

Preface

Silver Jubilee. What does it mean to you? To us, the Silver Jubilee is the definitive certification that something – a couple in love, an organisation, an event – has stood the test of time.

Let us consider, as an example, a scientific workshop. It may consolidate after ten years; it may be mature after fifteen. But only after twenty-five years one can assert that the workshop created a stable community around it, and that such a community in turn created scientific value. After twenty-five years, a workshop – and the community behind it – demonstrates that it had good ideas to share, and that they were shared in an effective way.

WOA is the Workshop on Objects and Agents. It is the most relevant event organised by the Working Group on Autonomous Agents and Multi-Agent Systems (MAS-AIIA) of the Italian Association for Artificial Intelligence, AIIA—initially, with the cooperation of TABOO, the Italian Association for Advanced Technologies based on Object-Oriented concepts. Its first edition dates back to 2000 and took place in Parma, co-organised by Antonio Corradi, Andrea Omicini, and Agostino Poggi. Since then, WOA has been run without interruptions up to now, with editions in Modena (2001), Milano (2002), Villasimius (2003), Torino (2004), Camerino (2005), Catania (2006), Genova (2007), Palermo (2008), Parma (2009), Rimini (2010), Rende (2011), Milano (2012), Torino (2013), Catania (2014), Napoli (2015), Catania (2016), Scilla (2017), Palermo (2018), Parma (2019), Bologna (2020, online), Bologna (2021), Genova (2022), and Roma (2023). The twenty-fifth edition took place in 2024, in the fairy fortress of Bard (Aosta): a great place for celebrating such an important achievement.

WOA welcomes seminal ideas that are worthy to be shared albeit not being necessarily ready for publication in more established venues, and its steering committee has always refrained – and will always refrain – from charging any participation fee: participation in WOA is for free. Being inclusive is indeed the workshop’s distinguishing feature.

Scientists attend WOA because it is the place where new ideas flow, and because it is welcoming. Despite the almost 100% acceptance rate, coherent with its inclusive nature, WOA has contributed to writing the history of agents and multi-agent systems, inside and outside Italy. At the time of writing, the 447 WOA papers indexed by Scopus, going from 2005 to 2024, received 1682 citations leading to an h-index of 17 for the workshop, comparable to the h-index of MAS-related events with similar size at European and international level. While we do not believe that bibliometric indicators are a complete and faithful representation of scientific impact, figures suggest that the ideas presented at WOA have often turned into significant results and shaped the research panorama.

This book tells the history of agents and multi-agent systems (MASs) through the lens of twenty-five WOA editions.

All the authors who published in a significant and continuous way at WOA were invited to prepare a chapter focussed on one specific subject, in order to ensure full

coverage of the WOA topics of interest and minimise overlaps. Most of them accepted enthusiastically, and collaborated to create the book that you are reading.

The book is divided into four parts for a total of 14 chapters. The first part covers *Languages and Models for Computational Agents* and contains three chapters. The second part also contains three chapters and covers *Designing Agent Interaction: Putting MAS Together*. The third part deals with *Software and System Engineering with Agents and MAS* and includes four chapters, while the fourth and last part discusses *Agents and MAS for Artificial Intelligence*, with four chapters as well.

Part I of the book is about agent languages and models, and starts with a chapter by Alina Vozna, Andrea Monaldini, Stefania Costantini, Giovanni De Gasperis, Pierangelo dell'Acqua, Andrea Formisano, and Andrea Rafanelli who discuss the *Evolution of Programming Languages in Agent Systems and the Role of Computational Logic*. The chapter illustrates the evolution over time of the programming languages adopted for implementing agents and MASs and presents the authors' DALI framework, some of its applications, and how Epistemic Logic may be exploited for modelling group dynamics.

In the second chapter, Domenico Rosaci and Giuseppe M. L. Sarnè aim at *Looking for the Best Partners* for interaction and cooperation. They illustrate different strategies, including trust, reliability, reputation, cyphering, and commitment, that provide means for assessing credibility, trustworthiness, and identity, for reducing information asymmetry between agents, and for preventing malicious activities.

The third chapter complements the second by analysing *Trust Evolution in Agent and Multi-Agent Systems: A Computational Modeling Perspective*. In this chapter, Alessandro Sapienza, Filippo Cantucci, Cristiano Castelfranchi, and Rino Falcone discuss the usage and evolution of the trust concept within MASs, with the objectives of understanding how and why the choice to utilise trust was made, and of outlining the future directions to be pursued and the challenges to be taken into account.

Part II of the book moves from the single-agent dimension to the multi-agent one. The first chapter in this part is by Matteo Baldoni, Cristina Baroglio, and Roberto Micalizio who overview *Interaction Protocols: from AUML to Social Commitments, from Artifacts to BSPL*. The chapter reviews the diverse approaches in the literature ranging from message passing (FIPA ACL, KQML, and AUML) to coordination via the environment (CARTAgO, JaCaMo), from normative multi-agent organisations (MOISE) to information-based interaction protocols (BSPL).

