

# Pietro Barbiero

[barbip@usi.ch](mailto:barbip@usi.ch)

AI researcher with 6+ years of experience in interpretable AI, neural-symbolic computing, and multi-agent interactions with evolutionary algorithms. **Working on:** interpretability, graph neural networks, and neural-symbolic AI.

## WORK EXPERIENCE

---

**Research assistant**, Università della Svizzera Italiana, Switzerland. 2023 – *present*

**Explainable AI researcher**, Tenyks, UK. Nov. 2022 – Dec. 2022

- Invented novel explainable data augmentation improving model performance by 15x on one class.
- Implemented state-of-the-art automatic model debugging.

**Machine learning scientist**, Amazon.com Inc., UK. Aug. 2022 – Oct. 2022

- Implemented and improved state-of-the-art explainable AI techniques in production environment.
- Realized explainable AI systems satisfying all the constraints for safe and ethical deployment.
- Managed and analyzed sensitive data.

**President**, Bactell Inc., US. Apr. 2021 – Oct. 2021

- Collaborated across a global diverse team (US, Italy, UK, and France).
- Engaged stakeholders and investors.
- Managed and delivered time-sensitive projects.

**Research assistant**, Artificial Intelligence Group, University of Cambridge, UK. 2020 – 2023

- Invented and implemented interpretable deep learning models.
- Published results at top international journals and conferences (e.g., NeurIPS, ICLR, AAAI, GECCO, IEEE).
- Supervised bachelor and master theses resulting in publications at top conferences (e.g., NeurIPS and AAAI).
- Collaborated in teaching master course on “Representation Learning on Graphs and Networks”.

**AI engineer & founding team member**, Bactell Inc., US. 2019 – 2021

- Invented and developed rapid, sensitive, and predictive methods to improve antimicrobial susceptibility testing.
- Participated in securing a \$40k seed round.
- Achieved results comparable with competitors and the current state of the art on two blind tests.

## EDUCATION

---

**Ph.D. in Computer Science**, Clare College, University of Cambridge, UK. 2020 – 2023

**M.Eng. in Mathematics**, Politecnico di Torino, Italy. 2017 – 2019

**B.Eng. in Computer Science**, Politecnico di Torino, Italy. 2013 – 2017

## RESEARCH

---

**40 papers, h-index 10, 406 citations.**

**Interpretable deep reasoning** [4, 5, 8, 10] [pytorch-explain@github](https://github.com/pytorch-explain)

- Invented and developed logic explained networks and high-dimensional concept representations.
- Solved the information bottleneck of concept-based models breaking the accuracy-vs-explainability trade-off.
- Invented interpretable unsupervised concept layer for graph neural networks.

**Biomedical digital twins** [14, 22] [digital-patient@github](https://github.com/digital-patient)

- Designed a modular framework integrating and extending pathophysiological models.
- Implemented a virtual patient combining organ, tissue, and cellular information with graph neural networks.
- Paper highlighted by the Deep Graph Library maintainers (<https://github.com/dmlc/dgl>).

## SOFTWARE AND LANGUAGE SKILLS

---

Fluent: Python, PyTorch, LaTeX.

Proficient: Ninja, MatLab, R, MySQL, MongoDB, Java, C, C++, HTML/CSS, Bash, Assembly 8086.

Bioinformatics: GATK, SPAdes, BLAST.

## AWARDS

---

Top-3 Paper Award — NeurIPS Workshop on Human in the Loop.	2022
Cambridge Academy of Therapeutic Sciences Sponsorship — Impulse Programme for Tech Innovators.	2021
EIA 2019 Top Team Award — European Innovation Academy 2019.	2019
DECON 2019 Travel Award — International Conference on Decision Economics.	2019
WIRN 2018 Travel Award — 28th Italian Workshop on Neural Networks.	2018
WIRN 2017 Travel Award — 27th Italian Workshop on Neural Networks.	2017

## INVITED TALKS AND SEMINARS

---

Invited at the Spring workshop on Mining and Learning, Sint-Michielsgestel, Netherlands. Title: <i>Interpretable Neural Symbolic Concept Reasoning</i>	2023
Massachusetts Institute of Technology, Cambridge, Massachusetts, US. Title: <i>Concept Embedding Models: Beyond the Accuracy Explainability Trade-Off</i>	2022
Université Côte d’Azur, Nice, France. Title: <i>Concept Embedding Models: Beyond the Accuracy Explainability Trade-Off</i>	2022
Workshop on “Language Models, Graph Embeddings and Neuro-Symbolic Learning”, Leuven, Belgium. Title: <i>Concept Embedding Models: Beyond the Accuracy Explainability Trade-Off</i>	2022
1st Nice Workshop on Interpretability, Nice, France. Title: <i>Concept Embedding Models: Beyond the Accuracy Explainability Trade-Off</i>	2022
Università degli Studi di Siena, Siena, Italy. Title: <i>Evolutionary coresets and machine learning epistemology</i>	2019

## LANGUAGES

---

English: C1 Italian: native speaker.

