Guest Editorial: Non-Euclidean Machine Learning

Stefanos Zafeiriou, *Member, IEEE*, Michael Bronstein, *Fellow, IEEE*, Taco Cohen, Oriol Vinyals, Le Song, Jure Leskovec, Pietro Liò, Joan Bruna, and Marco Gori, *Fellow, IEEE*

1 Introduction

Over the past decade, deep learning has had a revolutionary impact on a broad range of fields such as computer vision and image processing, computational photography, medical imaging and speech and language analysis and synthesis etc. Deep learning technologies are estimated to have added billions in business value, created new markets, and transformed entire industrial segments. Most of today's successful deep learning methods such as Convolutional Neural Networks (CNNs) rely on classical signal processing models that limit their applicability to data with underlying Euclidean grid-like structure, e.g., images or acoustic signals. Yet, many applications deal with non-Euclidean (graph- or manifold-structured) data. For example, in social network analysis the users and their attributes are generally modeled as signals on the vertices of graphs. In biology protein-to-protein interactions are modeled as graphs. In computer vision & graphics 3D objects are modeled as meshes or point clouds. Furthermore, a graph representation is a very natural way to describe interactions between objects or signals. The classical deep learning paradigm on Euclidean domains falls short in providing appropriate tools for such kind of data. Until recently, the lack of deep learning models capable of correctly dealing with non-Euclidean data has been a major obstacle in these fields. This special section addresses the need to bring together leading efforts in non-Euclidean deep learning across all communities.

From the papers that the special received twelve were selected for publication. The selected papers can naturally fall in three distinct categories: (a) methodologies that advance machine learning on data that are represented as graphs, (b) methodologies that advance machine learning on manifold-valued data, and (c) applications of machine learning methodologies on non-Euclidean spaces in computer vision and medical imaging.

1.1 Graph Neural Networks

The paper [1] proposes a novel approach to learning Graph Neural Networks (GNNs) formulated as a constrained optimisation under a Lagrangian framework. The model parameters, as well as the node states are jointly estimated by the optimisation process which convergence is expressed by a constraint satisfaction mechanism. The constrained learning procedure can be applied recursively multiple time. That way multiple levels of abstraction can be introduced similarly to multi-layer networks.

The main idea of the paper [2] revolves around the problem of attribute-missing graphs in GNNs. In particular, the paper proposes a novel distribution matching based GNN called Structure-Attribute Transformer (SAT) for learning with attribute-missing graphs. SAT leverages structures and attributes in a decoupled scheme and achieves the joint distribution modelling of structures and attributes by distribution matching techniques. Performance of the proposed methods is assessed on both node attribute completion, as well as link prediction tasks.

The paper [3] proposes a novel mixed inductive–transductive learning framework for GNNs. Typically the parameters of GNNs are computed under a supervised (i.e., inductive) learning framework. Nevertheless, due to the way information flows across the graph, GNNs can also take advantage of transductive learning (e.g. graph with "clustered" sets of nodes), using relationships among patterns. The experiments demonstrate some interesting properties for the proposed mixed model.

The main idea of the paper [4] revolves around the problem of graph node embedding. The method proposed, coined

- *Stefanos Zafeiriou is with the Department of Computing, Imperial College London, Queen's Gate SW7 2AZ, London, U.K. E-mail: s.zafeiriou@imperial.ac.uk.*
- *Michael Bronstein is with the Department of Computing, Imperial College London, Queen's Gate SW7 2AZ, London, U.K, and also with the Institute of Computational Science, Universita della Svizzera italiana, 6900 Lugano, Switzerland. E-mail: m.bronstein@imperial.ac.uk.*
- *Taco Cohen is with Qualcomm Technologies Netherlands B.V., 6546 Nijmegen, The Netherlands. E-mail: tacos@qti.qualcomm.com.*
- *Oriol Vinyals is with Google DeepMind, London AB101, U.K. E-mail: vinyals@google.com.*
- *Le Song is with the Mohamed bin Zayed University of Artificial Intelligence, Masdar City, Abu Dhabi 51133, UAE. E-mail: dasongle@gmail.com.*
- *Jure Leskovec is with Computer Science Department, Stanford University, Stanford, CA 94305-9040 USA. E-mail: jure@cs.stanford.edu.*
- *Pietro Liò is with Computer Laboratory, University of Cambridge, William Gates Building, Cambridge CB3 0FD, U.K. E-mail: pl219@cam.ac.uk.*
- *Joan Bruna is with the Courant Institute of Mathematical Sciences, New York University, New York, NY 10012-1185 USA E-mail: bruna@cs.nyu.edu.*
- *Marco Gori is with the Department of Information Engineering and Mathematics, University of Siena, 53100 Siena, SI, Italy. E-mail: bruna@cs.nyu.edu.*

