

1 CONFERENCE COMMITTEE

General Chair

Pierre Marquis (University of Artois, CNRS, CRIL–Institut
Universitaire de France, France)

Program Chairs

Magdalena Ortiz (TU Wien, Austria)

Maurice Pagnucco (The University of New South Wales,
Australia)

Local Arrangement Chairs

Long Tran-Thanh (University of Warwick, UK)

Thanh Van Dinh (East Asia University of Technology, Viet-
nam)

Van Dao Hong (East Asia University of Technology, Viet-
nam)

Tran Cao Son (New Mexico State University, USA)

Phi Le Nguyen (Hanoi University of Science and Technol-
ogy, Vietnam)

Workflow Chairs

Sanja Lukumbuza (TU Wien, Austria)

Anthony Tompkins (UNSW, Australia)

Demo Session Chair

Mario Alviano (University of Calabria, Italy)

Diversity and Inclusion Session Chairs

Shqiponja Ahmetaj (TU Wien, Austria)

Daxin Liu (University of Edinburgh, UK)

Renata Wassermann (Universidade de São Paulo, Brazil)

Doctoral Consortium Chairs

Camille Bourgaux (CNRS, ENS, Paris, France)

Johannes P. Wallner (TU Graz, Austria)

Funding and Sponsorship Chairs

Gianluca Cima (Sapienza University of Rome, Italy)

Zeynep Saribatur (TU Wien, Austria)

Publicity Chairs

Yazmin Angelica Ibañez-García (Cardiff University, UK)

Stavros Vassos (Helvia.ai, Greece).

Recent Published Research Track Chairs

Yongmei Liu (Sun Yat-sen University, China)

Mantas Šimkus (TU Wien, Austria)

Special Track: Reasoning, Learning and Decision- Making Track Chairs

Sebastijan Dumančić (TU Delft, The Netherlands)

Jean Christoph Jung (TU Dortmund University, Germany)

Masoumeh Mansouri (University of Birmingham, UK)

Special Track: KR in the Wild Track Chairs

Birte Glimm (University of Ulm, Germany)

Matti Järvisalo (University of Helsinki, Finland)

Tutorial and Workshop Chairs

Elena Botoeva (University of Kent, UK)

Abhaya Nayak (Macquarie University, Australia)

Video Track Chairs

Anni-Yasmin Turhan (Paderborn University, Germany)

Nico Potyka (Cardiff University, UK)

Webmaster

Tobias Geibinger (TU Wien, Austria)

2 AREA CHAIRS

Frank Wolter (University of Liverpool, UK)
Stefan Woltran (TU Wien, Austria)

Natasha Alechina (Open University Netherlands and Utrecht University, The Netherlands)
Leila Amgoud (CNRS,IRIT, France)
Ofer Arieli (Tel-Aviv Academic College, Israel)
Franz Baader (TU Dresden and SCADS.AI Dresden/Leipzig, Germany)
Chitta R. Baral (Arizona State University, USA)
Salem Benferhat (CRIL CNRS University of Artois, France)
Leopoldo Bertossi (Universidad San Sebastian, Santiago, Chile)
Meghyn Bienvenu (LaBRI—CNRS and University of Bordeaux, France)
Richard Booth (Cardiff University, UK)
Gerhard Brewka (University of Leipzig, Germany)
Pedro Cabalar (University of A Coruña, Spain)
Bernardo Cuenca Grau (University of Oxford, UK)
Giuseppe De Giacomo (University of Oxford, UK)
James Delgrande (Simon Fraser University, Canada)
Thomas Eiter (TU Wien, Austria)
Hélène Fargier (IRIT-CNRS and University of Toulouse, France)
Hector Geffner (RWTH Aachen University, Germany)
Victor Gutierrez Basulto (Cardiff University, UK)
Andreas Herzog (IRIT-CNRS, France)
Anthony Hunter (UCL London, UK)
Tomi Janhunen (Tampere University, Finland)
Gabriele Kern-Isberner (TU Dortmund, Germany)
Sébastien Konieczny (CRIL CNRS University of Artois, France)
Gerhard Lakemeyer (RWTH Aachen University, Germany)
João Leite (Universidade Nova de Lisboa, Portugal)
Yves Lesperance (York University, Canada)
Leonid Libkin (University of Edinburgh, UK and Relation-
alAI, France)
Alessio Lomuscio (Imperial College London, UK)
Carsten Lutz (Leipzig University, Germany)
Thomas Meyer (University of Cape Town and CAIR, South Africa)
Marie-Laure Mugnier (University of Montpellier, France)
Aniello Murano (University of Naples Federico II, Italy)
Rafael Peñaloza (University of Milano-Bicocca, Italy)
Pavlos Peppas (University of Patras, Greece)
Francesca Rossi (IBM T.J. Watson Research Lab, New York, USA)
Sebastian Rudolph (TU Dresden, Germany)
Uli Sattler (University of Manchester, UK)
Torsten Schaub (University of Potsdam, Germany)
Renate A. Schmidt (University of Manchester, UK)
Luciano Serafini (Fondazione Bruno Kessler, Italy)
Guillermo R. Simari (Universidad Nacional del Sur in Bahia Blanca, Argentina)
Michael Thielscher (The University of New South Wales, Australia)
Matthias Thimm (University of Hagen, Germany)
Francesca Toni (Imperial College London, UK)
Serena Villata (CNRS, France)
Renata Wassermann (Universidade de São Paulo, Brazil)

