

Explainable, Trustworthy, and Responsible AI and Multi-Agent Systems

8th International Workshop, EXTRAAMAS 2026, Paphos, Cyprus, May 25–26, 2026,
Revised Selected Papers

This book constitutes the refereed post-conference proceedings of the 8th International Workshop on Explainable, Trustworthy, and Responsible AI and Multi-Agent Systems, EXTRAAMAS 2026, held in Paphos, Cyprus during May 25–26, 2026. The 20 full papers and 2 short paper included in this book were carefully reviewed and selected from 34 submissions. They were organized in topical sections as follows: Foundations & Limits of Explainability in Agentic Systems; Trust, Ethics, and Human-Centred Explainable AI; Methods for Explainability: Argumentation, Learning, and Hybrid AI; Systems and Applications of Explainable Multi-Agent Systems.

This book constitutes the refereed post-conference proceedings of the 8th International Workshop on Explainable, Trustworthy, and Responsible AI and Multi-Agent Systems, EXTRAAMAS 2026, held in Paphos, Cyprus during May 25–26, 2026. The 20 full papers and 2 short paper included in this book were carefully reviewed and selected from 34 submissions. They were organized in topical sections as follows: Foundations & Limits of Explainability in Agentic Systems; Trust, Ethics, and Human-Centred Explainable AI; Methods for Explainability: Argumentation, Learning, and Hybrid AI; Systems and Applications of Explainable Multi-Agent Systems.



62,05 €

57,99 € (zzgl. MwSt.)

*vorbestellbar, Erscheinungstermin ca.
November 2026*

Artikelnummer: 9783032384355

Medium: Buch

ISBN: 978-3-032-38435-5

Verlag: Springer

Erscheinungstermin: 20.11.2026

Sprache(n): Englisch

Auflage: Erscheinungsjahr 2026

Serie: Lecture Notes in Artificial
Intelligence

Produktform: Kartoniert

Seiten: 439

Format (B x H): 155 x 235 mm

