Deep Learning for Physiological Signals (3380200)

Lecturer: Prof. Joachim Behar

Teaching assistants: Raphael Judkiewicz, Doron Hanuka

Introduction:

Continuous physiological signals are important in medicine because they enable real-time monitoring of patients' vital signs, such as heart rate and blood pressure. They help clinicians detect clinically relevant events or changes in a patient's condition and intervene when necessary. We define continuous physiological signals as parameters recorded without interruption, with very short intervals between samples (milliseconds or seconds).

Deep learning (DL) has been successfully applied to the analysis of these time series, including architectures originally developed for computer vision and natural language processing, as well as self-supervised methods, graph neural networks, and more.

This course aims to equip students with the skills needed to extract meaningful medical insights from continuous physiological data, an essential capability for advancing medical research and improving clinical care. Students will learn how machine learning is used in physiological time-series analysis. We will review common sources of continuous physiological signals in medical practice and explore deep learning techniques for analyzing them.

Topics will include Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs), such as Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) models, as well as Transformers for sequence modeling. We will also introduce recent approaches to self-supervised learning and discuss how resulting "foundation models" can be leveraged to improve performance. Finally, we will emphasize strategies for developing models with strong generalization capabilities.

Course plan:

Week	Topic
1	Physiological time series in medicine
2	High performance computing
3	Sequence modelling using RNN
4	Sequence modelling using Transformers
5	Domain adaptation and generalization
6	Self-supervised learning and foundation models
7	Time to event modelling
8	Reading group (five presentations)
9	Reading group (five presentations)
10	Reading group (five presentations)
11	Reading group (five presentations)
12	----
13	Final Projects Presentations.

Reading list:

Publications associated with the material presented in the lectures.

Sequence	Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. Advances in neural information processing systems, 30.
	Gulati, Anmol, et al. "Conformer: Convolution-augmented transformer for speech recognition." *arXiv preprint arXiv:2005.08100* (2020).
	Gu, Albert, and Tri Dao. "Mamba: Linear-time sequence modeling with selective state spaces." First conference on language modeling. 2024.
Generaliz.	Ganin, Yaroslav, and Victor Lempitsky. "Unsupervised domain adaptation by backpropagation." *International conference on machine learning*. PMLR, 2015.
	Dwibedi, Debidatta, et al. "With a little help from my friends: Nearest-neighbor contrastive learning of visual representations." *Proceedings of the IEEE/CVF international conference on computer vision*. 2021.
Self-supervised	Devlin J, et al. Bert: Pre-training of deep bidirectional transformers for language understanding. InProceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 2019 Jun (pp. 4171-4186).
	Chen, Ting, et al. "A simple framework for contrastive learning of visual representations." International conference on machine learning. PmLR, 2020.
	He, Kaiming, et al. "Masked autoencoders are scalable vision learners." Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. 2022.
	Caron, Mathilde, et al. "Emerging properties in self-supervised vision transformers." *Proceedings of the IEEE/CVF ICCV*. 2021. Also see later developments: DinoV2, DinoV3

List of publications for students' presentations:

This is a list of recent publications leveraging some of the deep learning topics covered in the lectures and modelling continuous physiological time series.

Sequence modelling	Phan, Huy, et al. "SeqSleepNet: end-to-end hierarchical recurrent neural network for sequence-to-sequence automatic sleep staging." IEEE Transactions on Neural Systems and Rehabilitation Engineering 27.3 (2019): 400-410.
	Phan, Huy, et al. "Sleeptransformer: Automatic sleep staging with interpretability and uncertainty quantification." *IEEE Transactions on Biomedical Engineering* 69.8 (2022): 2456-2467.
	Zhang, Andrew H., et al. "Mamba-based Deep Learning Approaches for Sleep Staging on a Wireless Multimodal Wearable System without Electroencephalography." *arXiv preprint arXiv:2412.15947* (2024).
Generalization	Ozyurt, Yilmazcan, Stefan Feuerriegel, and Ce Zhang. "Contrastive learning for unsupervised domain adaptation of time series." *arXiv preprint arXiv:2206.06243* (2022).
	Levy, Jeremy, et al. "DUDE: Deep Unsupervised Domain adaptation using variable nEighbors for physiological time series analysis." *Physiological Measurement* (2023).
Self-supervised	Ding, Zhengyao, et al. "AI modeling photoplethysmography to electrocardiography useful for predicting cardiovascular disease." *npj Digital Medicine* (2025).
	Zhang, Xiang, et al. "Self-supervised contrastive pre-training for time series via time-frequency consistency." *Advances in neural information processing systems* 35 (2022): 3988-4003.
	Li, Jun, et al. "An electrocardiogram foundation model built on over 10 million recordings with external evaluation across multiple domains." *arXiv preprint arXiv:2410.04133* (2024).

List of research projects

Research projects will focus on SSL or sequence modelling but should integrate consideration on generalization performance.

Focus	Topic	Starting point
SSL	Benchmark of 12-lead ECG foundation pre-trained models on a set of downstream tasks.	PhysioNet Challenge
SSL	Blood pressure estimation from photoplethysmography (PPG) leveraging PPG foundation models.	GitHub VitalDB
Sequence	EEG based sleep staging using RNN/Transformer/Mamba.	ISRUC-SLEEP Dataset
Sequence	AF detection algorithm from beat to beat using RNN/Transformer/Mamba.	SHDB-AF

Course evaluation:

Item	Percentage of course grade
Paper presentation (reading group)	30%
Research project presentation	30%
Research project report	40%