Proceedings of the 17th International Workshop on Neural-Symbolic Learning and Reasoning (NeSy)

Artur S. d'Avila Garcez[1], Tarek R. Besold[2], Marco Gori[3] and Ernesto Jiménez-Ruiz[1,4]

[1]City, University of London, UK
[2]Sony AI, Barcelona, Spain
[3]University of Siena, Italy
[4]SIRIUS, University of Oslo, Norway

Preface

NeSy is the annual meeting of the Neural-Symbolic Learning and Reasoning Association[1] and the premier venue for the presentation and discussion of the theory and practice of neural-symbolic computing systems.[2] Since 2005, NeSy has provided an atmosphere for the free exchange of ideas bringing together the community of scientists and practitioners that straddle the line between deep learning and symbolic AI.

Neural networks and statistical Machine Learning have obtained industrial relevance in a number of areas from retail to healthcare, achieving state-of-the-art performance at language modelling, speech recognition, graph analytics, image, video and sensor data analysis. Symbolic AI, on the other hand, is challenged by such unstructured data, but is recognised as being in principle transparent, in that reasoned facts from knowledge-bases can be inspected to interpret how decisions follow from input. Neural and symbolic methods also contrast in the problems that they excel at: scene recognition from images appears to be a problem still outside the capabilities of symbolic systems, for example, while neural networks are not yet sufficient for industrial-strength complex planning scenarios and deductive reasoning tasks.

Neurosymbolic AI aims to build rich computational models and systems by combining neural and symbolic learning and reasoning paradigms. This combination hopes to form synergies among their strengths while overcoming their complementary weaknesses.

NeSy 2023 was held in La Certosa di Pontignano, Siena, Italy, 3-5 July 2023.[3] NeSy welcomed

17th International Workshop on Neural-Symbolic Learning and Reasoning (NeSy 2023)
a.garcez@city.ac.uk (A. S. d. Garcez); tarek.besold@googlemail.com (T. R. Besold); marco.gori@unisi.it (M. Gori); ernesto.jimenez-ruiz@city.ac.uk (E. Jiménez-Ruiz)
https://www.city.ac.uk/about/people/academics/artur-davila-garcez (A. S. d. Garcez); https://sites.google.com/site/tarekbesold/ (T. R. Besold); https://www3.diism.unisi.it/~marco/ (M. Gori); https://www.city.ac.uk/about/people/academics/ernesto-jimenez-ruiz (E. Jiménez-Ruiz)
0000-0001-7375-9518 (A. S. d. Garcez); 0000-0002-8002-0049 (T. R. Besold); 0000-0001-6337-5430 (M. Gori); 0000-0002-9083-4599 (E. Jiménez-Ruiz)

[1]https://www.city-data-science-institute.com/nesy
[2]http://www.neural-symbolic.org/
[3]https://sites.google.com/view/nesy2023

submissions of the latest and ongoing research work on neurosymbolic AI for presentation at the workshop. Topics of interest included, but were not limited to:

- Knowledge representation and reasoning using deep neural networks;
- Symbolic knowledge extraction from neural and statistical learning systems;
- Explainable AI methods, systems and techniques integrating connectionist and symbolic AI;
- Enhancing deep learning systems through structured background knowledge;
- Neural-symbolic cognitive agents;
- Biologically-inspired neuro-symbolic integration;
- Integration of logics and probabilities in neural networks;
- Neural-symbolic methods for structure learning, transfer learning, meta, multi-task and continual learning, relational learning;
- Novel connectionist systems able to perform traditionally symbolic AI tasks (e.g. abduction, deduction, out-of-distribution learning);
- Novel symbolic systems able to perform traditionally connectionist tasks (e.g. learning from unstructured data, distributed learning);
- Embedding methods for structured information, such as knowledge graphs, mathematical expressions, grammars, knowledge bases, logical theories, etc.;
- Applications of neural-symbolic and hybrid systems, including in simulation, finance, healthcare, robotics, Semantic Web, software engineering, systems engineering, bioinformatics and visual intelligence.

NeSy received 41 regular submissions for peer-review; out of these, 33 papers were accepted for presentation in the workshop and inclusion within these proceedings. In addition, NeSy also received 26 extended abstracts summarising recently published papers at top conferences and journals; out of these, 15 extended abstracts were accepted.