## VOLUNTEER ACTIVITIES

---

Mentored next generations of STEM rising stars at <a href="#">LeadTheFuture</a> .	2022 – present
Reviewed papers for top conferences and journals (e.g., IEEE, NeurIPS, AAI, ICLR).	2020 – present

## REFERENCES

---

Pietro Lió ([pl219@cam.ac.uk](mailto:pl219@cam.ac.uk))—Full Professor, University of Cambridge, UK.  
Mateja Jamnik ([mj201@cam.ac.uk](mailto:mj201@cam.ac.uk))—Full Professor, University of Cambridge, UK.  
Giuseppe Marra ([giuseppe.marra@kuleuven.be](mailto:giuseppe.marra@kuleuven.be))— Assistant Professor, Katholieke Universiteit Leuven, Belgium.

## Publications

- [1] Dmitry Kazhdan, Botty Dimanov, Lucie Charlotte Magister, Pietro Barbiero, Mateja Jamnik, and Pietro Lio. ‘GCI: A (G) raph (C) oncept (I) nterpretation Framework’. In: *arXiv preprint arXiv:2302.04899* (2023).
- [2] Mateo Espinosa Zarlenga, Pietro Barbiero, Zohreh Shams, Dmitry Kazhdan, Umang Bhatt, Adrian Weller, and Mateja Jamnik. ‘Towards Robust Metrics for Concept Representation Evaluation’. In: *arXiv preprint arXiv:2301.10367* (2023).
- [3] Steve Azzolin, Antonio Longa, Pietro Barbiero, Pietro Liò, and Andrea Passerini. ‘Global explainability of gnns via logic combination of learned concepts’. In: *arXiv preprint arXiv:2210.07147* (2022).
- [4] Pietro Barbiero, Gabriele Ciravegna, Francesco Giannini, Pietro Lió, Marco Gori, and Stefano Melacci. ‘Entropy-based logic explanations of neural networks’. In: *Proceedings of the AAI Conference on Artificial Intelligence*. Vol. 36. 6. 2022, pp. 6046–6054.
- [5] Dobrik Georgiev, Pietro Barbiero, Dmitry Kazhdan, Petar Veličković, and Pietro Liò. ‘Algorithmic Concept-Based Explainable Reasoning’. In: *Proceedings of the AAI Conference on Artificial Intelligence*. Vol. 36. 6. 2022, pp. 6685–6693.
- [6] Rishabh Jain, Gabriele Ciravegna, Pietro Barbiero, Francesco Giannini, Davide Buffelli, and Pietro Lio. ‘Extending Logic Explained Networks to Text Classification’. In: *arXiv preprint arXiv:2211.09732* (2022).

