

CS-E4875 Research Project in Machine Learning, Data Science and Artificial Intelligence – Topics 2021

Topic #1: Bayesian optimisation for calibrating agent-based simulation model on urban mobility

Background: Bayesian optimisation is an iterative surrogate-based global optimisation approach that does not consider any assumptions on the functional form of a given optimisation problem, making it suitable for black-box functions. It comprises a surrogate model that maps the data to a utility, which guide the search through the objective space. As a consequence, bayesian optimisation is suitable for optimisation problems with expensive evaluation processes (in terms of resources like computation time, budget, etc), which highly correspond to agent-based simulation models on urban mobility. Calibration of such models for a city or region involves tuning from couples of hundreds to thousands parameters, along with computationally expensive simulations. Thus, number of simulations is the bottleneck towards convergence in global optima, and requires selection of both well-informative surrogate model and utility function that efficiently manage the exploration-exploitation trade-off.

The scope of the proposed research topic includes implementation/adoption of existing solutions for high-dimensional problems on calibrating mobility simulation model for a virtual city and real case-study - Tallinn. The research project aims at comparison of various approaches in terms of their efficiency in guiding the search through high-dimensional spaces.

Prerequisite:

- Programming skills and experience with R(preferable) or Python
- Basic knowledge in mathematical optimisation, machine learning and probability (bayesian paradigm)

Prerequisite:

- <https://ieeexplore.ieee.org/document/7352306>
- <https://arxiv.org/abs/1807.02811>
- <https://arxiv.org/abs/1810.03688>
- <https://www.sciencedirect.com/science/article/pii/S0167739X1830150X>
- <https://arxiv.org/abs/1902.10675>

Supervisor (name and email): Vladimir Kuzmanovski <vladimir.kuzmanovski@aalto.fi>, Jaakko Hollmén <jaakko.hollmen@aalto.fi>

Topic available: yes_one_instance

Topic available also for a group: yes (max 2)

Topic #2: Depth Guided Adaptation for Semantic Segmentation

Supervisor: Andrea Pilzer andrea.pilzer@aalto.fi

Background: Deep Learning based Computer Vision is a powerful tool nowadays, it can achieve high accuracy when trained on enough data. However, a model performance dramatically drops when deployed on data from a different domain. This is the case also for semantic segmentation models. Semantic segmentation is one of the fundamental tasks of computer vision, where each pixel of the image is assigned to the proper object/category (or in other words we are searching for sets of pixels with the same properties in the image). In addition to that, for semantic segmentation it is particularly costly to annotate all the pixels in modern HD images with the proper label.

To overcome the adaptation issue, researchers proposed several domain adaptation techniques. In this project, we will focus on the method proposed in a recent paper called DADA (<https://github.com/valeoai/DADA>). This method exploits guidance from depth estimation in order to adapt the semantic segmentation model to the new scene. The goal of the project is to familiarize with the problem and replicate some of the paper results using the code made publicly available by the authors.

Prerequisites: Interest in semantic segmentation and programming (PyTorch & Python)

Topic still available: yes_one_instance

Topic available also for a group: no

Topic #6: Scene coordinate regression for visual localization with uncertainty

Supervisor: Shuzhe Wang shuzhe.wang@aalto.fi (contact person) and Xiaotian Li xiaotian.li@aalto.fi

Background:

Visual localization aims at estimating precise six degree of freedom (6-DoF) camera pose with respect to a known environment. It is a fundamental component of many intelligent autonomous systems and applications in computer vision and robotics, e.g., augmented reality, autonomous driving, or camera-based indoor localization for personal assistants. Scene coordinate regression is one of the learning-based methods to directly regress 3D scene coordinates for visual localization without feature detection and description. It was experimentally shown that recent CNN-based scene coordinate regression methods achieve better localization performance on small-scale datasets compared to the state-of-the-art feature-based methods. In this project, we aim to incorporate an explicit model of confidence in one of the SoTA scene coordinate regression networks to predict for each pixel a distribution over the possible 3D location.

References:

- Hierarchical Scene Coordinate Classification and Regression for Visual Localization <https://arxiv.org/pdf/1909.06216.pdf>
- Learning Less is More – 6D Camera Localization via 3D Surface Regression: <https://arxiv.org/pdf/1711.10228.pdf>
- Differentiable RANSAC for Camera Localization: <https://arxiv.org/pdf/1611.05705.pdf>
- Supervising the new with the old: learning SFM from SFM: https://openaccess.thecvf.com/content_ECCV_2018/papers/Maria_Klodt_Supervising_the_new_ECCV_2018_paper.pdf

Prerequisites:

- Programming skills and experience with Python;
- Knowledge of machine learning, Computer Vision (CS-E4850 is highly recommended);
- **Bonus:** Experience with PyTorch, knowledge of statistics.

