

Preface

Christoph Düsing^{1,*}, Daniele Pala², Somayeh Kafaie³, Björn Eskofier⁴, Bruno Casella⁵, Ioanna Miliou⁶, Mirko Polato⁷, Enayat Rajabi⁸ and Giovanna Nicora⁹

¹Center for Cognitive Interaction Technology (CITEC), Bielefeld University, Germany

²University of Bergamo, Italy

³Saint Mary's University, Canada

⁴Institute of AI in Medicine, LMU Munich, Germany

⁵Department of Computer Science, University of Turin, Italy

⁶Department of Computer and Systems Sciences, Stockholm University, Sweden

⁷Department of Computer Science, University of Turin, Italy

⁸Cape Breton University, Canada

⁹University of Pavia, Italy

The three workshops represented in these joint proceedings were held on Friday, July 10, 2026, at the University of Ottawa in Ottawa, Canada, and were co-located with the 24th International Conference on Artificial Intelligence in Medicine (AIME 2026). AIME 2026 took place from July 7 to July 10, 2026. The Doctoral Consortium and tutorials were held on July 7, the main conference took place on July 8 and 9, and the workshops concluded the conference program on July 10.

The present joint proceedings bring together seven short papers from three AIME 2026 workshops. Although the workshops address different methodological perspectives, they share a common interest in developing trustworthy, evidence-based, and practically deployable artificial intelligence methods for medicine and public health.

1st International Workshop on Multicentric and Privacy-preserving Learning in Healthcare

The *1st International Workshop on Multicentric and Privacy-preserving Learning in Healthcare* (MULTIPLi Health) addressed the opportunities and challenges arising from the increasing availability of heterogeneous healthcare data distributed across multiple institutions. Although such data can support the development of more robust and generalizable artificial intelligence systems, privacy, governance, and regulatory constraints frequently prevent the direct sharing of patient-level data.

Multicentric and privacy-preserving learning paradigms, including federated learning, distributed learning, and federated analytics, provide mechanisms for collaborative analysis without requiring participating institutions to exchange their raw data. The workshop therefore focused on the methodological, technical, and practical challenges involved in designing trustworthy artificial intelligence systems for multicentric healthcare environments.

Topics of interest included federated learning architectures for healthcare; federated analytics and distributed statistical analysis; learning under institutional heterogeneity and class imbalance; differential privacy, secure aggregation, and secure multiparty computation; multicentric, multimodal, and longitudinal healthcare data; evaluation, benchmarking, and reproducibility; infrastructures and platforms for distributed healthcare artificial intelligence; clinical deployment; robustness and fairness; and explainable artificial intelligence in federated and distributed settings.

Joint Proceedings of the AIME 2026 Workshops: 1st International Workshop on Multicentric and Privacy-preserving Learning in Healthcare, Foundation Models for Public Health and Epidemiology: From Promise to Practice, and First International Workshop on Knowledge Graphs for Health, July 10, 2026, Ottawa, Canada

*Corresponding Editor.

✉ cdusing@techfak.uni-bielefeld.de (C. Düsing)



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

The submission deadline was May 22, 2026, and authors were notified of the acceptance decisions on June 5, 2026. The workshop received 11 submissions. Each submission was reviewed by at least two members of the dedicated program committee. Following review and discussion, four papers were accepted, corresponding to an acceptance rate of approximately 36.4%. Three of the accepted papers are included in this proceedings volume. The remaining accepted paper was not included because the authors did not submit a camera-ready version for publication.

Foundation Models for Public Health and Epidemiology: From Promise to Practice

The workshop on *Foundation Models for Public Health and Epidemiology: From Promise to Practice* examined the opportunities and limitations of large language models, multimodal generative systems, and other foundation models in population-level health applications. While these models have demonstrated considerable potential in biomedical and clinical research, their translation to public-health practice remains limited by challenges related to generalization, interoperability, evaluation, governance, and equity.

The workshop brought together researchers and practitioners to critically assess the current state of the field, identify barriers to real-world deployment, and formulate a research agenda connecting foundation-model capabilities with public-health requirements. It combined invited talks, contributed presentations, and structured discussions directed toward evaluation frameworks, hybrid modeling approaches, interoperable deployment, and community resources.

Topics of interest included foundation models for public-health surveillance and outbreak detection; multimodal generative models for epidemiological and environmental data; generalization and fairness across population groups; interoperability with public-health information systems and standards; evaluation and validation frameworks tailored to population-level impact; ethical, governance, and equity considerations; integration of agent-based models with generative artificial intelligence; and practical deployments within public-health agencies.