3 PROGRAM COMMITTEES

Main Track

- Erman Acar (University of Amsterdam, The Netherlands)
Shqiponja Ahmetaj (TU Wien, Austria)
Yamine Ait Ameer (IRIT-CNRS-Toulouse INP, France)
Gianvincenzo Alfano (University of Calabria, Italy)
José J. Alferes (Universidade Nova de Lisboa, Portugal)
Mario Alviano (University of Calabria, Italy)
Benjamin Aminof (TU Wien, Austria)
Alessandro Antonucci (Dalle Molle Institute for Artificial Intelligence USI-SUPSI, Switzerland)
Daan Apeldoorn (University Medical Center Mainz, Germany)
Theofanis Aravanis (University of the Peloponnese, Greece)
Carlos Areces (Universidad Nacional de Córdoba, Argentina)
Dörthe Arndt (TU Dresden, Germany)
Alessandro Artale (Free University of Bozen-Bolzano, Italy)
Mutsunori Banbara (Nagoya University, Japan)
Ringo Baumann (Leipzig University, Germany)
Peter Baumgartner (Data61, CSIRO and Australian National University, Australia)
Bartosz Bednarczyk (TU Dresden, Germany and University of Wrocław, Poland)
Francesco Belardinelli (Imperial College London, UK)
Stefano Bistarelli (University of Perugia, Italy)
Fernando Bobillo (University of Zaragoza, Spain)
Alexander Bochman (Holon Institute of Technology, Israel)
Thomas Bolander (Technical University of Denmark, Denmark)
Piero A. Bonatti (Università di Napoli Federico II, Italy)
Blai Bonet (University Pompeu Fabra, Spain)
Stefan Borgwardt (Technische Universität Dresden, Germany)
Zied Bouraoui (CRIL CNRS and University of Artois, France)
Pierre Bourhis (CRISAL, CNRS, University of Lille and INRIA of University of Lille)
Loris Bozzato (University of Insubria, Italy)
Maxime Buron (LIMOS; Clermont Auvergne University, France)
Marco Calautti (University of Milan, Italy)
David Carral (LIRMM, Inria, University of Montpellier, CNRS, France)
Giovanni Casini (National Research Council, Italy)
Balder ten Cate (University of Amsterdam, The Netherlands)
Federico Cerutti (University of Brescia, Italy)
Lukas Chrpá (Czech Technical University in Prague, Czechia)
Agata Ciabattoni (TU Wien, Austria)
Gianluca Cima (Sapienza University of Rome, Italy)
Jens Claßen (Roskilde University, Denmark)
Oana Cocarascu (King's College London, UK)
Marco Console (University of Edinburgh, UK)
Carl Corea (University of Koblenz, Germany)
Augusto B. Corrêa (University of Basel, Switzerland)
Fabio G. Cozman (Universidade de São Paulo, Brazil)
Carlos Damásio (Universidade Nova de Lisboa, Portugal)
Fabio Aurelio D'Asaro (University of Verona, Italy)
Mehdi Dastani (Utrecht University, The Netherlands)
Alexis de Colnet (TU Wien, Austria)
Ronald de Haan (University of Amsterdam, The Netherlands)
Warren Del-Pinto (University of Manchester, UK)
Antonio Di Stasio (University of Oxford, UK)
Federica Di Stefano (TU Wien, Austria)
Martin Dieguez (University of Angers, France)
Clare Dixon (University of Manchester, UK)
Dragan Doder (Utrecht University, The Netherlands)
Wolfgang Faber (Alpen-Adria-Universität Klagenfurt, Austria)
Jorge Fandinno (University of Nebraska Omaha, USA)
Eduardo Fermé (University of Madeira, Portugal)
Christian Fermüller (TU Wien, Austria)
Oliver Fernández Gil (TU Dresden, Germany)
Flavio Ferrarotti (Software Competence Center Hagenberg, Austria)
Johannes K. Fichte (Linköping University, Sweden)
Tommaso Flaminio (Artificial Intelligence Research Institute (IIIA-CSIC), Spain)
Giorgos Flouris (Foundation for Research and Technology - Hellas (ICS FORTH), Greece)
Gerhard Friedrich (Universität Klagenfurt, Austria)
Maurice Funk (Leipzig University, Germany)
Marco Garapa (University of Madeira, Portugal)
Luca Geatti (University of Udine, Italy)
Marcel Gehrke (University of Hamburg, Germany)
Tobias Geibinger (TU Wien, Austria)
Manolis Gergatsoulis (Ionian University, Greece)
Alessandro Gianola (INESC-ID/Instituto Superior Técnico, Universidade de Lisboa, Portugal)
Nina Gierasimczuk (Technical University of Denmark, Denmark)
Laura Giordano (Università degli Studi del Piemonte Orientale, Italy)
François Goasdoué (University of Rennes, France)
Lluís Godó Lacasa (Artificial Intelligence Research Institute IIIA - CSIC, Spain)
Lucía Gómez Álvarez, (Université Grenoble Alpes, Inria, CNRS, France)
Ricardo Gonçalves (Universidade Nova de Lisboa, Portugal)
Sergio Greco (University of Calabria, Italy)
Jasmin Grosinger (Örebro University, Sweden)
Paolo Guagliardo (University of Edinburgh, UK)
Jonas Philipp Haldimann (TU Wien, Austria and University of Cape Town, South Africa)
Markus Hecher (MIT, USA)
Jesse Heyninck (Open Universiteit, The Netherlands)
Till Hofmann (RWTH Aachen University, Germany)
Martin Homola (Comenius University of Bratislava, Slovakia)
Aaron Hunter (British Columbia Institute of Technology, Canada)
Giovambattista Ianni (University of Calabria, Italy)
Daniela Inglezan (Miami University, USA)

