



Brazil July 1-3, 2026

Conference Website: <https://xaiworldconference.com/2026>

About the Conference

<https://xaiworldconference.com/2026/call-for-papers/>

4th World Conference on eXplainable Artificial Intelligence

Call for papers

Artificial intelligence has undergone a significant shift in focus, with a growing emphasis on designing and developing intelligent systems that are both interpretable and explainable. This is due to the complexity of the models, built from data, and the legal requirements imposed by various national and international parliaments. This has been echoed both in the research literature and the press, attracting scholars worldwide and a lay audience. An emerging field in AI is Explainable Artificial Intelligence (xAI), which is devoted to producing intelligent systems that enable humans to understand their inferences, assessments, predictions, recommendations, and decisions. Initially devoted to designing post-hoc methods for explainability, eXplainable Artificial Intelligence (xAI) is rapidly expanding its boundaries to neuro-symbolic methods for producing self-interpretable models. Research has also shifted the focus to the structure of explanations and human-centred Artificial Intelligence, as the ultimate users of interactive technologies are humans.

The World Conference on Explainable Artificial Intelligence is an annual event that aims to bring together researchers, academics, and professionals, promoting the sharing and discussion of knowledge, new perspectives, experiences, and innovations in the field of Explainable Artificial Intelligence (xAI). This event is multidisciplinary and interdisciplinary, bringing together academics and scholars of different disciplines, including Computer Science, Psychology, Philosophy, Law and Social Science, to mention a few, and industry practitioners interested in the practical, social and ethical aspects of the explanation of the models emerging from the discipline of Artificial Intelligence (AI).

The conference organisation encourages submissions related to explainable AI and contributions from academia, industry, and other organisations discussing open challenges or novel research approaches related to the explainability and interpretability of AI systems. Topics include, and are not limited to:

Technical methods for XAI

Action Influence Graphs Agent-based explainable systems Ante-hoc approaches for interpretability Argumentative-based approaches for xAI Argumentation theory for xAI Attention mechanisms for xAI Automata for explaining RNN models Auto-encoders & latent spaces explainability Bayesian modelling for interpretability Black-boxes vs white-boxes Case-based explanations for AI systems Causal inference & explanations Constraints-based explanations Decomposition of NNET-models for XAI Deep learning & XAI methods Defeasible reasoning for explainability Evaluation approaches for XAI-based systems Explainable methods for edge computing Expert systems for explainability Sample-centric and dataset-centric explanations Explainability of signal processing methods Finite state machines for explainability Fuzzy systems & logic for explainability Graph neural networks for explainability Hybrid & transparent black box modelling Interpreting & explaining CNN Networks Interpretable representational learning Explainability & the Semantic Web Model-specific vs model-agnostic methods Neuro-symbolic reasoning for XAI

Natural language processing for explanations Ontologies & taxonomies for supporting XAI Pruning methods with XAI
Post-hoc methods for explainability Reinforcement learning for enhancing XAI Reasoning under uncertainty for explanations
Rule-based XAI systems Robotics & explainability Sample-centric & Dataset-centric explanations Self-explainable methods for XAI Sentence embeddings to xAI semantic features Transparent & explainable learning methods
User interfaces for explainability Visual methods for representational learning XAI Benchmarking XAI methods for neuroimaging & neural signals XAI & reservoir computing

Ethical Considerations for XAI

Accountability & responsibility in XAI Addressing user-centric requirements for XAI Trade-off model accuracy & interpretability
Explainable Bias & fairness of XAI systems Explainability for discovering, improving, controlling & justifying Moral Principles & dilemma for XAI
Explainability & data fusion Explainability/responsibility in policy guidelines Explainability pitfalls & dark patterns in XAI
Historical foundations of XAI Moral principles & dilemma for XAI Multimodal XAI approaches Philosophical consideration of synthetic explanations Prevention/detection of deceptive AI explanations Social implications of synthetic explanations
Theoretical foundations of XAI Trust & explainable AI The logic of scientific explanation for/in AI Expected epistemic & moral goods for XAI XAI for fairness checking XAI for time series-based approaches

Psychological Notions & concepts for XAI

Algorithmic transparency & actionability Cognitive approaches for explanations Cognitive relief in explanations
Contrastive nature of explanations Comprehensibility vs interpretability Counterfactual explanations Designing new explanation styles Explanations for correctability Faithfulness & intelligibility of explanations
Interpretability vs traceability explanations Interestingness & informativeness Irrelevance of probabilities to explanations
Iterative dialogue explanations Local vs. global interpretability & explainability Local vs global interpretability & explainability
Methods for assessing explanations quality Non-technical explanations in AI systems Notions and metrics of/for explainability
Persuasiveness & robustness of explanations Psychometrics of human explanations Qualitative approaches for explainability
Questionnaires & surveys for explainability Scrutability & diagnosis of XAI methods Soundness & stability of XAI methods

Social examinations of XAI

Adaptive explainable systems Backwards & forward-looking responsibility forms to XAI Data provenance & explainability
Explainability for reputation Epistemic and non-epistemic values for XAI Human-centric explainable AI
Person-specific XAI systems Presentation & personalization of AI explanations for target groups Social nature of explanations

Legal & administrative considerations of/for XAI

Black-box model auditing & explanation Explainability in regulatory compliance Human rights for explanations in AI systems
Policy-based systems of explanations The potential harm of explainability in AI Trustworthiness of XAI for clinicians/patients
XAI methods for model governance XAI in policy development XAI for situational awareness/compliance behavior

