

Tutorial Proposal on Sustainable Deep Learning for Time-series: Reservoir Computing

Andrea Ceni, Claudio Gallicchio, Xavier Hinaut, Gianluca Milano, and Abigail Morrison

Title: Sustainable Deep Learning for Time-series: Reservoir Computing

Abstract

Reservoir Computing (RC) is a promising approach for time-series forecasting that offers efficient and sustainable deep learning alternatives to traditional neural networks. This tutorial provides an introduction to sustainable deep learning with RC, including its theoretical foundations, practical implementation, and advantages over other approaches. The tutorial starts with an overview of the RC framework, including the Echo State Network (ESN) and Liquid State Machine (LSM) models. The fundamental principles of these models are explored, including their unique network structure, the learning mechanisms, and the fundamentals of RC mathematical background. The tutorial then proceeds to explore the strong links with computational neuroscience, and the exciting recent advances in neuromorphic hardware implementations of reservoirs. Finally, the tutorial covers practical implementation aspects, existing software libraries, as well as successful application examples in several domains including automotive, avionics, health, and robotics.

In addition to providing practical guidance, this tutorial also emphasizes the importance of sustainability in deep learning. The environmental impact of training deep neural networks is substantial, and RC offers a more energy-efficient approach that is particularly well-suited for time-series forecasting. By using RC, researchers can reduce the carbon footprint of their research while achieving state-of-the-art results.

Overall, this tutorial provides a comprehensive introduction to sustainable deep learning with RC for time-series forecasting, and it is suitable for both beginners and advanced researchers looking to explore this exciting area of machine learning.

Description of the tutorial

Reservoir Computing (RC) is a paradigm for designing recurrent neural models for sequential data processing, distinctly featured by efficient training algorithms. The approach is rooted in the analysis of the dynamical components of the neural architecture under the perspective of dynamical systems theory, and proposes to use fixed (possibly randomized) connections under asymptotic stability conditions. The resulting reservoir dynamics provide a rich encoding of the external input signals, enabling to reduce the learning problem to a simple regression performed by a linear model. The approach stands out in the context of Machine Learning and Deep Learning research for its distinctive and striking advantage in terms of computational efficiency, ensuring fast and sustainable training algorithms for time-series AI applications. Apart from the undoubted computational advantage, RC is appealing for a number of reasons, including its neurobiological plausibility, its amenability to implementations in unconventional neuromorphic hardware (e.g., in photonics), and the fact that it leads to clean mathematical analysis of the architectural biases of recurrent models. Intriguingly, RC gives a refreshed perspective over diverse topics in Deep Learning research, including learning beyond Backpropagation, stable neural architectures, the lottery ticket hypothesis, tiny ML, and neural architecture search.

The tutorial will offer a comprehensive view of the field, from basics to the most recent advances on the machine learning side, the starting and current trends of RC in Computational Neuroscience, the recent developments in neuromorphic hardware, and will survey current libraries of RC with a detailed introduction to ReservoirPy with Python Notebooks available online.

This tutorial is targeted to both researchers and practitioners that are interested in setting up fastly-trained recurrent neural networks for structured data. It will open their mind to a computing paradigm different from the most well-known ones, which offers state-of-the-art performances on several applications. Such a paradigm should be considered as its own and to implement mixtures of approaches in order to tackle seriously the environmental challenges. It is also recognized in neuroscience as an important principle to interpret and model brain computations. Basic concepts on Machine Learning and Deep Neural Networks are suggested prerequisites.

Relevance to the ECML community

This tutorial will provide an intriguing and broad overview on the emerging topic of efficiently trainable neural networks architectures, ranging from the enabling mathematical theory and the available software libraries, to computational neuroscience evidences and neuromorphic implementations in unconventional physical substrates. In this sense, the diverse backgrounds of the presenters, along with the designed format for the tutorial (featuring relatively short presentations), will be key factors in providing a stimulating and effective introduction to the subject.

The focus on energy expenditure, environmental and economic impact related to the development of deep learning algorithms is becoming increasingly important in the European and global landscape. We thereby believe that the issues addressed in this tutorial are extremely relevant to the ECML community.

Outline of the tutorial

This tutorial is structured in chapters, according to the following organization. Note that a 30 minutes break is expected between Chapters 3 and 4.