Jianmin Ji (University of Science and Technology of China, China)
 Magdalena Kacprzak (Bialystok University of Technology)
 Shakil M. Khan (University of Regina, Canada)
 Toryn Q. Klassen (University of Toronto and Vector Institute, Canada)
 Matthias Knorr (Universidade Nova de Lisboa, Portugal)
 Boris Konev (University of Liverpool, UK)
 Roman Kontchakov (Birkbeck, University of London, UK)
 Juha Kontinen (University of Helsinki, Finland)
 Patrick Koopmann (Vrije Universiteit Amsterdam, The Netherlands)
 Egor V. Kostylev (University of Oslo, Norway)
 Alisa Kovtunova (TU Dresden, Germany)
 Francesco Kriegel (Technische Universität Dresden, Germany)
 Markus Krötzsch (TU Dresden, Germany)
 Isabelle Kuhlmann (University of Hagen, Germany)
 Oliver Kutz (Free University of Bozen-Bolzano, Italy)
 Jürgen Landes (Ludwig Maximilian University of Munich, Germany)
 Davide Lanti (Free University of Bozen-Bolzano, Italy)
 Francesco Leofante (Imperial College London, UK)
 Vladimir Lifschitz (University of Texas, USA)
 Fangzhen Lin (Hong Kong University of Science and Technology, Hong Kong)
 Sebastian Link (University of Auckland, New Zealand)
 Francesca Alessandra Lisi (University of Bari “A. Moro”, Italy)
 Jorge Lobo (ICREA and University Pompeu Fabra, Spain)
 Anela Lolic (TU Wien, Austria)
 Emiliano Lorini (IRIT, France)
 Timothy S. Lyon (TU Dresden, Germany)
 Sagar Malhotra (TU Wien, Austria)
 Quentin Manière (Leipzig University, Germany)
 Marco Manna, (University of Calabria, Italy)
 Marco Maratea (University of Calabria, Italy)
 Lucas Martinelli Tabajara (Runtime Verification)
 Nicolas Maudet (Sorbonne Université, France)
 Andrea Mazzullo (Free University of Bozen-Bolzano, Italy)
 Kevin McAreevey (University of Bristol, UK)
 John-Jules Meyer (Utrecht University, The Netherlands)
 Loizos Michael (Open University of Cyprus, Cyprus)
 Munyque Mittelmann (University of Naples Federico II, Italy)
 Angelo Montanari (University of Udine, Italy)
 Leora Morgenstern (SRI International, USA)
 Filip Murlak (University of Warsaw, Poland)
 Nysret Musliu (TU Wien, Austria)
 Matthias Nickles (University of Galway, Ireland)
 Juan Carlos Nieves (Umeå University, Sweden)
 Andreas Niskanen (University of Helsinki, Finland)
 Johannes Oetsch (Jönköping University, Sweden)
 Cem Okulmus (Umeå University, Sweden)
 Nicola Olivetti (Aix-Marseille University (LIS), France)
 Andrea Orlandini (National Research Council, Italy)
 Piotr Ostropolski-Nalewaja (University of Wrocław, Poland and TU Dresden, Germany)
 Wassila Ouerdane (Paris-Saclay University, France)
 Ana Ozaki (University of Oslo, Norway)
 Aline Paes (Universidade Federal Fluminense, Brazil)
 Nina Pardal (University of Sheffield, UK)
 Xavier Parent (TU Wien, Austria)
 Francesco Parisi (University of Calabria, Italy)
 Timothy J. Parker (IRIT, France)
 Fabio Patrizi (Sapienza University of Rome, Italy)
 Guilherme Paulino-Passos (Imperial College London, UK)
 David Pearce (Universidad Politécnica de Madrid, Spain)
 Wojciech Penczek (Institute of Computer Science, PAS)
 Elise Perrotin (AIST, Japan)
 Laurent Perrussel (IRIT, University Toulouse Capitole, France)
 Liat Peterfreund (Hebrew University of Jerusalem, Israel)
 Ron Petrick (Heriot-Watt University, UK)
 Andreas Pieris (University of Edinburgh, UK and University of Cyprus, Cyprus)
 Ramón Pino Pérez (CRIL-CNRS, University of Artois, France)
 Sylwia Polberg-Riener (Cardiff University, UK)
 Axel Polleres (WU Wien, Austria)
 Sandro Preto (Federal University of ABC, Brazil)
 Badran Raddaoui (CNRS, France)
 Antonio Rago (Imperial College London, UK)
 Anna Rapberger (Imperial College London, UK)
 Maurício Reis (Universidade da Madeira, Portugal)
 Francesco Ricca (University of Calabria)
 Tjitze Rienstra (Maastricht University, The Netherlands)
 Ricardo O. Rodriguez (University of Buenos Aires, Argentina)
 Alessandro Ronca (University of Oxford, UK)
 Panos Rondogiannis (National and Kapodistrian University of Athens, Greece)
 Sasha Rubin (The University of Sydney, Australia)
 Chiaki Sakama (Wakayama University, Japan)
 Attila Sali (Alfréd Rényi Institute of Mathematics, Hungary)
 Zeynep G. Saribatur (TU Wien, Austria)
 Ken Satoh (Center for Juris-Informatics, Japan)
 Kai Sauerwald (University of Hagen, Germany)
 Domenico Fabio Savo (Università degli Studi di Bergamo, Italy)
 Federico M. Scafoglieri (Sapienza University of Rome, Italy)
 Konstantin Schekotihin (Alpen-Adria-Universität Klagenfurt, Austria)
 Steven Schockaert (Cardiff University, UK)
 Francois Schwarzentruher (École normale supérieure de Lyon, France)
 Dietmar A. Seipel (University of Würzburg, Germany)
 Baris Sertkaya (Frankfurt University of Applied Sciences, Germany)
 Paulo A. Shakarian (Arizona State University, USA)
 Gerardo I. Simari (Universidad Nacional del Sur in Bahía Blanca and CONICET, Argentina)
 Michael Sioutis (LIRMM, Université de Montpellier and CNRS, France)
 Kenneth Skiba (University of Hagen, Germany)
 Friedrich Slivovsky (University of Liverpool, UK)
 Marlo Souza (Federal University of Bahia, Brazil)

Mohan Sridharan (University of Edinburgh, UK)
 Roni Stern (Ben Gurion University, Israel)
 Frieder Stolzenburg (Harz University of Applied Sciences, Germany)
 Umberto Straccia (National Research Council, Italy)
 Hannes Strass (TU Dresden, Germany)
 Karim Tabia (University of Artois, France)
 David Tena Cucala (University of Oxford, UK)
 Andrea G. B. Tettamanzi (University of Côte d’Azur, France)
 Michael Thomazo (Inria; DI ENS; ENS; CNRS; University PSL, France)
 David Toman (University of Waterloo, Canada)
 Irina Trubitsyna (University of Calabria, Italy)
 Sara Ugolini (IIIA, CSIC, Spain)
 Markus Ulbricht (University of Leipzig, Germany)
 Mauro Vallati (University of Huddersfield, UK)
 Kees van Berkel (TU Wien, Austria)
 Leon van der Torre (University of Luxembourg, Luxembourg)
 Ivan Varzinczak (University of Paris 8, France)
 Srdjan Vesic (CNRS, France)
 Amanda Vidal (IIIA—CSIC, Spain)
 Jonni Virtema (University of Sheffield, UK)
 Domagoj Vrgoč (Pontificia Universidad Católica de Chile, Chile)
 Przemysław A. Wałęga (University of Oxford, UK)
 Kewen Wang (Griffith University, Australia)
 Yuheng Wang (Google)
 Christoph Wernhard (University of Potsdam, Germany)
 Emil Weydert (University of Luxembourg, Luxembourg)
 Marco Wilhelm (TU Dortmund, Germany)
 Xiang Yin (Imperial College London, UK)
 Marina Zanella (University of Brescia, Italy)
 Jessica Zangari (University of Calabria, Italy)
 Xiangfu Zhao (Yantai University, China)
 Shufang Zhu (University of Oxford, UK)
 Yoni Zohar (Bar Ilan University, Israel)

KR in the Wild Track Program Committee

Mario Alviano (University of Calabria, Italy)
 Chitta R. Baral (Arizona State University, USA)
 Gregor Behnke (University of Amsterdam, The Netherlands)
 Marc Denecker (KU Leuven, Belgium)
 Esra Erdem (Sabanci University, Turkiye)
 Johannes K. Fichte (Linköping University, Sweden)
 Martin Gebser (University of Klagenfurt, Austria)
 Joshua Hirschbrunn (University of Ulm, Germany)
 Ian Horrocks (University of Oxford, UK)
 Alexey Ignatiev (Monash University, Australia)
 Moritz Illich (University of Ulm, Germany)
 George Katsirelos (INRAE, France)
 Roman Kontchakov (Birkbeck, University of London)
 Markus Krötzsch (TU Dresden, Germany)
 Gerhard Lakemeyer (RWTH Aachen University, Germany)
 Daniel Le Berre (University of Artois and CNRS, France)
 Tuomo Lehtonen (Aalto University, Finland)
 Boris Motik (University of Oxford, UK)

Andreas Niskanen (University of Helsinki, Finland)
 Emilia Oikarinen (University of Helsinki, Finland)
 Bijan Parsia (University of Manchester, UK)
 Enrico Pontelli (New Mexico State University, USA)
 Nico Potyka (Cardiff University, UK)
 Guilin Qi (Southeast University, China)
 Francesco Ricca (University of Calabria, Italy)
 Jussi Rintanen (Aalto University, Finland)
 Baris Sertkaya (Frankfurt University of Applied Sciences, Germany)
 Heiner Stuckenschmidt (University of Mannheim, Germany)
 Matthias Thimm (University of Hagen, Germany)
 Francesca Toni (Imperial College London, UK)
 Mauro Vallati (University of Huddersfield, UK)
 Joost Vennekens (KU Leuven, Belgium)
 Guohui Xiao (University of Bergen, Norway)