A close focus on *Coordination of Software Agents: Models and Languages* is provided in the second chapter of Part II, written by Giacomo Cabri, Letizia Leonardi, Stefano Mariani, and Franco Zambonelli. Besides presenting agent coordination models and languages defined in the last 25 years, the authors discuss how agent coordination mechanisms have been applied in autonomic computing, self-organising systems, energy management, and autonomous vehicles.

Gianluca Aguzzi, Roberto Casadei, Danilo Pianini, and Mirko Viroli address *Self-Organisation with Aggregate Computing: A Reflection Under the Lenses of Multi-Agent Systems Engineering*. In this chapter, the authors overview results in the aggregate computing field and elaborate on how they can contribute to the design of distributed digital signs for environment-mediated agent coordination, collective plans for agents and swarms, and reinforcement learning frameworks for many-agent systems.

Part III is devoted to approaches for engineering agents and multi-agent systems, and starts with a chapter by Federico Bergenti and Stefania Monica who propose *Rethinking Software Agents as Building Blocks of Software Systems*. The chapter examines the implications of adopting cognitive agents within mainstream software development practice and presents a cognitive-agent level as a theoretical foundation for such an adoption.

The chapter *Designing Agent-Oriented Systems* by Massimo Cossentino, Luca Sabatucci, and Valeria Seidita explores the rich trajectory of Agent-Oriented Software Engineering (AOSE), delineates its path from conceptual foundations to contemporary applications, and presents a forward-looking perspective on AOSE, outlining innovative directions and potential research areas.

Alfredo Garro elaborates on *Agent-Based Computing for Science and Engineering: Back to the Future and Beyond...* The chapter explores the historical roots of agent-based approaches, tracing their evolution and renaissance in contemporary research, and highlights recent developments that extend the boundaries of agent-based approaches, pushing the paradigm “beyond” traditional confines.

The last chapter of Part III covers *Simulation and Agent-Based Simulation at WOA: 25 of Years of Simulation-Related Contributions, and an Outlook of What Awaits Us*. Therein, Giuseppe Vizzari and Daniela Briola discuss, through examples presented at WOA, how simulations may be carried out to demonstrate the working of an agent-based model or software system, to calibrate or evaluate it. They also show how agent-based approaches are particularly in tune with some of the characteristics of complex systems.

Part IV addresses agents and MASs for artificial intelligence (AI). The first chapter in this part presents *25 Years of Declarative Agent Technologies in Italy*. The authors – Davide Ancona, Daniela Briola, Angelo Ferrando, Maurizio Martelli, and Viviana Mascardi – illustrate declarative technologies for agents and MASs and how they can be exploited in the specification of the agents and MAS architecture, in the specification of their conversational flow, in their implementation, and in the representation of the agents’ knowledge.

With the rise, consolidation, and success of machine learning approaches and neural networks, a discussion on *Intelligent Agents from Symbolic to Neurosymbolic Systems: The Quest for Integration* inside the WOA community is needed. Andrea Agiollo, Roberta Calegari, Giovanni Ciatto, Matteo Magnini, Andrea Omicini, and Federico Sabbatini explore the development of rational agents and integration with machine learning techniques, discussing their transition from pure symbolic to subsymbolic and neurosymbolic systems.

Daniel-Costel Bouleanu, Marco Loaiza, Claudio Savaglio, Costin Bădică, Raffaele Gravina, and Giancarlo Fortino delve into the Internet of Things (IoT) and Agent-Based Computing (ABC) duo, in their chapter *From Objects to Agents, and Back to Smart Objects: Software Agents for Intelligent Internet of Things (IoT) Systems*. There, they discuss how ABC enables the development of advanced IoT systems from the design to the implementation and simulation phases, and present application examples, technologies and IoT platforms that make use of agents.

The book ends with the vision of Andrea Omicini, Alessandro Ricci, and Viviana Mascardi on the *AI Infrastructure: From Gigastructure to Edge Intelligence with Multi-Agent Systems*. In this chapter, the authors observe how the rise of generative AI and large

language models is reshaping the global landscape of computational infrastructures, and argue 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.

As this book shows, the WOA community has worked with continuity and coherence, and has a very clear vision of the road ahead. What the book cannot show is what the WOA community is for real, and which values it promotes. WOA has lovingly acted as a ‘nursery’ for many PhD students who gave their first talk there, and are now renowned scientists; we respect each other and are happy to meet, making the atmosphere at the workshop warm and relaxed; we look forward to moments of conviviality that are always cheerful and rewarding.

Actually, WOA’s long-lasting success is rooted in long-lasting friendship among community members that goes far beyond common scientific interests. Thanks to our tight-knit community many bright years are waiting for us, full of ideas, engaging discussions, and joy to share.

September 2025

Viviana Mascardi
Andrea Omicini