Topic still available: yes_one_instance

Topic available also for a group: no

Topic #7: Deep Learning for Natural Language Processing and Medical Applications

Background:

Recent years have witnessed many advances in natural language processing with deep learning, for example, large-scale pretrained language models, graph neural network-based models, and knowledge supervised methods. There still exist many challenges and unsolved tasks in the medical domain. Diagnosis notes contain complex diagnosis information, which includes a large number of professional medical vocabulary and noisy information such as non-standard synonyms and misspellings. Textual clinical notes are lengthy documents, usually from hundreds to thousands of tokens, leading to the difficulty of capturing long-term dependency. Thus, medical text understanding requires effective feature representation learning and complex cognitive process. This topic is about investigating deep learning-based NLP techniques for medical text understanding. Specific tasks include medical concept extraction, medical natural language inference, medical code assignment, medical entity recognition, and relation extraction. This project focuses on recent advances in medical NLP and will implement deep neural networks on one or two specific tasks to investigate the limitations and improve the performance with novel models. Besides, considering the sensitive nature of medical data, this project has a direction that uses the federated learning paradigm and improves the federated model aggregation methods. This project requires the candidate to have a solid background of deep learning and good programming skills with PyTorch and torch-based frameworks.

One of last year's research projects published a paper at the conference of ECML-PKDD (<https://arxiv.org/abs/2104.00952>)

Prerequisite: 1) knowledge of deep learning; 2) programming skills with deep learning frameworks (e.g., PyTorch); 3) experience with Latex typesetting and Linux servers.

Supervisor: [Shaoxiong Ji](mailto:shaoxiong.ji@aalto.fi) (shaoxiong.ji@aalto.fi)

Topic available: yes_one_instance

Topic available also for a group: no

Topic #8: eXplainable AI for MRI Images classification and explanation

In domains of scientific research such as medicine, forensics, finance it is important to explain the decisions made by the machine model in order to justify the outcome. This is especially true in the case of clinical decision-making. Here, in this research, the main objective is to develop a decision support system that will be able to help medical personal to detect and diagnose brain tumors using Brain MRI Images.

The research contains two steps.

- 1) Develop the Machine learning (ML) Model to detect Brain Tumor in the MRI image
- 2) Use of eXplainable Artificial method CIU (Contextual Importance or Utility) or Grad CAM to provide the explanation of the decision made by ML

Prerequisite: 1) knowledge of deep learning; 2) programming skills with deep learning frameworks (e.g., PyTorch Or Tensorflow)

Supervisor: Manik Madhikermi (Contact Person), Kary Främbling(manik.madhikermi@aalto.fi, kary.framling@aalto.fi)

Topic available: yes_one_instance

Topic available also for a group: no

Topic #9: Intrinsically Motivated Non-Player Characters via Coupled Empowerment Maximisation (CEM)

Background:

Non-Player Characters (NPCs) form an integral part of many videogames, but even the most advanced AI algorithms fall short of driving NPCs to support or challenge the player in open-ended ways. This is because existing techniques rely on the designer to specify the AI's goals across all possible situations - a task destined to fail in complex game worlds. Coupled Empowerment Maximisation (CEM) has been proposed as a formal principle to overcome this shortcoming and drive the behaviour of flexible, open-ended Non-Player Characters (NPC) across different types of videogames with minimal changes to the underlying algorithm [1,2]. The goal of the present project is to 1) re-implement the CEM algorithm in Python, provided a thorough formalisation and existing implementation in C++; and to 2) sanity-check the implementation by applying it to one testbed from published work reconstructed in the gridworld game Engine Griddly [3].

References:

[1] Guckelsberger, C., Salge, C. & Colton, S. (2016). Intrinsically Motivated General Companion NPCs via Coupled Empowerment Maximisation. Proc. Conference on Computational Intelligence and Games (CIG), 1–8

[2] Guckelsberger, C., Salge, C. & Togelius, J. (2018). New And Surprising Ways to be Mean: Adversarial NPCs with Coupled Empowerment Minimisation. Proc. Conference on Computational Intelligence and Games (CIG), 1–8.