Reasoning, Learning, and Decision Making Track Program Committee

Vaishak Belle (University of Edinburgh, UK)
 Pascal Bercher (The Australian National University, Australia)
 Blai Bonet (University Pompeu Fabra, Spain)
 Benjamin Bordais (RC Trust, Germany)
 Milan Češka (Brno University of Technology, Czechia)
 Dillon Z. Chen (LAAS-CNRS, France)
 Lukas Chrupa (Czech Technical University in Prague, Czechia)
 Alessandro Daniele (Fondazione Bruno Kessler, Italy)
 Hang Dong (University of Exeter, UK)
 Sébastien Destercke (Heudiasyc Laboratory, CNRS, France)
 Lei Fan (The University of New South Wales, Australia)
 Luis Galárraga (INRIA, France)
 Francesco Giannini (Scuola Normale Superiore, Italy)
 Julian Gutierrez (Monash University, Australia)
 Fredrik Heintz (Linköping University, Sweden)
 Daniel Höller (Saarland University, Germany)
 Yacine Izza (National University of Singapore, Singapore)
 Taylor T. Johnson (Vanderbilt University, USA)
 Sebastian Junges (Radboud University, The Netherlands)
 Bettina Könighofer (TU Graz, Austria)
 Nitesh Kumar (Cardiff University, UK)
 Ondřej Kuželka (Czech Technical University in Prague, Czechia)
 Matthias Lanzinger (TU Wien, Austria)
 Francesco Leofante (Imperial College London, UK)
 Songtuan Lin (The Australian National University, Australia)
 Robin Manhaeve (KU Leuven, Belgium)
 Joao Marques-Silva (ICREA, University of Lleida, Spain)
 Christian Meilicke (University of Mannheim, Germany)
 Daniel Neider (TU Dortmund University, Germany)
 David Penz (TU Wien, Austria)
 Ron Petrick (Heriot-Watt University, UK)
 Nico Potyka (Cardiff University, UK)
 David Rajaratnam (The University of New South Wales, Australia)

Savitha Sam Abraham (The University of Adelaide, Australia)
Francois Schwarzentruher (Ecole normale supérieure de Lyon, France, France)
Michael Sioutis (LIRMM, University of Montpellier and CNRS, France)
Yang Song (The University of New South Wales, Australia)
Mohan Sridharan (University of Edinburgh, UK)
Kostas Stathis (Royal Holloway, University of London, UK)
Charlie L. Street (University of Birmingham, UK)
Fabian Suchanek (Télécom Paris, France)
David Tena Cucala (University of Oxford, UK)
Anni-Yasmin Turhan (Paderborn University, Germany)
Mauro Vallati (University of Huddersfield, UK)
Xiaxia Wang (University of Oxford, UK)
Lianlong Wu (University of Cambridge, UK)
Pedro Zuidberg Dos Martires (Örebro University, Sweden)

Recently Published Research Track Program Committee

Loris Bozzato (Università degli Studi dell'Insubria, Italy)
Jens Claßen (Roskilde University, Denmark)
James Delgrande (Simon Fraser University, Canada)
Hector Geffner (RWTH Aachen University, Germany)
Jonas Philipp Haldimann (TU Wien, Austria and University of Cape Town, South Africa)
Martin Homola (Comenius University of Bratislava, Slovakia)
Timotheus Kampik (Umeå University, Sweden)
Matthias Knorr (Universidade Nova de Lisboa, Portugal)
Egor V. Kostylev (University of Oslo, Norway)
Francesco Kriegel (TU Dresden, Germany)
Davide Mario Longo (University of Calabria, Italy)
Carsten Lutz (Leipzig University, Germany)
Quentin Manière (Leipzig University, Germany)
Thomas Meyer (University of Cape Town and CAIR, South Africa)
Marie-Laure Mugnier (University of Montpellier)
Jandson S. Ribeiro (Cardiff University, UK)
Sebastian Sardina (RMIT University, Australia)
Leon van der Torre (University of Luxembourg, Luxembourg)
Johannes P. Wallner (TU Graz, Austria)
Xiaowang Zhang (Tianjin University, China)

Video Track Program Committee

Fabio G. Cozman (University of São Paulo, Brazil)
Annette Leßmöllmann (KIT Karlsruhe, Germany)
Carsten Lutz (University of Leipzig, Germany)
Nico Potyka (co-Chair) (Cardiff University, UK)
Anni-Yasmin Turhan (co-Chair) (Paderborn University, Germany)
Mary-Anne Williams (The University of New South Wales, Australia)

Additional Reviewers

Diego Aineto (University of Brescia, Italy)

Leila Amgoud (CNRS, IRIT, France)
Mark Amos (The University of New South Wales, Australia)
Aymeric Barbin (Sapienza University of Rome, Italy)
Timon Barlag (Leibniz University Hannover, Germany)
Davide Bizzaro, (Fondazione Bruno Kessler, Italy)
Zied Bouraoui (CRIL CNRS and University of Artois, France)
Camille Bourgaux (CNRS, ENS, Paris, France)
Sonja Bruckner (Software Competence Center Hagenberg, Austria)
Giovanni Buraglio, TU Wien, Austria
Angelos Charalambidis (Harokopio University, Greece)
Giannos Chatziagapis (National and Kapodistrian University of Athens, Greece)
Santiago Cifuentes (University of Buenos Aires)
Alessandro Cimatti (Fondazione Bruno Kessler, Italy)
Julien Corman (Free University of Bozen-Bolzano, Italy)
Mike Cruchten (University of Sheffield)
Tiziano Dalmonte, Free University of Bozen-Bolzano, Italy
Matthew Damigos (Ionian University, Greece)
Adam Dejl (Imperial College London, UK)
Jerome Delobelle (Université Paris Cité, France)
Vincent Derkinderen (KU Leuven, Belgium)
Devendra Dhami (TU Eindhoven, The Netherlands)
Daniele Theseider Dupré (Università del Piemonte Orientale, Italy)
Thorsten Engesser (IRIT, CNRS, France)
Katalin Fazekas (TU Wien, Austria)
Oskar Fiuk (University of Wrocław, Poland)
Maurice Funk (Leipzig University, Germany)
Luca Geatti (University of Udine, Italy)
Alessandro Gianola (INESC-ID/Instituto Superior Técnico and Universidade de Lisboa, Portugal)
Maksim Gladyshev (Utrecht University, The Netherlands)
Anton Gnatenko, Free University of Bozen-Bolzano, Italy
Larry Gonzalez (TU Dresden, Germany)
Umberto Grandi (IRIT, University Toulouse Capitole, France)
Sachin Grover (SRI, PARC, USA)
Alexander Hahn (TU Dortmund, Germany)
Andreas Herzig (IRIT-CNRS, France)
Tilman Hinnerichs (TU Delft, The Netherlands)
Till Hofmann (RWTH Aachen University, Germany)
Aidan Hogan (Universidad de Chile, Chile)
Xiaowei Huang (Liverpool University, UK)
Reijo Jaakkola (Tampere University, Finland)
Said Jabbour (CRIL CNRS UMR 8188, France)
Eleftherios Kalogeros (Ionian University, Greece)
Timotheus Kampik (Umeå University, Sweden)
Daumantas Kojelis (Manchester University, UK)
Jędrzej Kołodziejski (TU Dortmund, Germany)
Babis Kostopoulos (Harokopio University, Greece)
Francesco Kriegel (TU Dresden, Germany)
Isabelle Kuhlmann (University of Hagen, Germany)
Alex Ivliev (TU Dresden, Germany)
Magnus Liebl (University of Würzburg, Germany)
Xinghan Liu (TU Wien)
Vítor Lourenço (Universidade Federal Fluminense, Brazil and Amazon, UK)