Keynote and Invited Talks

NeSy 2023 featured the following keynotes and invited talks:

Keynote talk
Leslie Valiant, University of Harvard, USA
University of Siena Laurea ad Honorem Lecture
Yann LeCun, Meta AI and New York University, USA
Invited talks
Fosca Giannotti, Scuola Normale Superiore, Pisa, Italy
Alvaro Velasquez, DARPA Programme Manager, Assured Neuro-Symbolic Learning and Reasoning, USA
Murray Campbell, IBM Research, USA

Leslie Valiant

Title: Augmenting Learning with Reasoning

Abstract: The question we ask is whether one can build on the success of machine learning to address the broader goals of artificial intelligence. We regard reasoning as the major component of cognition that needs to be added. We suggest that the central challenge therefore is to unify our understanding of these two phenomena, learning and reasoning, into a single framework with a common semantics. In such a framework one would aim to learn rules with the same success that predicates can be learned by means of machine learning, and, at the same time, to reason with the rules with guarantees analogous to those of standard logic. We discuss how Robust Logic fulfils the role of such a theoretical framework. We also discuss the challenges of realizing this on a significant scale for tasks where the performance offered exceeds that achievable by learning alone.

Yann LeCun

Title: Towards Machines that can Learn, Reason, and Plan.

Abstract: How could machines learn as efficiently as humans and animals? How could machines learn how the world works and acquire common sense? How could machines learn to reason and plan? Current AI architectures, such as Auto-Regressive Large Language Models fall short. I will propose a modular cognitive architecture that may constitute a path towards answering these questions. The centerpiece of the architecture is a predictive world model that allows the system to predict the consequences of its actions and to plan a sequence of actions that optimize a set of objectives. The world model employs a Hierarchical Joint Embedding Predictive Architecture (H-JEPA) trained with self-supervised learning. The JEPA learns abstract representations of the percepts that are simultaneously maximally informative and maximally predictable. The corresponding working paper is available here: https://openreview.net/forum?id=BZ5a1r-kVsf

Murray Campbell

Title: Towards a Neuro-Symbolic Benchmark

Abstract: Benchmarks are the primary tool we use for assessing our progress in AI. However, benchmarking as currently practiced is problematic in several ways. This talk will review these problems and discuss some of the principles and best practices for overcoming these issues. The talk will then consider how these principles could be applied to the development of a neuro-symbolic benchmark, i.e., a benchmark where we expect neuro-symbolic approaches to have an advantage of purely neural or symbolic methods.

Fosca Giannotti

Title eXplainable AI (XAI): paradigms in support of synergistic human-machine interaction and collaboration

Abstract: Black box AI systems for automated decision-making, often based on machine learning over (big) data, map a user's features into a class or a score without exposing the

reasons why. This is problematic not only for the lack of transparency but also for possible biases inherited by the algorithms from human prejudices and collection artifacts hidden in the training data, which may lead to unfair or wrong decisions. The future of AI lies in enabling people to collaborate with machines to solve complex problems. Like any efficient collaboration, this requires good communication, trust, clarity, and understanding. Explainable AI addresses such challenges and for years different AI communities have studied such topics, leading to different definitions, evaluation protocols, motivations, and results. This lecture provides a reasoned introduction to the work of Explainable AI (XAI) to date and surveys the literature. A special focus will be on paradigms in support of synergistic human-machine interaction and collaboration to improve joint performance in high-stake decision-making as for example methods aimed at engaging users with factual and counterfactual or other high-level explanations encoding domain knowledge and user background, methods focusing on conversational explainable AI, methods for the understanding the impact of explanation on expert users' information-seeking strategies, mental model updating, and trust calibration and the steps needed for new paradigms that can promote collaboration and seamless interaction maintaining the human responsibility of the choice through a progressive disclosure to prevent cognitive overload.