Digital Object Identifier no. 10.1109/TPAMI.2021.3129857

PINE, introduces a novel notion of partial permutation invariant set function, to capture dependencies. PINE can learn an arbitrary form of the representation function from the neighborhood, without losing any potential dependence structures, and automatically decide the significance of neighbor nodes at different distance for both homogeneous and heterogeneous graphs. Theoretical guarantees are provided for general homogeneous and heterogeneous graphs. Empirical results demonstrate the capability of PINE to produce node vectors for various learning tasks in both homogeneous and heterogeneous graph.

The paper [5] proposes a new method for learning features for graph classification problems. The method transforms arbitrary-sized graphs into fixed-sized backtrackless aligned grid structures and define a new spatial graph convolution operation associated with the grid structures. That way it proposes an alternative to bridge the theoretical gap between traditional CNN models and spatially-based GNN models. Experiments on many graph datasets demonstrate the effectiveness of the proposed model.

1.2 Manifold-valued Networks

The first paper that deals with learning on manifold-valued data is [6] which proposes a novel theoretical framework for developing deep neural networks to cope with grids of manifold-valued data inputs. The novel architecture coined ManifoldNet define manifold-valued data convolutions using the weighted Fréchet Mean (wFM). The hidden layers of ManifoldNet compute wFMs of their inputs, where the parameters of the network are to be learned. This ensures that the data remain manifold-valued as they propagate through the hidden layers. Analogous to the equivariance of convolution in Euclidean space to translations, the paper proves that wFM is equivariant to the action of the group of isometries admitted by the Riemannian manifold on which the data reside.

The next paper [7] demonstrates how ideas from stochastics on manifolds and, in particular, horizontal frame bundle flows and non-linear bridge sampling schemes, can be used in the cases of convolutions with manifold domain, and convolutions with manifold target. The paper demonstrates how horizontal flows in the frame bundle provides a direct way of quantifying the role of curvature in the non-commutativity of the convolution when using parallel transport along minimizing geodesics. The paper also discusses numerical implementations and computational aspects of the algorithms.

The last paper of this category [8] proposes a novel higher order Volterra convolutional neural network, coined VolterraNet, for data defined as samples of functions on Riemannian homogeneous spaces. Analogous to the result for traditional convolutions, the paper proves that the Volterra functional convolutions are equivariant to the action of the isometry group admitted by the Riemannian homogeneous spaces, and under some restrictions, any non-linear equivariant function can be expressed as our homogeneous space Volterra convolution, generalizing the non-linear shift equivariant characterization of Volterra expansions in Euclidean space. The efficacy of VolterraNets are demonstrated in several real data sets and in a variety of applications.

1.3 Computer Vision and Medical Imaging Applications

The first paper that falls in vision and medical imaging applications category is [9] which proposes CNN architectures for omni-directional images. A spherical polyhedron representation is introduced to sample pixels on a non-Euclidean spherical surface enabling the use of CNN structures used in conventional Euclidean 2D images. The proposed method is tested in many tasks including monocular depth estimation task where it outperforms other state-of-the-art methods applied on omni-directional images.

The next paper [10] proposes to use a novel manifold-valued Generative Adversarial Network (GAN) for generating videos displaying facial expressions. In particular, the proposed GAN generates motion on the hypersphere by learning the distribution of facial expression dynamics of the different classes of facial expressions. The experimental results demonstrate the effectiveness of the approach in generating realistic videos with continuous motion, realistic appearance and identity preservation.

The paper [11] proposes a new learnable graph pooling method for learning from multiple surface-valued data with applications to brain surface classification tasks. The proposed method introduces an intrinsic aggregation of graph nodes based on graph spectral embedding. The efficacy of the proposed method is demonstrated in a battery of tasks including subject-sex classification, regression of cortical region sizes and classification of the stages of Alzheimer's disease.

The last paper in this category is [12] which tries to answer an important question: Can "generated" data, i.e., sampled from GANs, be used for performing statistical group difference tests in cases-versus-controls group analysis on graph measurements. The paper performs such group difference analysis on representations from 3D brain image volumes such as the cortical thickness measurements on the cortical surface mesh. The empirical results demonstrate that (a) GANs may indeed serve as a useful substitute for real data for testing cases-versus-controls groups and (b) exploiting the geometric nature of the data is beneficial from both theoretical and practical perspectives.