Matthew Lynn (RWTH Aachen University, Germany)
Jan Maly (WU Wien, Austria)
Paulo Mann (Universidade do Estado do Rio de Janeiro,
Brazil)
Cristinel Mateis (Austrian Institute of Technology)
Arne Meier (Leibniz University Hannover, Germany)
Agnieszka Mensfelt (Royal Holloway, University of Lon-
don, UK)
Gian Carlo Milanese (University of Milano-Bicocca, Italy)
Paolo Morettin (University of Trento, Italy)
Matthew Morris (University of Oxford, UK)
Nima Motamed (Utrecht University, The Netherlands)
Nysret Musliu (TU Wien, Austria)
Hidetomo Nabeshima (University of Yamanashi, Japan)
Christoph Ohrem (University of Münster, Germany)
Nicola Paoletti (King's College London, UK)
David Parker (University of Oxford, UK)
Pekka Parviainen (University of Bergen, Norway)
Aleksandar Pavlović (TU Wien, Austria)
Gabriella Pigozzi (Université Paris-Dauphine)
Anna Rapberger (Imperial College London, UK)
Aldo Ricioppo (University of Cyprus, Cyprus)
Miguel Romero (Universidad Católica de Chile, Chile)
Alessandro Ronca (University of Oxford, UK)
Vladislav Ryzhikov (Birkbeck, University of London, UK)
Cesar Sanchez (IMDEA Software Institute, Spain)
Steven Schockaert (Cardiff University, UK)
Claudia Schon (Trier University of Applied Sciences, Ger-
many)
Nicolas Schwind (AIST, Japan)
Ian Shillito (ANU, Australia)
Takehide Soh (Kobe University, Japan)
Andre Thevapalan (TU Dortmund, Germany)
Michael Thielscher (The University of New South Wales,
Australia)
Matthias Thimm (University of Hagen, Germany)
Son Tran (Deakin University, Australia)
Timothy van Bremen (University of Huddersfield, UK)
Tarik Viehmann (RWTH Aachen University, Germany)
Marcel Waleska (University of Würzburg, Germany)
Daniel Weidner (University of Würzburg, Germany)
Marco Wilhelm (TU Dortmund, Germany)
Bo Xiong (University of Stuttgart, Germany)
Xiang Yin (Imperial College London, UK)
Michael Zakharyashev (Birkbeck, University of London,
UK)

4 PREFACE

Knowledge Representation and Reasoning (KR&R) is a cornerstone of Artificial Intelligence (AI) and a broad and prolific field of research with a vibrant international community. A fundamental premise in KR&R is that an intelligent agent's knowledge can be explicitly represented in a symbolic, declarative form, suitable for processing by dedicated reasoning engines. KR&R has contributed to the theory and practice of most areas of AI, including automated planning, computer vision, machine learning, multi-agent systems, natural language understanding, and robotics. Beyond AI, its impact can be seen in fields as diverse as computational biology, cybersecurity, databases, optimization, software engineering, and verification. The International Conference Series on Principles of Knowledge Representation and Reasoning, known as KR, is the leading forum for timely in-depth presentation of progress in the theory and principles underlying the representation and computational management of knowledge.

The twenty-first edition of KR was held from Saturday 2 November to Friday 8 November, 2024, in Hanoi, the ancient and lively capital city of Vietnam. The scientific program included four keynotes with an outstanding line-up of internationally renowned invited speakers, over a hundred technical talks, tutorials, workshops, software demonstrations, poster sessions, and a doctoral consortium. Social events allowed participants to experience the best of Vietnamese culture and cuisine. In the 2024 edition, we focused on heralding new approaches and developments and on stimulating reflection on the role of KR&R in the contemporary AI landscape, while maintaining the outstanding quality of the conference. We introduced new tracks that were very successful and we revised the conference subject/topic areas to make them more contemporary. The KR conference remains at the forefront of AI and has a very bright future.

Conference Tracks and Events. The conference featured various tracks and special events, summarized below.

- *Main Track:* As in the past, the Main Track focused on papers presenting novel results on the principles of KR&R, which clearly contribute to the formal foundations of the field, or show the applicability of KR&R techniques to implemented or implementable systems. It also welcomed papers from other areas that demonstrate clear use of, or contributions to, the principles or practice of KR&R.

The Main Track received 242 paper submissions and accepted 66 (27.3%). 62 of the accepted papers were full papers and 4 were short papers.

- *Demo Track:* This track was introduced in KR 2024 to showcase demonstrations of KR&R software. Instead of calling for direct submissions, the authors of selected submissions from the main and special tracks were invited to demonstrate their software based on suggestions from reviewers and Area Chairs.

The Demo Track exhibited 9 demonstrations at the conference on the afternoon of Thursday 7 November.

- *Recent Published Research Track:* The KR 2024 Recently Published Research Track provided a platform to discuss recent research on KR&R-related topics that had been published in relevant venues but not yet presented to the wider KR community. The track invited submissions of extended abstracts for papers published or accepted in journals or conference proceedings since 2022.

The Recently Published Research Track received 21 submissions, 17 of which were presented at the conference. Extended abstracts of accepted papers were shared on the conference web site.

- *Special Track – KR in the Wild:* This track was introduced in KR 2024 as a complement to the traditional KR Main Track that focuses on theoretical advances. In contrast, *KR in the Wild* aimed to showcase successful deployments of KR&R formalisms in all types of application domains as well as recent state-of-the-art developments in automated reasoning systems which form the basis for successful deployment of declarative problem solving in an ever-increasing number of practical settings.

Towards these goals, the KR in the Wild Track sought contributions from areas that are sometimes considered not core KR&R research. This included state-of-the-art reasoning systems and solvers developed in the various vibrant declarative programming communities as well as insightful applications of such systems. As further goals, the track aimed to foster interactions between practical and theoretical advances, and to encourage the discussion of new ideas, research experiences, emerging results and open challenges that can inspire novel research directions and influence the future of KR& research.