The final submission deadline was May 17, 2026, and authors were notified of the acceptance decisions on May 25, 2026. The workshop received six submissions. Each submission was reviewed by at least two members of the dedicated program committee. Following review and discussion, five papers were accepted, corresponding to an acceptance rate of approximately 83.3%. Three of the accepted papers are included in this proceedings volume. The remaining two accepted papers were not included because the authors withdrew their papers or did not submit camera-ready versions for publication.

First International Workshop on Knowledge Graphs for Health

The *First International Workshop on Knowledge Graphs for Health* (KG4Health) addressed the intersection of knowledge graphs and artificial intelligence in biomedical and healthcare applications. Although modern generative artificial intelligence systems can process large amounts of clinical text and biomedical information, concerns regarding factual accuracy, explainability, interoperability, and fragmented data sources continue to limit their safe use in high-stakes medical settings.

Knowledge graphs provide structured and interpretable representations that can ground artificial intelligence systems in established biomedical and clinical knowledge. The workshop offered a forum for discussing how such representations can support evidence-based, transparent, interoperable, and clinically reliable artificial intelligence.

Topics of interest included knowledge-graph and large-language-model integration; neuro-symbolic artificial intelligence and retrieval-augmented generation for medical question answering; automated construction and enrichment of medical knowledge graphs; interoperability with established ontologies and standards such as SNOMED CT, UMLS, and FHIR; graph-based explanations for clinical predictions; and applications including drug repurposing, rare-disease diagnosis, and personalized treatment pathways.

The final submission deadline was May 11, 2026, and authors were notified of the acceptance decisions on May 25, 2026. The workshop received four submissions. Each submission was reviewed by at least two members of the dedicated program committee. Following review and discussion, one paper was accepted and is included in this proceedings volume, corresponding to an acceptance rate of 25%.

Acknowledgments

The proceedings editors thank the organizers of AIME 2026 and the University of Ottawa for hosting the conference and its workshops. We are especially grateful to the members of the three program committees for contributing their time and expertise to the review process. We also thank all authors, invited speakers, presenters, and workshop participants for their contributions and for the discussions that shaped the three workshop programs.

MULTIPLi Health Organizing Committee

- Christoph Düsing, Bielefeld University, Germany
- Bruno Casella, University of Turin, Italy
- Ioanna Miliou, Stockholm University, Sweden
- Björn Eskofier, LMU University Hospital Munich, Germany
- Mirko Polato, University of Turin, Italy

MULTIPLi Health Program Committee

- Alessio Barbaro Chisari, University of Catania
- Amr Abourayya, TU Dortmund University and University Hospital Essen
- Francesco Piccialli, University of Naples Federico II
- Gabriele Tolomei, Sapienza University of Rome
- Gianluca Mittone, University of Turin
- Marco Aldinucci, University of Turin
- Panagiotis Papapetrou, Stockholm University
- Roberto Esposito, University of Turin
- Sara Bouchenak, INSA

Foundation Models for Public Health and Epidemiology Organizing Committee

- Daniele Pala, University of Bergamo, Italy
- Giovanna Nicora, University of Pavia, Italy

Foundation Models for Public Health and Epidemiology Program Committee

- John H. Holmes, University of Pennsylvania, USA
- Arianna Dagliati, University of Pavia, Italy
- Enrico Longato, University of Padua, Italy
- Ettore Lanzarone, University of Bergamo, Italy

KG4Health Organizing Committee

- Reza (Enayat) Rajabi, Cape Breton University, Canada
- Somayeh Kafaie, Saint Mary's University, Canada

KG4Health Program Committee

- Fatemeh Bagheri, Saint Mary's University, Canada
- Zhendong Sha, Cleveland Clinic Research, USA
- Ahmet Soylu, Kristiania University College, Norway
- Hannah Kim, Temple University, USA
- Akhil Chaudhary, Cape Breton University, Canada
- Faezeh Ensan, Toronto Metropolitan University, Canada
- Christophe Debruyne, University of Liège, Belgium
- Majid Ziaratban, Saint Mary's University, Canada
- Jaber Rad, Dalhousie University, Canada
- Allard Oelen, German National Library of Science and Technology, Germany
- Hande Küçük McGinty, Kansas State University, USA
- Zaynab Mousavian, Emory University, USA