Preface

The Workshop “From Objects to Agents” (WOA) is the reference event for Italian researchers active in the Agents and Multi-Agent Systems research domain. Since its very first edition in 2000, located in Parma (Italy), WOA was conceived as a meeting occasion for researchers and practitioners from the working group on MAS of AIxIA and from the TABOO association (Advanced Technologies Based on Concepts from Object-Orientation). After that, WOA was held on a yearly basis in many different Italian locations, from north to south, gaining conspicuous success and succeeding in gathering researchers and practitioners from various research fields, thanks to its format. Despite stemming from an Italian initiative, WOA is an international workshop where presenters and participants exchange opinions and discuss ongoing works in a friendly yet rigorous setting.

The 24th edition of the workshop was held on November 7–8, 2023 in Roma. During these two days, 22 speakers joined the workshop, almost 40 attendees joined in presence, and 2 joined remotely. The travel and accommodation of two students were partly supported by AIxIA. This edition was structured in a keynote speech, a round table with three speakers, and five technical sessions. The five technical sessions hosted the presentation of 18 papers, 17 of which are collected in this virtual volume published by CEUR. The topics discussed in the papers covered some of the hottest topics laying under the umbrella of “Cognition: an outdated goal or a permanent challenge for the new AI paradigm?”, as requested by the call for papers. The choice of this theme was deliberate.

The challenge that AI systems are representing for society is of enormous importance: their competence and generality of support is growing and consequently their pervasiveness in each area of our existence: individual and social. They will change not only our material and institutional world/reality but our mind.

The new AI systems that are spreading respond to the problems submitted to them, with results that are often completely indistinguishable from those that humans themselves can achieve. At the same time, the real intelligent capabilities of these models are debated and questioned. In fact, even if we are comforted by some of their amazing results, they appear to proceed according to strictly data-driven approaches, which are far from obvious with respect to the performances expressed. It therefore becomes interesting to ask whether an approach purely guided by the relationships that emerge from data is sufficient to reproduce some typical attitudes of humans in the exercise of their intelligent performances and interactions. And what sense does it make today to build architectures, top-down, that make use of "cognitive modeling". Also in the Agents and MAS domain, as we know, this direction has often been followed mainly using the symbolic approach. In fact, between the 80s and 90s of the last century, the cross-fertilization with the social and cognitive sciences and the encounter with the MAS favored a more cognitively plausible vision of agents (for example with the Belief Desire intention (BDI) model), as autonomous systems endowed with the ability to represent, planning and social action.

With the development and significant acquisitions of machine learning, especially deep learning, data-driven science is imposing a substantial reversal of scientific investigation. We can say that BigData Science determines a dominance of the predictive scientific function over the explanatory one. In fact, the adoption of statistical models and data-mining techniques for the extraction of information from data leads to the identification of regular trends, which in some cases allow the projection on the future of the same trends without an investigation into the producing causes, also neglecting the possibility of indicating and suggesting the guidelines or principles to which the results of these models should comply. A direct consequence is the scarce role of social and cognitive mechanisms: the more the ability to extract

regularities from vast databases is refined, the less the individual mechanisms that generate them are investigated, since these mechanisms are judged irrelevant in the prediction of phenomena.

The “Fabio Bellifemine” keynote speech was given by Cristiano Castelfranchi (Institute of Cognitive Sciences and Technologies, ISTC-CNR), who discussed the topic “Eliminativism for AI? No longer ‘cognitive representations’ and ‘minds’ in artificial intelligences?”. In his discussion, the speaker reconstructed the “eliminativistic” attack to “Cognitive Science” aimed at eliminating notions and models of “mental representations” (“beliefs”, “intentions”, reasoning, ..) and then “mind”. He has investigated whether this attack concerns just in psychology and Cognitive Science or it will propagate - thanks to the “Generative AI” approach - in AI, against the AI models of “intelligence”, of autonomous “Agents”, of their reciprocal “influencing” based on changing the other’s beliefs and goals; of their cooperation based on “mind reading” and goal-adoption. “Cognitive modeling” is on the one side a crucial “scientific” (not just technical) contribution of AI to Cognitive Science, on the other side it is a fundamental approach/instrument for building “artificial” intelligences, Agents, and hybrid societies. Moreover “transparency” - a crucial condition for human trust in Agents and robots - is not transparency of the underlying “algorithms”; it is transparency (i) of the hidden “values” and interests; and (ii) of the “minds” of the agent (explainability): the “reasons” of its choice and behavior; its assumptions/beliefs and its goals. That is, “mind reading”. The speaker argued that we need to build real minds/intelligences not simply to fake them; “anthropomorphized” machines.

As far as it concerns the round table, the three participants Mario De Caro (Università Roma3, Dipartimento di Filosofia), Chiara Gallese (Università di Torino, Dipartimento di Giurisprudenza), Dario Guarascio (Università Sapienza, Dipartimento di Diritto ed Economia delle attività produttive), together with the discussant Rino Falcone (Institute of Cognitive Sciences and Technologies, ISTC-CNR), argued about “Progress and Complexity of AI: A Cross-Sector Exploration”.

The 18 papers collected in this issue were organized into five thematic sessions. The final versions also include the outcomes of the discussions that followed the presentations at the workshop. The authors’ contributions cover extremely relevant research areas that include (i) Learning and Advancements in Multi-Agent Systems (ii) Cognition, trust and reputation (iii) Theoretical approaches (iv) Applicative studies (v) The role of Explainability.

In the end, the Organizing Scientific Committee gratefully thanks all those who, with their work and their enthusiasm, have contributed to the success of this edition of WOA: the members of the Program Committee, the Institute of Cognitive Sciences and Technologies (ISTC-CNR), the National Research Council (CNR), AIxIA, the speakers of the workshop sessions, the speakers of the “Fabio Bellifemine” keynote speech and of the round table, and all collaborators who participated in the organization.

Overall, they would like to thank the lively, creative, and sometimes volcanic community that has been regularly meeting for 24 years at the workshop.

The Organizing Committee

Rino Falcone, Cristiano Castelfranchi, Alessandro Sapienza, and Filippo Cantucci