Acknowledgements

This special section would not have been possible without the efforts and interests from all the authors who contributed their submissions. The authors would like to take this opportunity to thank them and are grateful to the reviewers for their careful and valuable comments on the submitted manuscripts and for their detailed and helpful suggestions for improvement. S. Zafeiriou is supported from EPSRC project DEFORM under Grant EP/S010203/1.

References

[1] M. Tiezzi, G. Marra, S. Melacci, and M. Maggini, "Deep constraint-based propagation in graph neural networks," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 2, pp. 727–739, Feb. 2022.

[2] X. Chen, S. Chen, J. Yao, H. Zheng, Y. Zhang, and I. W. Tsang, "Learning on attribute-missing graphs," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 2, pp. 740–757, Feb. 2022.

[3] G. Ciano, A. Rossi, M. Bianchini, and F. Scarselli, "On inductive-transductive learning with graph neural networks," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 2, pp. 758–769, Feb. 2022.
[4] S. Gui *et al.*, "Pine: Universal deep embedding for graph nodes via partial permutation invariant set functions," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 2, pp. 770–782, Feb. 2022.
[5] L. Bai, L. Cui, Y. Jiao, L. Rossi, and E. Hancock, "Learning backtrackless aligned-spatial graph convolutional networks for graph classification," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 2, pp. 783–798, Feb. 2022.
[6] R. Chakraborty, J. Bouza, J. Manton, and B. C. Vemuri, "Manifoldnet: A deep neural network for manifold-valued data with applications," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 2, pp. 799–810, Feb. 2022.
[7] S. Sommer and A. M. Bronstein, "Horizontal flows and manifold stochastics in geometric deep learning," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 2, pp. 811–822, Feb. 2022.
[8] M. Banerjee, R. Chakraborty, J. Bouza, and B. C. Vemuri, "Volterranet: A higher order convolutionalnetwork with group equivariance forhomogeneous manifolds," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 2, pp. 823–833, Feb. 2022.
[9] Y. Lee, J. Jeong, J. Yun, W. Cho, and K.-J. Yoon, "SpherePHD: Applying CNNs on 360° images with non-euclidean spherical polyhedron representation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 2, pp. 834–847, Feb. 2022.
[10] N. Otberdout, M. Daoudi, A. Kacem, L. Ballihi, and S. Berretti, "Dynamic facial expression generation on hilbert hypersphere with conditional wasserstein generative adversarial nets," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 2, pp. 848–863, Feb. 2022.
[11] K. Gopinath, C. Desrosiers, and H. Lombaert, "Learnable pooling in graph convolution networks for brain surface analysis," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 2, pp. 864–876, Feb. 2022.
[12] T. Q. Dinh *et al.*, "Performing group difference testing on graph structured data from Gans: Analysis and applications in neuroimaging," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 2, pp. 877–889, Feb. 2022.

Stefanos Zafeiriou (Member, IEEE) is currently a professor of machine learning and computer vision with the Department of Computing, Imperial College London, London, U.K. Between 2016–2020, he was a distinguishing research fellow with the University of Oulu under Finish Distinguishing Professor Programme. He cofounded and exited two startups, including Facesoft and Ariel AI. He has coauthored more than 80 journal papers mainly on novel machine learning methodologies applied to various domains published in the most prestigious journals in his field of research, including *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *International Journal of Computer Vision*, and many papers in top conferences, including CVPR, ICCV, ECCV, and ICML. He has more than 17K+ citations to his work, h-index 60. He was an associate editor and the guest editor in more then eight journals. He was the general chair of BMVC 2017. He was the guest editor of more than ten journal special issues and co-organised more than twenty workshops or special sessions on specialised computer vision and machine learning topics in top venues, including CVPR or FG or ICCV or ECCV or NeurIPS. He is currently the area chair in the top venues of his field. His students are frequent recipients of very prestigious and highly competitive fellowships, such as the Google Fellowship, the Intel Fellowship, and the Qualcomm Fellowship. He was the recipient of the Prestigious Junior Research Fellowships from Imperial College London in 2011, the President's Medal for Excellence in Research Supervision for 2016, Google Faculty Research Awards, and an Amazon AWS ML Research Award. He is an EPSRC Early Career research fellow.