- **Chapter 1: Introduction** (10 minutes)
Preliminaries on Deep Learning for sequential data and Recurrent Neural Networks (RNNs); green AI and the problem of designing sustainable Deep Learning algorithms.
- **Chapter 2: Basics** (30 minutes)
Randomization in Deep Neural Networks; Echo State Networks, Liquid State Machines, and deep architectures; quality of dynamics and network topologies; training algorithms; Reservoir Computing for graphs.
- **Chapter 3: Mathematical foundations** (40 minutes)
Recurrent Neural Networks as input-driven dynamical systems. Stability of recurrent neural networks; Memory-nonlinearity trade-off; Edge-of-chaos dynamics.
- **Chapter 4: Computational Neuroscience of Reservoir Computing** (40 minutes)
Recent advances in Reservoir Computing in Neuroscience; Paradigm shift in the last decade from brain data analysis towards the RC paradigm of “decoding” complex non-linear computations.
- **Chapter 5: Neuromorphic Hardware implementations** (40 minutes)
Principles of neural networks in unconventional and neuromorphic hardware; Physical substrates for neuromorphic reservoirs: photonics, memristors, spintronics, and mechanical systems.
- **Chapter 6: Software libraries and applications** (40 minutes)
Python frameworks for Reservoir Computing; hyperparameter exploration; building complex architectures; example applications (e.g., sound classification with implicit segmentation) and tricks of the trade. Short hands-on lab demonstration (ReservoirPy): Implementing Echo State Networks in Python for time-series classification.
- **Chapter 7: Quo vadis?** (10 minutes)
Conclusions and outlook.

Instructors

Andrea Ceni

Affiliation: Department of Computer Science, University of Pisa, Italy;

email: andrea.ceni@di.unipi.it

Presenter for: Chapter 3 (Mathematical Foundations)

Short bio: Andrea Ceni received the M.Sc. degree in Mathematics cum laude from Università degli Studi di Firenze, Italy, in 2017, and the Ph.D. degree from the Department of Computer Science of the University of Exeter, UK, in 2021. He has been a Postdoctoral Research

Associate at CEMPS, University of Exeter. Currently, he is a Research Fellow at the Department of Computer Science of the University of Pisa, Italy. His research interests include recurrent neural networks, deep learning, computational neuroscience, chaos theory, and complex systems.

Claudio Gallicchio

Affiliation: Department of Computer Science, University of Pisa, Italy;

email: claudio.gallicchio@unipi.it

Presenter for: Chapters 1 (Introduction), 2 (Basics), 6 (Quo vadis?)

Short bio: Claudio Gallicchio is a Tenured-track Assistant Professor of Machine Learning at the Department of Computer Science of the University of Pisa, Italy. He received his PhD from the University of Pisa, where he focused on Reservoir Computing models and theory for structured data. His research is based on the fusion of concepts from Deep Learning, Recurrent Neural Networks, and Randomized Neural Systems. He is the founder and former chair of the IEEE CIS Task Force on Reservoir Computing, and a member of the IEEE CIS Data Mining and Big Data Analytics Technical Committee, of the IEEE CIS Task Force on Deep Learning, and of the IEEE Task Force on Learning for structured data. He (co-)authored several publications at major ML conferences, including ICML, AAAI, and IJCNN/WCCI. He has (co-)organized several events (special sessions, workshops, and tutorials) focused on Reservoir Computing and Randomized Neural Networks at ML international conferences (including AAAI, IJCNN/WCCI, ICANN, ESANN). He serves as a member of several program committees of conferences and workshops in ML and AI (including NeurIPS, ICML, AISTATS, ICLR, IJCAI-PRICAI, ECML-PKDD, ICONIP, IJCNN/WCCI). He served as guest editor for Special Issues on topics related to Reservoir Computing in ML journals including IEEE Transactions on Neural Networks and Learning Systems (TNNLS) and Cognitive Computation (Springer). Currently, he is an Academic Editor of PLOS ONE, and an Associate Editor of TNNLS.