The KR in the Wild Special Track received 32 paper submissions with 8 papers (25%) selected for presentation at the conference.

- *Special Track – Reasoning, Learning and Decision Making:* This track was introduced in KR 2024 although similar tracks have appeared in recent editions of the conference. This special track emphasized the synergistic interactions between KR&R and complementary fields of AI, such as machine learning, robotics, multiagent systems, and decision-making in autonomous systems. The focus was on contributions that advance the state-of-the-art at the intersection of KR&R and these fields. Examples include integrating knowledge representation and reasoning into complete robotic systems, demonstrating how autonomous systems benefit from KR&R, employing common-sense reasoning and explicit world models for decision-making, and extending KR&R techniques to address challenges in decision-making across various environments or when interacting with the physical world. The track also sought submissions that combine aspects of KR&R and machine learning (ML) research, such as applying KR&R methods to solve ML challenges, using ML

methods to address KR&R problems, integrating learning and reasoning at the modelling or solving level, and leveraging combined KR&R and ML approaches for real-world applications.

The Reasoning, Learning and Decision Making Special Track received 98 paper submissions with 17 papers (17.4%) selected for presentation at the conference.

- *Video Track*: Another innovation in KR 2024 was the introduction of a video track. Its aim was twofold. First, it aimed to encourage KR groups to make their work more accessible to a wider audience and to promote best practice in KR&R outreach. Second, it provided a prominent platform for researchers to showcase interesting applications and insights and make them available to a wide audience from academia, industry and the general public. Eight videos were received and four of them were selected for screening at the conference. Each keynote presentation was preceded by one of these four videos. They are available on the KR Inc. website (<https://kr.org/>) and on the KR Conference Series YouTube channel.
- *Doctoral Consortium*: In the doctoral consortium, students had the opportunity to present their work during a poster session, which was preceded by a lightning talk in the conference plenary. Each participating student was also assigned a senior member of the KR community as a mentor and took part in a mentoring lunch. Ten students applied and were accepted to participate in the doctoral consortium.
- *Diversity and Inclusion (D&I) Event*: In a panel discussion, selected members of our community shared insights on diversity and inclusion in the KR community. This was followed by an interactive Q&A session for participants to engage with the panelists. After the panel, participants transitioned to a lunch activity, where interactive elements enhanced networking and discussion. Attendees were assigned random seating to encourage diverse interactions, along with conversation starters aligned with the panel themes to foster meaningful dialogue.

Submissions and Reviewing. We received 372 full and short paper submissions to the main technical program, including the KR Main Track, and the two special tracks *KR in the Wild* and *Reasoning, Learning and Decision Making*. In a thorough single-blind review process, submissions received between three and four reviews, with an average of 3.67 reviews per paper. In the Main Track and the Reasoning, Learning and Decision Making Track, reviewed papers also received a meta-review by an Area Chair. A total of 91 papers (24.5%) were accepted for presentation and inclusion in the conference proceedings. This excellent scientific program is the result of the hard work of 48 Area Chairs, 275 Program Committee members and 94 additional reviewers who provided a total of 1118 reviews and 277 meta-reviews. In addition, 16 PC members selected the accepted contributions to the Recently Published Research Track and the new Video Track. We are deeply grateful to all reviewers

and meta-reviewers for their hard work and dedication to the outstanding scientific program of KR 2024.

Submission Topics. The most popular submission topics of KR 2024 were Argumentation; Learning Logical Representations; Neuro-Symbolic AI; Description Logics; KR and Language Models; Belief Change; as well as Combinations of different methods (probabilistic, temporal, qualitative, ...). Other popular topics included Explanation, abduction and diagnosis; Reasoning about actions and change, action languages; Autonomous Decision Making; KR and Embeddings; Logic programming and answer set programming; and Non-monotonic logics.

Best Paper Awards. We are very glad to reveal the winners and runners-up for this year's best paper awards.

- *Ray Reiter Best Paper Award*: "Contractions Based on Optimal Repairs", by Franz Baader (TU Dresden and SCADS.AI Dresden/Leipzig, Germany) and Renata Wassermann (Universidade de São Paulo, Brazil).
- *Ray Reiter Best Paper Award Runner-Up*: "Shapley Value Computation in Ontology-Mediated Query Answering", by Meghyn Bienvenu, Diego Figueira and Pierre Lafourcade (LaBRI—CNRS and University of Bordeaux, France).
- *Marco Cadoli Best Student Paper Award*: "Qiana: A First-Order Formalism to Quantify over Contexts and Formulas", by Simon Coumes (Telecom Paris, Institut Polytechnique de Paris), Pierre-Henri Paris (Telecom Paris, Institut Polytechnique de Paris), François Schwarzentruher (Université de Rennes, IRISA, CNRS) and Fabian M. Suchanek (Telecom Paris, Institut Polytechnique de Paris).

We are very grateful to James P. Delgrande (Simon Fraser University, Canada), Gabriele Kern-Isberner (TU Dortmund, Germany), Renate Schmidt (University of Manchester, UK) and Frank Wolter (University of Liverpool, UK) for contributing their time, expertise, and effort to the selection of the KR 2024 best paper awards. There was a very competitive field of papers considered for these awards. Their insightful and considered deliberations have led to our selection of these worthy winners.

We also want to honour the following outstanding contributions to the Special Tracks:

- *Special Track: KR in the Wild Best Paper Award*: "Nemo: Your Friendly and Versatile Rule Reasoning Toolkit", by Alex Ivliev, Lukas Gerlach, Simon Meusel, Jakob Steinberg and Markus Krötzsch from TU Dresden.
- *Special Track: Reasoning, Learning and Decision Making Honorable Mention Award*: "Planning Domain Model Acquisition from State Traces without Action Parameters", by Tomas Balyo, Martin Suda, Lukas Chrupa, Dominik Safranek, Stephan Gocht, Filip Dvorak, Roman Bartak and G. Michael Youngblood from Filuta AI.

- *Special Track: Reasoning, Learning and Decision Making Honorable Mention Award*: “Logical Distillation of Graph Neural Networks”, by Alexander Pluska, Pascal Welke, Thomas Gärtner and Sagar Malhotra from TU Wien.

The *Best Video Award* went to Matti Berthold and Quentin Manière from Leipzig University for their video “A Gentle Introduction to the Stable Model Semantics for Logic Programs” by unanimous vote of the committee. The award came with a 500 USD prize from KR Inc. A *Peoples’ Choice Award* for the Video Track of KR 2024 was selected at the conference and announced after the publication of these proceedings.

KR Inc. Awards. The KR 2024 program also included a session to announce and honor the winners of three KR Inc. awards.