[3] Griddly: A cross platform grid-based research environment that is designed to be able to reproduce grid-world style games. <https://github.com/Bam4d/Griddly>

Prerequisites:

Mandatory: Good knowledge of Python/Numpy/SciPy, and some interest in videogame AI.

Optional: Ability to read modern C++ (11x) is of advantage to migrate the algorithm.

Supervisor: Christian Guckelsberger (christian.guckelsberger@aalto.fi)

Topic available: yes_one_instance

Topic available also for a group: no

Topic #10: Missing Data Imputation for Supervised Learning with Variational Autoencoder (VAE)

Background:

Many real-world datasets come with missing values for various reasons. However, missing values are usually treated roughly, e.g., by removing samples/features that contain missing values or imputing missing values with a constant, which may lose crucial information, especially when the *missingness is not at random*. For example, in electronic health records, the patient with lots of missing data is likely to be healthier than the rest because the patient may not go for lab tests often. Here, the main objectives are: 1. develop sophisticated missing data imputation algorithms that can take both label and feature information into account with variational autoencoder; 2. study the identifiability of the missing mechanism. You will work on real-world datasets such as MIMIC-III or UCI datasets, and aim at an article with cooperative researchers.

Related materials:

1. How to deal with missing data in supervised deep learning? (<https://openreview.net/forum?id=jEXxzPUMYVZ>)
2. MIWAE: deep generative modelling and imputation of incomplete data sets (<http://proceedings.mlr.press/v97/mattei19a.html>)
3. not-MIWAE: Deep Generative Modelling with Missing not at Random Data (<https://openreview.net/forum?id=tu29GQT0JFy>)

Prerequisite: 1) familiar with latent variable model and Bayesian inference (e.g., CS-E4820); 2) knowledge of deep learning; 3) programming skills with deep learning frameworks (e.g., PyTorch)

Supervisor: Tianyu Cui (tianyu.cui@aalto.fi, contact person), Zhiheng Qian (zhiheng.qian@aalto.fi)

Topic available: yes_one_instance

Topic available also for a group: no

Topic #11: Multilabel Classification for Music and Audio using Deep Learning

Background:

In this project, you will work with audio and music data to develop multi-label classification systems that are robust to acoustically challenging conditions. Some applications include audio tagging, music recommendation, sound event detection, music genre classification, and other audio tasks. These tasks are challenging due to the large size of data, and the non-linear human perception of sound which is influenced by psycho-acoustics.

The goal of the project is to explore different deep learning architectures (CNNs, RNNs, Wavenet, Capsule nets, etc), as well as data augmentation techniques, to improve the performance of specific tasks for acoustically challenging scenarios. These scenarios include poor quality loudspeakers (e.g. mobile phones) or highly reverberant rooms. The methods will be evaluated on common datasets such as MSD, FMA, Jamendo or Freesound (among others). Ideally, the solutions developed will beside to participate in competitions such as DCASE, MediaEval, or MIREX.

This is a vast project with many possible tasks for enthusiastic students. Some activities could include literature search, coding of different parts (pre-processing, signal processing, deep learning, evaluation, auralization, acoustic simulations), or running experiments and collecting results.

Related materials:

- Purwins, H., B. Li, T. Virtanen, J. Schlüter, S. Chang, and T. Sainath. "Deep Learning for Audio Signal Processing." <https://doi.org/10.1109/JSTSP.2019.2908700>.
- Nam, J., K. Choi, J. Lee, S. Chou, and Y. Yang. "Deep Learning for Audio-Based Music Classification and Tagging: Teaching Computers to Distinguish Rock from Bach." <https://doi.org/10.1109/MSP.2018.2874383>.
- Choi, Keunwoo, György Fazekas, Kyunghyun Cho, and Mark Sandler. "A Tutorial on Deep Learning for Music Information Retrieval." <http://arxiv.org/abs/1709.04396>.
- Pons, Jordi, Oriol Nieto, Matthew Prockup, Erik Schmidt, Andreas Ehmann, and Xavier Serra. "End-to-End Learning for Music Audio Tagging at Scale." <http://arxiv.org/abs/1711.02520>.
- H. Daolang and Falcon-Perez R., "SSELDNET: A FULLY END-TO-END SAMPLE-LEVEL FRAMEWORK FOR SOUND EVENT LOCALIZATION AND DETECTION," DCASE2021 Challenge, Nov. 2021.
- P.-A. Grumiaux, S. Kiti, L. Girin, and A. Guérin, "A review of sound source localization with deep learning methods," arXiv [cs.SD], 2021.
- T. Kim, J. Lee and J. Nam, "Comparison and Analysis of SampleCNN Architectures for Audio Classification," in IEEE Journal of Selected Topics in Signal Processing, vol. 13, no. 2, pp. 285-297, May 2019, doi: 10.1109/JSTSP.2019.2909479. <https://ieeexplore.ieee.org/document/8681654>
- Park, D.S., Chan, W., Zhang, Y., Chiu, C., Zoph, B., Cubuk, E.D., Le, Q.V. (2019) SpecAugment: A Simple Data Augmentation Method for Automatic Speech Recognition. Proc. Interspeech 2019, 2613-2617, DOI: 10.21437/Interspeech.2019-2680. <https://arxiv.org/abs/1904.08779>