Xavier Hinaut

Affiliation: Inria, Bordeaux, France;

email: xavier.hinaut@inria.fr

Presenter for: Chapter 6 (Software libraries and applications)

Short bio: Xavier Hinaut is Research Scientist in Computational Neuroscience at Inria, Bordeaux, France. His work is at the frontier of neurosciences, machine learning, robotics and linguistics: from the modeling of human sentence processing to the analysis of birdsongs and their neural correlates. He manages the DeepPool ANR project on human sentence modeling with Deep Reservoirs architectures. He leads ReservoirPy development: a new Python library for Reservoir Computing. <https://github.com/reservoirpy/reservoirpy>

Gianluca Milano

Affiliation: INRiM, Torino, Italy;

email: g.milano@inrim.it

Presenter for: Chapter 5 (Neuromorphic Hardware implementations)

Short bio: Gianluca Milano is currently a permanent researcher at the Italian National Institute of Metrological Research (INRiM). He received a Ph.D. in Physics cum laude from Politecnico di Torino, Italy, in collaboration with the Italian Institute of Technology (IIT). His main research interests and activities focus on: *i)* the investigation of electronic and ionic transport properties and physicochemical phenomena in nanodevices and low dimensional

systems; and *ii*) memristive devices and architectures for memory and neuromorphic applications, from material synthesis to device characterization, modeling, and implementation of unconventional and brain-inspired computing paradigms in neuromorphic architectures. He is coordinator of the european project EMPIR MEMQuD (<https://memqud.inrim.it/>) that involves 15 european partners, including Universities, research centers and industries, that focus on the development of memristive devices working in the quantum regime for quantum and neuromorphic applications. He is member of the Editorial advisory Board for the journal APL Machine Learning. For his work on *in-materia* implementation of reservoir computing in self-organizing networks of nano objects he has received the NEST prize for Nanoscience 2021.

Abigail Morrison

Affiliation: Institute of Neuroscience and Medicine (INM-6), Research Centre Jülich, and Computer Science 3 - Software Engineering, RWTH Aachen, Germany

email: morrison@fz-juelich.de

Presenter for: Chapter 4 (Computational Neuroscience of Reservoir Computing)

Short bio: Abigail Morrison has a background in physics and AI, which led her ineluctably to the field of computational neuroscience. Her primary research focus is uncovering the brain's computational principles, with particular emphasis on representation and learning, through the development of spiking neural network models. Her recent work examines the limits and potential of reservoir computing as a model for understanding cortical computation. A secondary focus is on high-performance simulation technology and its associated software ecosystem. Following research positions at the Bernstein Center Freiburg (Germany) and RIKEN Brain Science Institute (Japan), she held professorships at the University of Freiburg and the Ruhr University of Bochum (Germany), before being appointed to the Department of Computer Science at RWTH Aachen in 2020.

Previous Venues

Similar tutorials/seminars, covering a part of the topics of the proposed one for ECML, have been previously given in the following venues.

[AAAI 2021] “Deep Randomized Neural Networks”, Tutorial of the 35th AAAI Conference on Artificial Intelligence (AAAI), 2-9 February 2021, Virtual Conference. Speaker: Claudio Gallicchio. Audience size: around 50 people. Duration: 1.5 hours. Conference program link and tutorial material: <https://aaai.org/Conferences/AAAI-21/aaai21tutorials/#mq3> and https://sites.google.com/site/cgallicch/resources/tutorial_DRNN

[IJCNN 2021] “Reservoir Computing: Randomized Recurrent Neural Networks”, Tutorial of the International Joint Conference on Neural Networks (IJCNN), 18-22 July 2021, Virtual Conference. Speaker: Claudio Gallicchio. Audience size: around 50 people. Duration: 2 hours. Conference program link: <https://www.ijcnn.org/tutorials> Tutorial recording on YouTube: <https://www.youtube.com/watch?v=XJg7VdN7g-0>

[NMP2021 - 11th Optoelectronics and Photonics Summer School] “Neuromorphic Computing for Neural Networks”, 20-26 June 2021, Monte Bondone, Italy. Speaker: Claudio

Gallicchio. Audience size: around 100 people. Duration: 2 hours. Program of the school: <https://event.unitn.it/nmp2021/>.

Technical equipment

The tutorial requires standard conference equipment, including a video projector, microphones, and usb-c or hdmi cables for slides presentation.