The Test of Time Award, initiated in 2020, is awarded for a paper published at the KR conference at least 15 years ago whose contributions have made a significant and lasting impact on the KR&R academic discipline, the application or adoption of KR&R techniques in industry, and/or for the betterment of society. This year the KR Test of Time Award goes to Adnan Darwiche for the KR 2002 paper “A Logical Approach to Factoring Belief Networks”. This paper developed an innovative approach to probabilistic inference based upon compiling Bayesian networks into tractable Boolean circuits. In addition to enabling exact inference for some models (ones with very large treewidth) that had been beyond the scope of existing approaches, this paper has had a lasting impact on the research community by popularizing the use of symbolic reasoning for exact probabilistic inference.

The Early Career Award, introduced in 2021, recognizes recent PhD students who have already made significant research contributions to the theory and/or practice of KR&R and who demonstrated great promise for future impactful contributions. This year the Early Career Award goes to Anna Rapberger for her work on computational argumentation. During her PhD studies (2018-2023) and subsequent postdoctoral research, Anna Rapberger has produced an impressive and innovative body of work centred on the theoretical and practical aspects of computational argumentation. Her research contributions were highlighted during the short plenary talk she delivered at KR 2024.

The Distinguished Service Award has previously only been awarded once (to Peter Patel-Schneider in 2016). This year the Awards Committee is delighted to revive this award to recognize Anthony (Tony) Cohn for his exceptional and sustained service to the Knowledge Representation and Reasoning community. Professor of Automated Reasoning in the School of Computing, University of Leeds, Tony Cohn is well-known for his ground breaking contributions to KR&R which led to numerous awards, including best paper awards and several fellowships (AAAI, EurAI, AISB, BCS, ...). Tony has shaped the KR community and its organization

KR Inc. for more than 25 years. He acted as PC co-chair of KR 1998 in Trento and as general chair of KR 2000 in Breckenridge. From 2000 to 2002 he was president of KR Inc. and continued as a member of KR Inc.’s board of directors until 2023. He was highly instrumental in turning KR Inc. into the organization it is today and the KR conference into the world’s leading event in its area.

We are very grateful to the members of this year’s KR Inc. Awards Committee: Meghyn Bienvenu (LaBRI—CNRS and University of Bordeaux, France), Gerhard Brewka (University of Leipzig, Germany), Tommie Meyer (University of Cape Town, South Africa)—supplemented by Diego Calvanese (Free University of Bozen-Bolzano, Italy) and Tony Cohn (University of Leeds, UK) for the Early Career Award. We thank them for contributing their time, expertise, and effort in selecting the winners.

Tutorials and Workshops. KR 2024 had an open call for tutorials and workshops. We had a strong tutorial and workshop program on a variety of interesting and timely topics covering the whole spectrum of KR&R. Seven tutorials were held at KR 2024:

- *An Introduction to Approximation Fixpoint Theory*, by Jesse Heyninck and Hannes Strass;
- *Large Language Models are Human-like Annotators*, by Manish Gupta, Mounika Marreddy, Subba Reddy Oota and Lucie Flek;
- *Argumentation and Machine Learning*, by Antonio Rago and Kristijonas Cyras;
- *Probing Machine Learning Models in Angluin’s Style*, by Ana Ozaki;
- *Fundamental Problems in Statistical Relational AI*, by Sagar Malhotra;
- *Formal Aspects of Strategic Reasoning in Multi-Agent Systems*, by Munyque Mittelmann and Aniello Murano; and,
- *Iterated Belief Change*, by Richard Booth and Jake Chandler.

The workshop program comprised the following events:

- *5th Workshop on Explainable Logic-Based Knowledge Representation (XLoKR 2024)*
Nico Potyka, Franz Baader, Bart Bogaerts, Joerg Hoffmann, Thomas Lukasiewicz and Francesca Toni;
- *1st International Workshop on Next-Generation Language Models for Knowledge Representation and Reasoning (NeLaMKRR 2024)*
Ha Thanh Nguyen, Randy Goebel, Francesca Toni, Ken Satoh and Kostas Stathis; and,
- *1st Workshop on Symbolic and Neuro-Symbolic Architectures for Intelligent Robotics Technology (SYNERGY 2024)*
Francesco Fabiano and Marcello Balduccini;

The tutorials and workshops were held over a period of three days preceding the main conference.

In addition, KR 2024 was strengthened and enhanced by the co-location of the *22nd International Workshop on Non-Monotonic Reasoning* (NMR 2024).

Keynote Presentations. The conference program included four keynote presentations from internationally recognized researchers aimed at cutting edge topics in KR&R:

- Meghyn Bienvenu (LaBRI—CNRS & University of Bordeaux, France): *KR Meets Data Quality*
- Subbarao Kambhampati (Arizona State University, USA): *Can LLMs Really Reason & Plan?*
- Nina Narodytska (VMware Research by Broadcom, USA): *Logic-based Explainability of ML Models.*
- Sheila McIlraith (University of Toronto, Canada) & Murray Shanahan (Imperial College London, UK): *Sheila McIlraith and Murray Shanahan in conversation with Joe Halpern.*

Conference Organization. We wish to thank those people who helped in the many aspects that were needed to make KR 2024 a hugely successful event. The Area Chairs, Track Chairs, Special Session Chairs, Program Committee members and additional reviewers were very helpful and supportive in their thorough and timely review of the submissions. The local arrangements teams, especially Professor Long Tran-Thanh and Ms. Hà Nguyễn Thu have gone above and beyond to ensure an outstanding conference. We cannot thank them enough.

The Workflow Chairs, Sanja Lukumbuzya (TU Wien, Austria) and Anthony Tompkins (UNSW, Australia) made sure the conference processes ran like a well oiled machine. In particular, they helped immensely with the move to CMT as the conference platform as well as ensuring that all aspects of the conference program were managed in a timely and efficient manner. Our KR 2024 Web Master Tobias Geibinger (TU Wien, Austria) developed an engaging web site and was always prompt in updating the site. He has set a high standard for others to follow. We were very fortunate in assembling this team and are very grateful for their outstanding contributions to the conference.

In closing, we thank our KR colleagues and the KR community for their significant contributions to KR 2024 which appear in the following pages. The amazing spirit and unwavering support of the KR community is what makes our conference special.

Vienna (Austria), Sydney (Australia), Lens (France)
October 2024

Magdalena Ortiz and Maurice Pagnucco
Program Chairs

Pierre Marquis
General Chair

5 INVITED TALKS

KR Meets Data Quality

Meghyn Bienvenu (LaBRI—CNRS and University of Bordeaux, France)

Real-world data notoriously suffers from various forms of imperfections (missing facts, erroneous facts, duplicates, etc.), which can limit its utility and lead to flawed analyses and unreliable decision making. This makes data quality an issue of paramount importance across application domains, and one which I’ll argue can both benefit from KR research and serve as a testbed for KR techniques. Indeed, while recent years have seen increasing interest in machine learning-based approaches, declarative approaches to improving data quality remain highly relevant, due to their better interpretability. In this talk, I will illustrate the synergy between data quality and KR by giving an overview of some of my recent work on querying inconsistent data using repair-based semantics and on rule-based approaches to entity resolution, highlighting the insights gained and directions for future research.

