Danilo P. Mandic, IEEE Fellow (2023-2025)

Professor of Signal Processing
Department of Electrical and Electronics Engineering
Imperial College
Homepage https://www.imperial.ac.uk/people/d.mandic
e-mail: d.mandic@imperial.ac.uk

Topic 1: Machine Intelligence on Graphs

Abstract: The area of Machine Intelligence on graphs promises a paradigm shift in Data Analtyics, as we approach information processing of new classes of data which are typically acquired on irregular but structured domains (such as social networks, various ad-hoc sensor networks). Yet, despite the long history of Graph Theory, current approaches tend to focus on aspects of optimisation of graphs themselves rather than on eliciting strategies relevant to the objective application of the graph paradigm, such as detection, estimation, statistical and probabilistic inference, clustering and separation from signals and data acquired on graphs. In order to bridge this gap, we first revisit graph topologies from a Data Analytics point of view, to establish a taxonomy of graph networks through a linear algebraic formalism of graph topology (vertices, connections, directivity). This serves as a basis for spectral analysis of graphs, whereby the eigenvalues and eigenvectors of graph Laplacian and adjacency matrices are shown to convey physical meaning related to both graph topology and higher-order graph properties, such as cuts, walks, paths, and neighbourhoods. We also review main trends in graph neural networks (GNN) and graph convolutional networks (GCN) from the perspective of graph signal filtering. Particular insight is given to the role of diffusion processes over graphs, to show that GCNs can be understood from the graph diffusion perspective. Given the largely heuristic nature of the existing GCNs, their treatment through graph diffusion processes may also serve as a basis for new designs of GCNs. Tensor representation of lattice-structured graphs is next considered, and it is shown that tensors (multidimensional data arrays) can be treated as a special class of graph signals, whereby the graph vertices reside on a high-dimensional regular lattice structure. Finally, the concept of graph tensor networks is shown to provide a unifying framework for learning of big data on irregular domains. Examples on dimensionality reduction, image clustering, metro traffic modelling and financial modelling support the material.

Topic 2: Machine Intelligence for eHealth: Hearables for 24/7 Doctorless Hospitals?

Abstract: Commercial wellbeing and gaming applications, together with future health systems, require the means to assess the neural and physiological function of a user from readily available data. Ideally, this should be achieved in a 24/7 fashion, in the community, and in a self-administred, discreet, and unobtrusive fashion. The Hearables paradigm, that is, in-ear sensing of neural function and vital signs is such an emerging solution. The talk introduces our own Hearables device, which is based on an earplug with the embedded electrodes, optical, acoustic, mechanical

and temperature sensors. We show how such a miniaturised embedded system can be can used to reliably measure the Electroencephalogram (EEG), Electrocardiogram (ECG), pulse, respiration, temperature, blood oxygen levels, and behavioural cues. Unlike standard wearables, such an inconspicuous Hearables earpiece benefits from the relatively stable position of the ear canal with respect to vital organs to operate robustly during daily activities. However, this comes at a cost of weaker signal levels and exposure to noise. This opens novel avenues of research in Machine Intelligence for eHealth, with a number of challenges and opportunities for algorithmic solutions. We further demonstrate how combining data from multiple sensors within such an integrated wearable device improves both the accuracy of measurements and the ability to make sense from artefacts in real-world scenarios. The ability to stream neural and physiological data in real time also makes Hearables a viable solution for the integration with smart environments and in future eHealth. The material is supported by a Hearables-specific "end-to-end" approach which revolves around fully interpretable domain knowledge, starting from the Biophysics of the generation and propagation of physiological signals on human body, through to the sensor-skin-hardware interface, and high-level decision making with a user in the loop.

Topic 3: Tensor Decompositions for Big Data Applications: Low-complexity Deep Neural Networks

Abstract: The exponentially increasing availability of big and streaming data comes as a direct consequence of the rapid development and widespread use of multi-sensor technology. The quest to make sense of such large volume and variety of data has both highlighted the limitations of standard flat-view matrix models and the necessity to move toward more versatile data analysis tools. One such model which is naturally suited for data of large volume, variety and veracity are multi-way arrays or tensors. The associated tensor decompositions have been recognised as a viable way to break the "Curse of Dimensionality", an exponential increase in data volume with the tensor order. Owing to a scalable way in which they deal with multi-way data and their ability to exploit inherent deep data structures when performing feature extraction, tensor decompositions have found application in a wide range of disciplines, from very theoretical ones, such as scientific computing and physics, to the more practical aspects of machine learning and signal processing. It is therefore both timely and important for a wider Data Analytics community to become acquainted with the fundamentals of such techniques. We show that for data which exhibit an underlying structure, this can lead to orders of magnitude savings in the storage and computational complexity (super-compression), together with exploring natural data separability by virtue of multilinear tensor algebra. This allows us to reconsider deep neural networks, which are often composed of thousands of nodes and millions of learning parameters, from the perspective of supercompression that is inherent to tensor networks, and in this way dramatically reduce their computational complexity, at a negligible loss of accuracy. Finally, this talk exploits the convergence of concepts and ideas in tensor decompositions and data analytics on graphs. It is demonstrated how their conjoint treatment provides feasible means to mitigate the curse of dimensionality associated with both the big data and deep learning paradigms – yielding reduced parameter optimised neural networks.

Topic 4: Interpretable Convolutional NNs and Graph CNNs: Role of Domain Knowledge

Abstract: The success of deep learning (DL) and convolutional neural networks (CNN) has also highlighted that NN-based analysis of signals and images of large sizes poses a considerable challenge, as the number of NN weights increases exponentially with data volume – the so called Curse of Dimensionality. In addition, the largely ad-hoc fashion of their development, albeit one reason for their rapid success, has also brought to light the intrinsic limitations of CNNs - in particular those related to their black box nature. To this end, we revisit the operation of CNNs from first principles and show that their key component – the convolutional layer – effectively performs matched filtering of its inputs with a set of templates (filters, kernels) of interest. This serves as a vehicle to establish a compact matched filtering perspective of the whole convolution-activation-pooling chain, which allows for a theoretically well founded and physically meaningful insight into the overall operation of CNNs. This is shown to help mitigate their interpretability and explainability issues, together with providing intuition for further developments and novel physically meaningful ways of their initialisation. Such an approach is next extended to Graph CNNs (GCNNs), which benefit from the universal function approximation property of NNs, pattern matching inherent to CNNs, and the ability of graphs to operate on nonlinear domains. GCNNs are revisited starting from the notion of a system on a graph, which serves to establish a matched-filtering interpretation of the whole convolution-activation-pooling chain within GCNNs, while inheriting the rigour and intuition from signal detection theory. This both sheds new light onto the otherwise black box approach to GCNNs and provides well-motivated and physically meaningful interpretation at every step of the operation and adaptation of GCNNs. It is our hope that the incorporation of domain knowledge, which is central to this approach, will help demystify CNNs and GCNNs, together with establishing a common language between the diverse communities working on Deep Learning and opening novel avenues for their further development.