- [7] Marta Lovino, Vincenzo Randazzo, Gabriele Ciravegna, Pietro Barbiero, Elisa Ficarra, and Giansalvo Cirrincione. ‘A survey on data integration for multi-omics sample clustering’. In: *Neurocomputing* 488 (2022), pp. 494–508.
- [8] Lucie Charlotte Magister, Pietro Barbiero, Dmitry Kazhdan, Federico Siciliano, Gabriele Ciravegna, Fabrizio Silvestri, Pietro Liò, and Mateja Jamnik. ‘Encoding Concepts in Graph Neural Networks’. In: *arXiv e-prints* (2022), arXiv–2207.
- [9] Han Xuanyuan, Pietro Barbiero, Dobrik Georgiev, Lucie Charlotte Magister, and Pietro Lió. ‘Global Concept-Based Interpretability for Graph Neural Networks via Neuron Analysis’. In: *arXiv preprint arXiv:2208.10609* (2022).
- [10] Mateo Espinosa Zarlenga, Pietro Barbiero, Gabriele Ciravegna, Giuseppe Marra, Francesco Giannini, Michelangelo Diligenti, Zohreh Shams, Frederic Precioso, Stefano Melacci, Adrian Weller, Pietro Liò, and Mateja Jamnik. ‘Concept Embedding Models’. In: *Advances in neural information processing systems* (2022). [In print].
- [11] Mateo Espinosa Zarlenga, Pietro Barbiero, Zohreh Shams, Dmitry Kazhdan, Umang Bhatt, Adrian Weller, and Mateja Jamnik. ‘On The Quality Assurance Of Concept-Based Representations’. In: *Advances in neural information processing systems* (2022). [Under review]. URL: <https://openreview.net/pdf?id=Ehkh6jyas6v>.
- [12] Pietro Barbiero, Gabriele Ciravegna, Dobrik Georgiev, and Francesco Giannini. ‘Pytorch, Explain! A Python library for Logic Explained Networks’. In: *arXiv preprint arXiv:2105.11697* (2021).
- [13] Pietro Barbiero, Giovanni Squillero, and Alberto Tonda. ‘Predictable Features Elimination: An Unsupervised Approach to Feature Selection’. In: *International Conference on Machine Learning, Optimization, and Data Science*. Springer, 2021, pp. 399–412.
- [14] Pietro Barbiero, Ramon Viñas Torné, and Pietro Lió. ‘Graph Representation Forecasting of Patient’s Medical Conditions: Toward a Digital Twin’. In: *Frontiers in genetics* 12 (2021).
- [15] Gabriele Ciravegna, Pietro Barbiero, Giansalvo Cirrincione, Giovanni Squillero, and Alberto Tonda. ‘Discovering hierarchical neural archetype sets’. In: *Progresses in Artificial Intelligence and Neural Systems*. Springer, 2021, pp. 255–267.
- [16] Gabriele Ciravegna, Pietro Barbiero, Francesco Giannini, Marco Gori, Pietro Lió, Marco Maggini, and Stefano Melacci. ‘Logic explained networks’. In: *arXiv preprint arXiv:2108.05149* (2021).
- [17] Francesca Nonis, Pietro Barbiero, Giansalvo Cirrincione, Elena Carlotta Olivetti, Federica Marcolin, and Enrico Vezzetti. ‘Understanding abstraction in deep CNN: an application on facial emotion recognition’. In: *Progresses in Artificial Intelligence and Neural Systems*. Springer, 2021, pp. 281–290.
- [18] Pietro Barbiero, Andrea Bertotti, Gabriele Ciravegna, Giansalvo Cirrincione, and Elio Piccolo. ‘DNA Microarray Classification: Evolutionary Optimization of Neural Network Hyper-parameters’. In: *Neural Approaches to Dynamics of Signal Exchanges*. Springer, 2020, pp. 305–311. ISBN: 978-981-13-8949-8. DOI: [10.1007/978-981-13-8950-4\\_28](https://doi.org/10.1007/978-981-13-8950-4_28).
- [19] Pietro Barbiero, Andrea Bertotti, Gabriele Ciravegna, Giansalvo Cirrincione, Elio Piccolo, and Alberto Tonda. ‘Understanding Cancer Phenomenon at Gene Expression Level by using a Shallow Neural Network Chain’. In: *Neural Approaches to Dynamics of Signal Exchanges*. Springer, 2020, pp. 281–290. ISBN: 978-981-13-8949-8. DOI: [10.1007/978-981-13-8950-4\\_26](https://doi.org/10.1007/978-981-13-8950-4_26).
- [20] Pietro Barbiero, Gabriele Ciravegna, Alberto Tonda, and Giovanni Squillero. ‘Generating Neural Archetypes to Instruct Fast and Interpretable Decisions’. In: *Decision Economics: Complexity of Decisions and Decisions for Complexity*. Springer, 2020.
- [21] Pietro Barbiero, Giansalvo Cirrincione, Maurizio Cirrincione, Elio Piccolo, and Francesco Vaccarino. ‘Neural Epistemology in Dynamical System Learning’. In: *Neural Approaches to Dynamics of Signal Exchanges*. Springer, 2020, pp. 213–221. ISBN: 978-981-13-8950-4. DOI: [10.1007/978-981-13-8950-4\\_20](https://doi.org/10.1007/978-981-13-8950-4_20).
- [22] Pietro Barbiero and Pietro Lió. ‘The Computational Patient has Diabetes and a COVID’. In: *arXiv preprint arXiv:2006.06435* (2020).
- [23] Pietro Barbiero, Marta Lovino, Mattia Siviero, Gabriele Ciravegna, Vincenzo Randazzo, Elisa Ficarra, and Giansalvo Cirrincione. ‘Unsupervised multi-omic data fusion: The neural graph learning network’. In: *International Conference on Intelligent Computing*. Springer, 2020, pp. 172–182.
- [24] Pietro Barbiero, Giovanni Squillero, and Alberto Tonda. ‘Modeling generalization in machine learning: A methodological and computational study’. In: *arXiv preprint arXiv:2006.15680* (2020).
- [25] Pietro Barbiero, Giovanni Squillero, and Alberto Tonda. ‘Uncovering Coresets for Classification With Multi-Objective Evolutionary Algorithms’. In: *arXiv preprint arXiv:2002.08645* (2020).
- [26] Pietro Barbiero and Alberto Tonda. ‘Making Sense of Economics Datasets with Evolutionary Coresets’. In: *Decision Economics: Complexity of Decisions and Decisions for Complexity*. Springer, 2020.
- [27] Gabriele Ciravegna, Giansalvo Cirrincione, Federica Marcolin, Pietro Barbiero, Nicole Dagnes, and Elio Piccolo. ‘Assessing discriminating capability of geometrical descriptors for 3D face recognition by using the GH-EXIN neural network’. In: *Neural Approaches to Dynamics of Signal Exchanges*. Springer, 2020, pp. 223–233.
- [28] Giansalvo Cirrincione, Pietro Barbiero, Gabriele Ciravegna, and Vincenzo Randazzo. ‘Gradient-based Competitive Learning: Theory’. In: *arXiv preprint arXiv:2009.02799* (2020).
- [29] Giansalvo Cirrincione, Gabriele Ciravegna, Pietro Barbiero, Vincenzo Randazzo, and Eros Pasero. ‘The GH-EXIN Neural Network for Hierarchical Clustering’. In: *Neural Networks* 118.34 (2020).
- [30] Giansalvo Cirrincione, Rahul R Kumar, Ali Mohammadi, Shahin Hedayati Kia, Pietro Barbiero, and Jacopo Ferretti. ‘Shallow versus Deep Neural Networks in Gear Fault Diagnosis’. In: *IEEE Transactions on Energy Conversion* (2020).
- [31] Jacob Deasy, Emma Rocheteau, Katharina Kohler, Daniel J. Stubbs, Pietro Barbiero, Pietro Liò, and Ari Ercole. ‘Forecasting ultra-early intensive care strain from COVID-19 in England’. In: *medRxiv* (2020).