Meghyn Bienvenu is a senior researcher (directrice de recherche) at the CNRS (French National Center for Scientific Research), based at the LaBRI research lab in Bordeaux, France. Her research interests span a range of topics in knowledge representation and reasoning and database theory, but she is best known for her contributions to ontology-mediated query answering and to the study of logic-based methods for handling inconsistent data. Bienvenu’s research has been recognized by an invited Early Career Spotlight talk at IJCAI 16, the 2016 CNRS Bronze Medal in computer science, and together with her coauthors, a 2023 ACM PODS Alberto Mendelzon Test-of-Time Award. She has taken on numerous responsibilities within the AI, KR, and database theory communities, notably serving as PC co-chair for KR 2021 and associate editor of Artificial Intelligence Journal.

Can LLMs Really Reason & Plan?

Subbarao Kambhampati (Arizona State University)

Large Language Models (LLMs) are on track to reverse what seemed like an inexorable shift of AI from explicit to tacit knowledge tasks. Trained as they are on everything ever written on the web, LLMs exhibit “approximate omniscience”—they can provide answers to all sorts of queries, but with nary a guarantee. This could herald a new era for knowledge-based AI systems—with LLMs taking the role of (blowhard?) experts. But first, we have to stop confusing the impressive form of the generated knowledge for correct content, and resist the temptation to ascribe reasoning, planning, self-critiquing etc. powers to approximate retrieval by these n-gram models on steroids. We have to focus instead on LLM-Modulo techniques that complement the unfettered idea generation of LLMs with careful vetting by model-based AI systems. In this talk, I will reify this vision and attendant caveats in the context of the role of LLMs in planning tasks.

Subbarao Kambhampati is a professor of computer science at Arizona State University. Kambhampati studies fundamental problems in planning and decision making, motivated in particular by the challenges of human-aware AI systems. He is a fellow of Association for the Advancement of Artificial Intelligence, American Association for the Advancement of Science, and Association for Computing Machinery, and was an NSF Young Investigator. He served as the president of the Association for the Advancement of Artificial Intelligence, a trustee of the International Joint Conference on Artificial Intelligence, the chair of AAAS Section T (Information, Communication and Computation), and a founding board member of Partnership on AI. Kambhampati’s research as well as his views on the progress and societal impacts of AI have been featured in multiple national and international media outlets. He can be followed on Twitter @rao2z.

Logic-based Explainability of ML Models

Nina Narodytska (VMware Research by Broadcom, USA)

Machine learning models are among the most successful artificial intelligence technologies making an impact in a variety of practical applications. However, many concerns were raised about the ‘magical’ power of these models. It is disturbing that we are clearly lacking in understanding of the decision making process behind this technology. Therefore, a natural question is whether we can trust decisions that neural networks make. There are a large number of research approaches to explainability. One popular family of methods, so-called ad-hoc explainability methods, uses heuristic-based solutions. These methods are among the most practical due to their scalability but they do not provide any guarantees on the quality of explanations. To address this issue, we propose a formal approach where explainability is formalized as a logical problem and solved using optimization tools, like SMT, SAT, ILP. Using these techniques we are able to compute provably correct explanations for smaller ML models. We consider several techniques to scale logic-based methods to larger ML models. We will highlight an interesting connection between explainability and robustness of ML models.

Nina Narodytska is a staff researcher at VMware Research by Broadcom. Prior to VMware, she was a researcher at Samsung Research America. She completed postdoctoral studies in the Carnegie Mellon University School of Computer Science and the University of Toronto. She received her PhD from the University of New South Wales. She was named one of “AI’s 10 to Watch” researchers in the field of AI in 2013. She has presented invited talks and tutorials at FMCAD 18, CP 19, AAAI 20, IJCAI 20, LMML 22, CP 22 and ESSAI 23.

Great Ideas from KR: drawing on the past to shape the future

Sheila McIlraith (University of Toronto, Canada) and Murray Shanahan (Imperial College London, UK) in conversation with Joe Halpern (Cornell University, USA)

Sheila McIlraith and Murray Shanahan, in conversation with Joe Halpern, will reflect on KR’s noble intellectual legacy, drawing on some of their personal favourite ideas from the KR canon. They will also speculate on the role of KR in this era of large language models, including how KR will shape the future of AI.

Sheila McIlraith is a Professor in the Department of Computer Science at the University of Toronto, a Canada CIFAR AI Chair (Vector Institute), and an Associate Director and Research Lead at the Schwartz Reisman Institute for Technology and Society. McIlraith’s research is in the area of AI sequential decision making broadly construed, with a focus on human-compatible AI. McIlraith is a Fellow of the ACM and the Association for the Advancement of Artificial Intelligence (AAAI). With co-authors McIlraith has been honoured with the 2011 SWSA 10-year Award, recognizing the highest impact paper from ISWC 10 years prior; in 2022, the ICAPS 22 Influential Paper Award, a 10-year test of time award; and most recently the 2023 IJCAI-JAIR Paper Prize, awarded annually to an outstanding paper published in JAIR in the preceding five years.

Murray Shanahan is a principal research scientist at Google DeepMind and Professor of Cognitive Robotics at Imperial College London. His publications span artificial intelligence, robotics, machine learning, logic, dynamical systems, computational neuroscience, and philosophy of mind. He is active in public engagement, and was scientific advisor on the film *Ex Machina*. His books include “Embodiment and the Inner Life” (2010) and “The Technological Singularity” (2015).

Joseph (Joe) Halpern is the Joseph C. Ford Professor of the Computer Science Department at Cornell University. His major research interests are in reasoning about knowledge and uncertainty, security, distributed computation, decision theory, and game theory. He has coauthored 6 patents, three books (“Reasoning About Knowledge”, “Reasoning about Uncertainty”, and “Actual Causality”), and over 400 technical publications. Halpern is a Fellow of AAAI, AAAS (American Association for the Advancement of Science), the American Academy of Arts and Sciences, ACM, IEEE, the Game Theory Society, the National Academy of Engineering, and SAET (Society for the Advancement of Economic Theory). Among other awards, he received the Kampe de Fariet Award in 2016, the ACM SIGART Autonomous Agents Research Award in 2011, the Dijkstra Prize in 2009, the ACM/AAAI Newell Award in 2008, the Godel Prize in 1997, was a Guggenheim Fellow in 2001-02, and a Fulbright Fellow in 2001-02 and 2009-10.

6 SPONSORS

The KR 2024 conference was supported by the following organisations.

Platinum Sponsors

Artificial Intelligence Journal
KR Inc. (Principles of Knowledge Representation and Reasoning)

Gold Sponsors

Faculty of Engineering, The University of New South Wales,
Sydney, Australia

Silver Sponsors

Potassco Solutions
VinIF Institute, Vietnam