Michael Bronstein (Fellow, IEEE) received the PhD degree with distinction from the Technion –Israel Institute of Technology in 2007. He joined the Department of Computing as a professor in 2018. He has been a professor with USI Lugano, Switzerland since 2010 and held visiting positions with Stanford, Harvard, MIT, TUM, and Tel Aviv University. He has has authored more than 200 papers, a book, and holds more than 35 granted patents.His research interests include theoretical and computational geometric methods for machine learning and data science, computer vision and pattern recognition to geometry processing, computer graphics, and biomedicine. He is currently a PI and ML lead with Project CETI, a TED Audacious Prize winning collaboration aimed at understanding the communication of sperm whales. During 2017–2018, he was a fellow with the Radcliffe Institute for Advanced Study, Harvard University, and since 2017, he has been a Rudolf Diesel fellow with TU Munich. He was invited as a young scientist to the World Economic Forum, an honor bestowed on 40 world's leading scientists under the age of 40. His industrial experience includes technological leadership in multiple startup companies, including Novafora, Videocites, Invision (acquired by Intel in 2012), and Fabula AI (acquired by Twitter in 2019). He joined Twitter as the head of Graph Learning Research. He was the principal engineer with Intel Perceptual Computing (2012–2019) and was one of the key developers of the Intel RealSense 3D camera technology. He was the recipient of five ERC Grants, two Google Faculty Research Awards, two Amazon ML Research Awards, Facebook Computational Social Science Award, Dalle Molle Prize, Royal Society Wolfson Merit Award, and the Royal Academy of Engineering Silver Medal. He is a member of Academia Europaea, IAPR, ELLIS, and BCS, alumnus of the Technion Excellence Program and the Academy of Achievement, and ACM distinguished speaker.

Taco Cohen received the BSc degree in theoretical computer science from Utrecht University, the MSc degree in artificial intelligence, and the PhD degree in machine learning with prof. Max Welling from the University of Amsterdam. He is currently a machine learning research scientist with Qualcomm AI Research, Amsterdam, and the co-director with ELLIS Geometric Deep Learning program. He was the co-founder of Scyfer, a company focussed on active deep learning, acquired by Qualcomm in 2017. His research interests include equivariant networks and geometric deep learning, causality, and interactive learning. He was an intern with Google Deepmind (working with Geoff Hinton) and OpenAI. He was the recipient of the 2014 University of Amsterdam thesis Prize, a Google PhD Fellowship, ICLR 2018 Best Paper Award for "Spherical CNNs", and was named one of 35 innovators under 35 in Europe by MIT in 2018.

Oriol Vynials received the PhD degree in electrical engineering and computer sciences from the University of California, Berkeley. He is currently a principal scientist with DeepMind and a team lead of the Deep Learning group. His research interests include artificial intelligence, with particular emphasis on machine learning, deep learning and reinforcement learning, and science. He was with Google Brain team. His research has been featured multiple times at the New York Times, Financial Times, WIRED, and BBC His articles have been cited over 125000 times. Some of his contributions, or TensorFlow are used in Google Translate, Text-To-Speech, and Speech recognition, serving billions of queries every day. He was the lead researcher with AlphaStar project, creating an agent that defeated a top professional at the game of StarCraft, achieving Grandmaster level, also featured as the cover of Nature. He was also a contributor to the research and papers relating to the recent AlphaFold project, an AI system providing the solution to the protein folding problem (a 50-year old grand challenge in biology). He was the recipient of the 2016 MIT TR35 innovator Award.

