. Zambonelli, F., Omicini, A.: Challenges and research directions in agent-oriented software engineering. Auton. Agents Multi-Agent Syst. **9**(3), 253–283 (2004). https://doi.org/10.1023/B:AGNT.0000038028.66672.1e. special Issue: Challenges for Agent-Based Computing
102. Zelikman, E., Wu, Y., Mu, J., Goodman, N.D.: STaR: self-taught reasoner bootstrapping reasoning with reasoning. In: Proceedings of the 36th International Conference on Neural Information Processing Systems (NIPS 2022). Curran Associates Inc., Red Hook, NY, USA (2022). https://dl.acm.org/doi/10.5555/3600270.3601396