Prerequisites:

- Good knowledge of deep learning, programming skills in python.
- Bonus: Experience with PyTorch.
- Bonus: Familiarity with audio signal processing concepts such as filtering, time-frequency transforms, impulse responses, sampling frequency, pitch, etc ...

Supervisor: Ricardo, Falcon Perez (ricardo.falconperez@aalto.fi)

Topic available: yes_many_instances

Topic available also for a group: yes (max 2)

Topic #12: Quantifying uncertainty for adversarial samples for medical images

Description: Deep Learning has shown to be successful in many applications involving natural signals such as image, speech etc where traditional methods struggle to provide practical performance. However, the black box nature and lack of explainability has created a barrier in using such systems in critical applications such as medical domain. Further, recent studies show that a small change (as small as the value of one pixel) could change the predictions of DL systems with high confidence. These changes are relevant for medical imaging applications as data acquisition can always be noisy. Applying Bayesian methods to neural networks has been a widely successful solution to quantify the uncertainty of model predictions. In this project, the aim is to explore the effects of adversarial samples in deep learning based medical imaging systems and develop novel uncertainty quantifying methods for such scenarios.

References:

- Apostolidis, Kyriakos D., and George A. Papakostas. "A Survey on Adversarial Deep Learning Robustness in Medical Image Analysis." *Electronics* 10.17 (2021): 2132.
- Ma, Xingjun, et al. "Understanding adversarial attacks on deep learning based medical image analysis systems." *Pattern Recognition* 110 (2021): 107332.
- Kompa, Benjamin, Jasper Snoek, and Andrew L. Beam. "Second opinion needed: communicating uncertainty in medical machine learning." *NPJ Digital Medicine* 4.1 (2021): 1-6.
- Hirano, Hokuto, Akinori Minagi, and Kazuhiro Takemoto. "Universal adversarial attacks on deep neural networks for medical image classification." *BMC medical imaging* 21.1 (2021): 1-13.

Requirements: Knowledge in Python or Julia, Knowledge in any deep learning framework (PyTorch/Tensorflow/Flux.jl)

Supervisor: Vishnu Raj (vishnu.raj@aalto.fi)

Topic available: yes_many_instances

Topic available also for a group: yes (max 2)

Topic#13: Multimodal fusion and classification for medical applications

Descriptions: Multimodal fusion and classification is a new and emerging area. Several existing works where unimodal data are fused into the multimodal framework to capture these different modalities into a single framework. Such fusion allows making a decision very effectively. However, these methods are less explored in the medical domain. For example in the medical dialogue, the patient describes the pain or symptom using natural language (text), audio, and gesture (image), so the multi-modal information is also important to better grasp the patient situation. The aim of this project is to explore the different fusion approaches and develop end to end framework for detection tasks in medical applications.

References:

- Li, Qiuchi, et al. "Quantum-inspired neural network for conversational emotion recognition." *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 35. No. 15. 2021.

- Tan, W., Tiwari, P., Pandey, H. M., Moreira, C., & Jaiswal, A. K. (2020). Multimodal medical image fusion algorithm in the era of big data. *Neural Computing and Applications*, 1-21.
- Xu, Zhen, David R. So, and Andrew M. Dai. "MUFASA: Multimodal Fusion Architecture Search for Electronic Health Records." *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 35. No. 12. 2021.
- <https://arxiv.org/pdf/1906.00295.pdf>

Requirements:

- Experience in Python
- Experience with processing multimodal data (audio, text, visual)
- Experience with deep learning frameworks (PyTorch, Tensorflow)

Supervisor: Prayag Tiwari (prayag.tiwari@aalto.fi)