Organisation

General Chair

Tarek Besold	Sony AI, Barcelona, Spain

Programme Chairs

Artur d'Avila Garcez	City, University of London
Ernesto Jiménez Ruiz	City, University of London
Marco Gori	University of Siena, Italy

Local Organisation

Michelangelo Diligenti	University of Siena, Italy
Francesco Giannini	University of Siena, Italy
Marco Maggini	University of Siena, Italy
Stefano Melacci	University of Siena, Italy

Challenge Chair

Pranava Madhyastha	City, University of London, UK

Applied Session Chairs

Dragos Margineantu	Boeing, Seattle, USA
Alberto Speranzon	Lockheed Martin, USA

Program Committee

Hanna Abi Akl	Data ScienceTech Institute, Paris, France
Mehwish Alam	Télécom Paris, Institute Polytechnique de Paris, France
Vito Walter Anelli	Politecnico di Bari, Italy
Vaishak Belle	The University of Edinburgh, UK
Davide Beretta	Università degli Studi di Parma, Italy
Matthew Brown	University of California Los Angeles, US
Jiaoyan Chen	University of Manchester, UK
Claudia d'Amato	University of Bari, Italy
Wang-Zhou Dai	Nanjing University, China
Alessandro Daniele	Fondazione Bruno Kessler, Italy
Devendra Dhami	Technische Universität Darmstadt, Germany
Elvira Domínguez	Universidad Politécnica de Madrid, Spain
Ivan Donadello	Free University of Bozen-Bolzano, Italy
Vasilis Efthymiou	ICS-FORTH, Greece
Vijay Ganesh	University of Waterloo, Canada
Leilani Gilpin	University of California, Santa Cruz, US
Eleonora Giunchiglia	TU Wien, Austria
Oktie Hassanzadeh	IBM Research, US
Janna Hastings	University of Zurich, Switzerland
Nelson Higuera	TU Wien, Austria
Robert Hoehndorf	King Abdullah University of Science and Technology, Saudi Arabia
Andreas Holzinger	Medical University and Graz University of Technology, Austria
Xiaowei Huang	University of Liverpool, UK
Filip Ilievski	University of Southern California, US
Azanzi Jiomekong	University of Yaounde, Cameroon
Moa Johansson	Chalmers University of Technology, Sweden
Kristian Kersting	TU Darmstadt, Germany
Muhammad Jaleed Khan	University of Galway, Ireland
Egor V. Kostylev	University of Oslo, Norway
Kai-Uwe Kuehnberger	University of Osnabrck, Germany
Sofoklis Kyriakopoulos	City, University of London, UK

Luis Lamb	Federal University of Rio Grande do Sul, Brazil
Francesca Alessandra Lisi	Università degli Studi di Bari, Italy
Thomas Lukasiewicz	TU Wien, Austria
Pranava Madhyastha	City, University of London, UK
Robin Manhaeve	Katholieke Universiteit Leuven, Belgium
Lia Morra	Politecnico di Torino, Italy
Raghava Mutharaju	IIIT-Delhi, India
Francesca Rossi	IBM Research, US
Alessandra Russo	Imperial College London, UK
Alessandro Oltramari	Bosch Research and Technology Center, US
Catia Pesquita	Universidade de Lisboa, Portugal
Alina Petrova	Thomson Reuters Labs, UK
Md Kamruzzaman Sarker	University of Hartford, CT, US
Gaia Saveri	University of Pisa, Italy
Ute Schmid	University of Bamberg, Germany
Howard Schneider	Sheppard Clinic North, Canada
Luciano Serafini	Fondazione Bruno Kessler, Italy
Daniel L. Silver	Acadia University, Canada
Alessandro Sperduti	University of Padua, Italy
Mihaela Stoian	University of Oxford, UK
David Tena Cucala	University of Oxford, UK
Alberto Testolin	University of Padova, Italy
Son Tran	University of Tasmania, Australia
Efthymia Tsamoura	Samsung AI Research, UK
Frank Van Harmelen	Vrije Universiteit Amsterdam, Netherlands
Antonio Vergari	University of Edinburgh, UK
Tillman Weyde	City, University of London, UK
Gerson Zaverucha	Federal University of Rio de Janeiro, Brazil

Additional Reviewers

Piyush Jha	University of Waterloo, Canada
Fernando Zhapa-Camacho	King Abdullah University of Science and Technology, Saudi Arabia
John Lu	University of Waterloo, Canada
Francesco Manigrasso	Politecnico di Torino, Italy

Acknowledgements

We thank all members of the program committee, additional reviewers, keynote speakers, authors and local organizers for their efforts. We would also like to acknowledge that the work of the workshop organisers was greatly simplified by using the EasyChair conference management system and the CEUR open-access publication service.