

Preface

Artificial Intelligence (AI) continues to evolve at a remarkable pace, shaping industries, societies, and our daily lives. As AI systems become more pervasive and autonomous, their transparency, trustworthiness, and societal impact raise urgent questions. The need for explainable, responsible, and fair AI is at the forefront of both academic research and public discourse, especially as new applications challenge our understanding of ethics, accountability, and human-AI collaboration.

The International Workshop on EXplainable, Trustworthy, and Responsible AI and Multi-Agent Systems (EXTRAAMAS) has, since its inception in 2019, served as a vibrant forum for advancing research and fostering dialogue on these critical issues. Now in its seventh edition, EXTRAAMAS 2025 convened in Detroit, USA, bringing together a diverse community of researchers and practitioners dedicated to the foundational and applied challenges of explainable AI (XAI) and multi-agent systems.

This year's workshop built on the strong legacy of previous editions, with a renewed focus on four key areas: (i) XAI Fundamentals, (ii) Applied XAI showcases, (iii) Cross-Disciplinary Perspectives, and (iv) Responsible and Fair AI.

The 2025 program featured two distinguished keynote speakers who embody the workshop's interdisciplinary ethos: Simona Tiribelli gave a talk on "Substantial Fairness in AI Ethics: The Path Forward," challenging us to move beyond procedural definitions of fairness and to embrace more substantive, justice-oriented approaches to AI design—particularly in sensitive domains like healthcare. Munindar P. Singh presented "Explanations in Responsible Autonomy: Trustworthy AI, Norm Deviation, and Consent," offering philosophical and practical insights into how explanations underpin responsible AI behavior and trustworthiness, drawing from both theory and legal case studies.

This year, EXTRAAMAS received 23 submissions, reflecting the growing interest and dynamism in the field. Each paper underwent a rigorous single-blind peer-review process with between 3 and 5 reviews per submission, ensuring the highest standards of quality and relevance. The 14 accepted papers (13 full length and one short), presented in person in Detroit, span a wide range of topics—from foundational advances in XAI to innovative applications in law, ethics, and human-AI interaction. With the authors' consent, these contributions are made available on the EXTRAAMAS website, furthering our commitment to open and accessible knowledge dissemination.

We extend our heartfelt gratitude to the special track chairs, publicity chairs, and the Program Committee for their dedication and expertise. We are also deeply thankful to the authors, presenters, and participants whose enthusiasm and engagement make EXTRAAMAS a unique and inspiring event year after year.

As AI continues to transform our world, the EXTRAAMAS community remains committed to shaping a future where AI systems are not only powerful and intelligent but also transparent, trustworthy, and aligned with human values.

June 2024

Davide Calvaresi
Amro Najjar
Andrea Omicini
Kary Främling