Topic available: yes_one_instances

Topic available also for a group: yes (max 1)

Topic #14: Explaining graph neural networks

Description: Graph neural networks (GNNs) have recently led to breakthroughs in many relevant tasks, such as antibiotic design, time-travel prediction, and simulation of physical systems. However, explaining predictions of GNNs remains an open challenge. This project aims to build explaining methods for GNNs, focusing on the recent class of temporal GNNs. We are interested in approaches for explaining node/edge-level predictions in both transductive and inductive settings. We expect this project to increase the transparency of (temporal) GNNs for high-stake tasks, such as drug design and personalized healthcare. Notably, the results of this project may represent a significant step towards alleviating fairness risks in graph representation learning.

Prerequisites: Good programming skills and a solid background in machine learning.

Supervisor name and email: Amauri Souza (amauri.souza@aalto.fi, contact person), Samuel Kaski (samuel.kaski@aalto.fi)

Topic available: yes_many_instances

Topic available also for a group: yes

Topic #15: Differential Privacy in Practical Implementations

Description: Applying machine learning algorithms on data containing sensitive personal information carries the risk of memorisation of specific personal information in the trained model. Examination of the model or its prediction outputs can then leak personal information. Differential privacy is a mathematical framework which allows to establish strict theoretical bounds on this leakage and has become the de-facto standard in research and practical application. However, formulations of differential privacy rarely account for the fact that real-world implementations rely on imperfect approximations of ideal math and have to contend with issues such as imperfect sources of randomness and finite-precision number representations. In this thesis you will review and summarise such problems arising from implementing ideal formulations of differential privacy on real-world computers and possible mitigation strategies.

Prerequisites: Basic understanding of probability, details of number representation, basic concepts of cryptography

Supervisor name and email: lukas Prediger (lukas.m.prediger@aalto.fi, contact person); Samuel Kaski (samuel.kaski@aalto.fi)

Topic available: yes_many_instances

Topic available also for a group: yes

Topic #16: Application of Gaussian process modeling and inference

Description: Gaussian processes (GP), seen as infinite-dimensional generalizations of multivariate normal distributions, are useful distributions on functions. Gaussian process modeling imposes a GP as a prior, which generates flexible nonlinear predictions along with their uncertainties. In this thesis, we will focus on a specific application (such as time-series data and other applications that require nonlinearity and uncertainty measures), and design sensible kernel forms to model and predict real-world data. Additionally, this thesis project also accepts a literature study, as well as practical experiments about the inference of Gaussian process and efforts to scale up Gaussian process models to large datasets.

Prerequisites: Understanding of probability, statistics and programming skills.

Supervisor name and email: Zheyang Shen (zheyang.shen@aalto.fi), contact person); Samuel Kaski (samuel.kaski@aalto.fi)

Topic available: yes_many_instances

Topic available also for a group: yes

Topic #17: Human and Machine decisions

Description: How can we improve interaction between people and decision-making AIs? The interdisciplinary field of human-machine interaction attempts to understand the process of human decision-making, builds predictive models for it and helps people to make good decisions. The need in human-machine interaction arises in behavioral economics, psychology and neuroscience with applications ranging from public policies to daily routines. For this topic, you will review papers related to human-machine interactions, focusing on challenges, applications and existing solutions for improving machine learning models with theories of decision-making and vice versa.

Prerequisites: basic understanding of reinforcement learning and probabilities.

Supervisor name and email: Alex Aushev (alexander.aushev@aalto.fi, contact person); Samuel Kaski (samuel.kaski@aalto.fi)

Topic available: yes_many_instances

Topic available also for a group: yes

Topic #18: Neural temporal point process

Description: Temporal point processes (TPPs) can be used to model discrete events occurring over continuous-time horizons. Compared to the classical TPP models, which are usually defined to handle simple event patterns, the recent works attempt to define more expressive models by combining the ideas from deep learning. This project on so-called "Neural temporal point processes" will involve surveying the current literature, discussing main design choices, and providing a comparative study of related works. Optionally, one can also perform experiments on small datasets using existing methods and their implementations.

Prerequisites: Basic understanding of probability and statistics, machine learning and programming skills.

Supervisor name and email: Pashupati Hedge (pashupati.hegde@aalto.fi, contact person); Samuel Kaski (samuel.kaski@aalto.fi)

Topic available: yes_many_instances

Topic available also for a group: yes

Topic #19: Deep learning in speech and language processing

Background: Deep learning is changing the ways how speech and language data can be processed and represented. Several specific topics are available either for experimenting with new model architectures in real-word data or applications, such as automatic speech and speaker recognition, translation and language learning. The topic can be selected together with the student.