Safety & security approaches for XAI

Adversarial attacks explanations Explanations for risk assessment Explainability of federated learning Explainable IoT malware detection Privacy & agency of explanations XAI for Privacy-Preserving Systems

XAI techniques of stealing attack & defence XAI for human-AI cooperation XAI & models output confidence estimation

Applications of XAI-based systems

Application of XAI in cognitive computing Dialogue systems for enhancing explainability

Explainable methods for medical diagnosis Business & Marketing XAI systems for healthcare

Explainable methods for HCI Explainability in decision-support systems Explainable recommender systems

Explainable methods for finance & automatic trading systems Explainability in agricultural AI-based methods Explainability in transportation systems

Explainability for unmanned aerial vehicles Explainability in brain-computer interfaces Interactive applications for XAI

Manufacturing chains & application of XAI Models of explanations in criminology, cybersecurity & defence XAI approaches in Industry 4.0

XAI systems for health-care XAI technologies for autonomous driving XAI methods for bioinformatics

XAI methods for linguistics/machine translation XAI methods for neuroscience XAI models & applications for IoT

XAI methods for XAI for terrestrial, atmospheric, & ocean remote sensing XAI in sustainable finance & climate finance XAI in bio-signals analysis

Submission

Submitted manuscripts must be novel and not substantially duplicate existing work. Manuscripts must be written using Springer's Lecture Notes in Computer Science (LNCS) in the format provided here. Latex and word files are admitted; however, the former is preferred. All submissions and reviews will be handled electronically. The conference has a no dual submission policy, so submitted manuscripts must not be currently under review at another publication venue.

Articles must be submitted using the easy-chair platform here.

While registering on the platform, the contact author must provide the following information: paper title, all author names, affiliations, postal address, e-mail address, and at least three keywords.

The conference will not require a strict page number, as we believe authors have different writing styles and would like to produce scientific material differently. However, the following types of articles are admitted:

full articles between 14 and 24 pages (including references)

short articles between 10 and 14 pages (including references)

Full articles should report on original and substantial contributions of lasting value, and the work should concern the theory and/or practice of Explainable Artificial Intelligence (xAI). Moreover, manuscripts showcasing the innovative use of xAI methods, techniques, and approaches and exploring the benefits and challenges of applying xAI-based technology in real-life applications and contexts are welcome. Evaluations of proposed solutions and applications should be commensurate with the claims made in the article. Full articles should reflect more complex innovations or studies and have a more thorough discussion of related work. Research procedures and technical methods should be presented sufficiently to ensure scrutiny and reproducibility. We recognise that user data may be proprietary or confidential; therefore, we encourage sharing (anonymized, cleaned) data sets, data collection procedures, and code. Results and findings should be communicated clearly, and implications of the contributions for xAI as a field and beyond should be explicitly discussed.

Shorter articles should generally report on advances that can be described, set into context, and evaluated concisely. These articles are not 'work-in-progress' reports but complete studies focused on

smaller but complete research work, simple to describe. For these articles, the discussion of related work and contextualisation in the wider body of knowledge can be smaller than that of full articles.

Appendixes and supplemental material

Appendices and supplemental material must be placed within the article and the maximum number of pages mentioned above.

Special session articles

The article submitted to the special sessions follows the submission procedure of the main track and must be submitted via EasyChair, as mentioned above. The types of articles admitted are full and shorter, as described above. The authors of an article to be associated with a special session must select the name of such special session in the list of topics in EasyChair, along with other relevant topics.

Authors commit to reviewing

By submitting to the conference, each senior manuscript author (holding at least a PhD) volunteers to be added to the pool of potential PC members/reviewers for the conference and may be asked to review manuscripts. This does not apply to authors who have already agreed to contribute to the conference in some capacity (e.g., as PC/SPC members of the main conference or special tracks, area chairs, or members of the organising committee) and authors who are not qualified to be on the programme committee.

Ethical & Human Subjects Considerations

The conference organisers expect authors to discuss the ethical considerations and the impact of the presented work and/or its intended application, where appropriate. Additionally, all authors must comply with the ethical standards and regulatory guidelines associated with human subjects research, including the use of personally identifiable data and research involving human participants. Manuscripts reporting on human subjects research must include a statement identifying any regulatory review the research is subject to (and identifying the form of approval provided) or explaining the lack of required review.

Submission and publication of multiple articles

Each author is limited to a combined maximum of 4 submissions to the main conference track, and authors may not be added or deleted from papers following submission.

Use of Generative AI

Generative AI models such as LLMs, including ChatGPT, BARD, LLaMA, and similar, are against the criteria for authorship of scientific manuscripts submitted and published in the conference. If authors use any of these tools while writing their manuscript, they assume full responsibility for all content. This includes verifying its correctness and assessing plagiarism of any part of their work. Suppose the text generated by the above generative AI models is the subject of scientific inquiry as part of the manuscript's methodology or analysis. In that case, it must be adequately described, documented and made explicit in the paper.

Topics of Interest

2 topics

Research papers are invited in, but not limited to, the following areas:

- [Uncategorized](#)

- [Artificial Intelligence & Machine Learning](#)

Important Dates

Jul 01

Conference Date

July 1-3, 2026

© 2026 CallForPaper.org - All Rights Reserved

Providing global research dissemination and event management services.