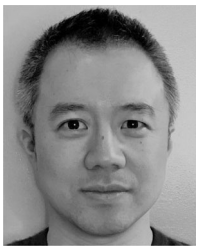

Le Song received the PhD degree in computer science from the University of Sydney and National ICT Australia. He is currently a professor and the deputy department chair with Mohamed bin Zayed University of Artificial Intelligence. Till January 2021, he was an associate professor of computational science and engineering and the associate director of Center for Machine Learning, Georgia Institute of Technology, USA. As an educator, he received honors and awards for his research and articles as a principal investigator and co-principal investigator. He developed machine learning methods and algorithms for complex and dynamic data, machine learning and cross-campus multi-disciplinary research, large scale machine learning package for Internet data, nonparametric probabilistic graphical models for complex social and biological data, and analyzed sensor time series data using kernel methods. He authored or coauthored more than 160 papers in peer-reviewed top machine learning conferences and journals, including NeurIPS, ICML, ICLR, AISTATS, and JMLR for 15 years. He was the recipient of several Best Paper Awards at the ACM Conference on Recommendation System (Recsys) in 2016, Artificial Intelligence and Statistics (AISTATS) in 2016, IEEE International Parallel Distributed Processing Symposium (IPDPS) in 2015, Neural Information Processing Systems (NeurIPS) in 2013, International Conference on Machine Learning (ICML) in 2010, the recipient of the National Science Foundation CAREER Award in 2014, Outstanding Junior Faculty Research Award in 2014, and Lockheed Martin Inspirational Young Faculty Award in 2014. He was invited as the area chair in many international conferences. He is also an active member of several regional and international groups in the field of machine learning, artificial intelligence, and statistics, sincluding being a board member of the International Conference on Machine Learning.

Jure Leskovec received the bachelor's degree in computer science from University of Ljubljana, Slovenia and the PhD degree in machine learning from Carnegie Mellon University. He did postdoctoral training with Cornell University. is currently an associate professor of computer science with Stanford University where he is a member of the InfoLab and the AI Lab. He is also an investigator with Chan Zuckerberg Biohub, where he focus on developing new methods for analysis of biomedical data. His research interests include applied machine learning, data science for large interconnected systems, modeling complex, richly-labeled relational structures, graphs, and networks for systems at all scales, from interactions of proteins in a cell to interactions between humans in a society, applications include commonsense reasoning, recommender systems, computational social science, and computational biology with an emphasis on drug discovery. This research was the recipient of several awards including a Lagrange Prize, Microsoft research faculty fellowship, the Alfred P. Sloan fellowship, and numerous best paper and test of time awards. It has also been featured in popular press outlets, including the New York Times and the Wall Street Journal.

Pietro Liò received the PhD degree in complex systems and non linear dynamics from the School of Informatics, dept of Engineering, University of Firenze, Italy and the PhD degree in theoretical genetics from the University of Pavia, Italy. He is currently a professor of computational biology with the Department of Computer Science and Technology, University of Cambridge and a member of the Artificial Intelligence Group. He is also a member of the Cambridge Centre for AI in medicine. His research interests include systems biology, multi omic, and clnical data integration methodology, developing artificial intelligence and computational biology models to detect biomarkers for diseases complexity and address personalised, and precision medicine.

Joan Bruna received the BS degree degree in mathematics and electrical engineering in 2002 from the Universitat Politècnica de Catalunya, Barcelona, Spain, the MS degree in applied mathematics in 2005 from the Ecole normale supérieure Cachan, France, and the PhD degree in applied mathematics in 2013 from the Ecole Polytechnique, Palaiseau, France. He was a postdoctoral researcher with Courant Institute, New York University (NYU) and a postdoctoral fellow with Facebook AI Research, Menlo Park, California. In 2015, he became an assistant professor with the University of California, Berkeley Statistics Department and starting in the fall of 2016, he joined the Courant Institute, NYU, as an assistant professor in computer science, data science, and mathematics (affiliated), where currently is an associate professor. His research interests include invariant signal representations, high-dimensional statistics and stochastic processes, and deep learning and its applications to signal processing.

Marco Gori (Fellow, IEEE) received the PhD degree in 1990 from Università di Bologna, Italy. He is currently with the School of Computer Science, McGill University, Montreal. In 1992, he became an associate professor of computer science with Università di Firenze and in November 1995, he joined the Università di Siena, where he is currently a full professor of computer science, where he is leading the Siena Artificial Intelligence Lab. He has authored or coauthored the monograph "*Machine Learning: A constraint-based approach*," (MK, 560 pp., 2018), which contains a unified view of his approach. His research interests include machine learning with applications to pattern recognition, web mining, game playing, and bioinformatics. His pioneering role in neural networks has been emerging especially from the recent interest in Graph Neural Networks, that he contributed to introduce in the seminal paper "Graph Neural Networks," IEEE-TNN, 2009. He was the chair of the Italian Chapter of the IEEE Computation Intelligence Society and the president of the Italian Association for Artificial Intelligence. He was one the first people involved in European project on Artificial Intelligence CLAIRE and he is currently a fellow of Machine Learning association ELLIS. He is currently with the scientific committee of ICAR-CNR and is the president of the Scientific Committee of FBK-ICT. He is currently the International 3IA chair with Université Côte d'Azur. He is a fellow of EurAI and IAPR.