Prerequisite: One of Aalto's basic course in speech recognition or natural language processing or corresponding knowledge. Knowledge in deep learning. Experience in scientific programming, e.g. in Python.

Supervisor name and email: Mikko Kurimo mikko.kurimo@aalto.fi and his research group

Topic available: yes_multiple_instances

Topic available also for a group: yes (max 3)

Topic #20: Visual and Multimodal Transformers and BERTs

Background: In this project, the task is to reproduce one of the recently published results relating to visual and multimodal Transformers and BERTs. The topics will be selected together with the students.

Prerequisite: knowledge of deep learning and PyTorch/TensorFlow

Supervisor name and email: Jorma Laaksonen and research group, jorma.laaksonen@aalto.fi

Language: English

Topic available: yes_multiple

Topic available also for a group: yes (max 2 students)

Topic #21: Project in Bayesian approximate inference and diagnostics

Background: In this project, the task is to implement and test one of the recently published Bayesian approximate inference methods, inference diagnostics, or develop improved visualizations for the said diagnostics. The topic can be selected together with the student. Some example topics are importance sampling diagnostics, diagnosing funnel and banana shaped posteriors, Monte Carlo standard error for arbitrary function, analysis of dynamic Hamiltonian Monte Carlo behavior, analysis of low rank black box variational inference, visualization of results from projective predictive model selection.

Prerequisites: knowledge of Bayesian methods, and R or Python

Supervisor name and email: Aki Vehtari (aki.vehtari@aalto.fi)

Topic available: yes_many_instances (max 4 instances)

Topic available also for a group: yes

Topic #22: Optical character recognition and document analysis for technical drawings

Background: For decades, technical drawings relating to built environments, such as houses and streets, have been archived in paper form by municipal authorities. Indexing and accessing these documents has been difficult as long as they have been stored in such analog format. The municipalities are transforming the documents to digital format, which makes accessing the documents easier, but as such it has not yet solved the indexing issue. Evolved methods are needed for automatically finding the particular contents of the drawings which are needed for indexing them based on the type of the drawing and on the geographical location to which it relates.

Description: The task is to study and develop methods that can be used for optical character recognition and document content analysis of technical drawings. Test material will be provided by the city of Espoo. The students are expected to find and experiment with available software and services that can be used for the task. These will include but not be limited to Tesseract. Both quantitative and qualitative results and analyses are expected.

Prerequisites: CS-E4850 Computer Vision or other related knowledge

Supervisor name and email: Jorma Laaksonen (jorma.laaksonen@aalto.fi)

Topic available: yes_many_instances (max 2 instances)

Topic available also for a group: yes (max 2 students in the group)

Topic #23: Understanding the Expressivity of Neural Network Control Policies in Reinforcement Learning

Background: Using multi-layered neural network policies is the de facto standard practice for learning-based motion control solutions. However, the relationship between the architecture of neural network backbones of control policies and their performance is poorly understood both theoretically and empirically. The student will focus on (1) designing toy problems for simple reinforcement learning setup, (2) exploring visualization techniques for understanding control granularity within state space, and (3) exploring architectures and hyperparameters of a neural network model while taking sensitivity to randomization into account. This work aims to be a hands-on introduction to deep reinforcement learning (DRL) research where the student will have an opportunity to implement state-of-the-art algorithms as well as examine various numerical footprints of neural network models and their resulting trajectories.

Related Work:

- On the Expressive Power of Deep Neural Networks. <https://proceedings.mlr.press/v70/raghu17a.html>.
- On the number of response regions of deep feed forward networks with piece-wise linear activations. <https://arxiv.org/abs/1312.6098>
- On the Number of Linear Regions of Deep Neural Networks. <https://arxiv.org/abs/1402.1869>

Supervisor: Nam Hee Gordon Kim (namhee.kim@aalto.fi)

Prerequisite: Familiar with Python for differentiable programming (e.g. PyTorch) and Pythonic ways to do object-oriented programming. Visualization techniques (e.g. Matplotlib). Experiment tracking (e.g. Weights & Biases, Neptune) and managing cluster compute jobs (e.g. SLURM).

Language: English

Topic available: yes_one_instance

Topic available also for a group: yes (max 2 students in the group)