- [32] Jacopo Ferretti, Pietro Barbiero, Vincenzo Randazzo, Giansalvo Cirrincione, and Eros Pasero. ‘Towards uncovering feature extraction from temporal signals in deep CNN: the ECG case study’. In: *2020 International Joint Conference on Neural Networks (IJCNN)*. IEEE. 2020, pp. 1–7.
- [33] Pietro Barbiero, Evelyne Lutton, Giovanni Squillero, and Alberto Tonda. ‘A Novel Outlook on Feature Selection as a Multi-objective Problem’. In: *International Conference on Artificial Evolution (Evolution Artificielle)*. Springer. 2019, pp. 68–81.
- [34] Pietro Barbiero, Giovanni Squillero, and Alberto Tonda. ‘Beyond coreset discovery: evolutionary archetypes’. In: *Proceedings of the Genetic and Evolutionary Computation Conference Companion*. ACM. 2019, pp. 47–48. ISBN: 978-1-4503-6748-6.
- [35] Pietro Barbiero, Giovanni Squillero, and Alberto Tonda. ‘Evolutionary discovery of coresets for classification’. In: *Proceedings of the Genetic and Evolutionary Computation Conference Companion*. ACM. 2019, pp. 1747–1754. ISBN: 978-1-4503-6748-6.
- [36] Pietro Barbiero and Alberto Tonda. ‘Fundamental Flowers: Evolutionary Discovery of Coresets for Classification’. In: *International Conference on the Applications of Evolutionary Computation (Part of EvoStar)*. Springer. 2019, pp. 550–564. ISBN: 978-3-030-16692-2.
- [37] Rita Berto, Giuseppe Barbiero, Pietro Barbiero, and Giulio Senes. ‘An individual’s connection to nature can affect perceived restorativeness of natural environments. Some observations about biophilia’. In: *Behavioral Sciences* 8.3 (2018), p. 34.
- [38] Pietro Barbiero, Andrea Bertotti, Gabriele Ciravegna, Giansalvo Cirrincione, Eros Pasero, and Elio Piccolo. ‘Supervised gene identification in colorectal cancer’. In: *Quantifying and Processing Biomedical and Behavioral Signals*. Springer. 2017, pp. 243–251. ISBN: 9783319950945. DOI: [10.1007/978-3-319-95095-2\\_23](https://doi.org/10.1007/978-3-319-95095-2_23).
- [39] Pietro Barbiero, Andrea Bertotti, Gabriele Ciravegna, Giansalvo Cirrincione, Eros Pasero, and Elio Piccolo. ‘Unsupervised Gene Identification in Colorectal Cancer’. In: *Quantifying and Processing Biomedical and Behavioral Signals*. 2017. ISBN: 9783319950945. DOI: [10.1007/978-3-319-95095-2\\_21](https://doi.org/10.1007/978-3-319-95095-2_21).
- [40] Pietro Barbiero, Andrea Bertotti, Gabriele Ciravegna, Giansalvo Cirrincione, and Elio Piccolo. ‘Neural Biclustering in Gene Expression Analysis’. In: *2017 International Conference on Computational Science and Computational Intelligence (CSCI)*. Vol. 11. 2017. ISBN: 978-1-5386-2652-8.