EPFL-CIS & RIKEN-AIP Joint Online Workshop on Machine Learning
(September 7 – 8, 2022)

Abstract DAY 1: Sep. 7, 2022

Switzerland 10:00 / Japan 17:00 Jingfeng Zhang (RIKEN-AIP)
Title: Adversarial robustness: from basic science to applications

Abstract:
When we deploy models trained by standard training (ST), they work well on natural test data. However, those models cannot handle adversarial test data (also known as adversarial examples) that are algorithmically generated by adversarial attacks. An adversarial attack is an algorithm which applies specially designed tiny perturbations on natural data to transform them into adversarial data, in order to mislead a trained model and let it give wrong predictions. Adversarial robustness is aimed at improving the robust accuracy of trained models against adversarial attacks, which can be achieved by adversarial training (AT). What is AT? Given the knowledge that the test data may be adversarial, AT carefully simulates some adversarial attacks during training. Thus, the model has already seen many adversarial training data in the past, and hopefully it can generalize to adversarial test data in the future. AT has two purposes: (1) correctly classify the data (same as ST) and (2) make the decision boundary thick so that no data lie nearby the decision boundary. In this talk, I will introduce how to leverage adversarial attacks/training for evaluating/enhancing reliabilities of AI-powered tools.

Bio:
Jingfeng Zhang is a researcher in RIKEN-AIP at “Imperfect Information Learning Team” supervised by Prof. Masashi Sugiyama. Prior to RIKEN-AIP, Jingfeng obtained his Ph.D. degree (in 2020) under Prof. Mohan Kankanhalli at School of Computing in the National University of Singapore, and his Bachelor's Degree (in 2016) at Taishan College in Shandong University, China. Jingfeng is the receiver of Strategic Basic Research Programs ACT-X 2021-2023 funding, JSPS Grants-in-Aid for Scientific Research (KAKENHI), Early-Career Scientists 2022-2023, and the RIKEN Ohbu Award 2022 Jingfeng serves as a reviewer for prestigious ML conferences such as ICLR, ICML, NeurIPS, etc. Jingfeng’s long-term research interest is making artificial intelligence safe for human beings.

Switzerland 10:30 / Japan 17:30 Guillermo Ortiz-Jimenez (EPFL-CIS)
Title: Catastrophic overfitting is a bug but also a feature

Abstract:
Despite clear computational advantages in building robust neural networks, adversarial training (AT) using single-step methods is unstable as it suffers from catastrophic overfitting (CO): Networks gain non-trivial robustness during the first stages of adversarial training, but suddenly reach a breaking point where they quickly lose all robustness in just a few iterations. Although some works have succeeded at preventing CO, the different mechanisms that lead to this remarkable failure mode are still poorly understood. In this work, however, we find that the interplay between the structure of the data and the dynamics of AT plays a fundamental role in CO. Specifically, through active interventions on typical datasets of natural images, we establish a causal link between the structure of the data and the onset of CO in single-step AT methods. This new perspective provides important insights into the mechanisms that lead to CO and paves the way towards a better understanding of the general dynamics of robust model construction.

Bio:
Guillermo Ortiz-Jimenez is a fourth-year PhD student at EPFL working under the supervision of Pascal Frossard. His research focuses on understanding deep learning using empirical methods with a focus on robustness and generalization. During his PhD, Guillermo has visited the University of Oxford as part of the ELLIS PhD program, where he is co-supervised by Philip Torr. He is currently a research intern at Google in Zurich. Before starting his PhD, Guillermo received his MSc. in Electrical Engineering from TU Delft, Netherlands, and his BSc. in Telecommunications Engineering from Universidad Politécnica de Madrid, Spain. He ranked first at both institutions.

Abstract DAY 1: Sep. 7, 2022

Switzerland 11:00 / Japan 18:00 Vo Nguyen Le Duy (RIKEN-AIP)
Title: Exact Statistical Inference for the Wasserstein Distance by Selective Inference

Abstract:
The Wasserstein distance (WD), which is a metric used to compare the probability distributions, has attracted significant attention and is being used more and more in statistics and machine learning. When the WD calculated from noisy data is used for various decision-making problems, it is necessary to quantify its statistical reliability, e.g., in the form of confidence interval (CI). Several studies have been proposed in the literature, but almost all of them are based on asymptotic approximation and do not have finite-sample validity. In this study, we propose an exact (non-asymptotic) inference method for the WD inspired by the concept of conditional Selective Inference (SI). We will show that, by conditioning on the optimal coupling, the exact sampling distribution of the WD can be derived, which enables us to construct CI that has finite-sample coverage guarantee.

Bio:
Vo Nguyen Le Duy is currently a PhD student at Nagoya Institute of Technology and Junior Research Associate at RIKEN Center for Advanced Intelligence Project, Japan. He received the B.S. degree from the Danang University of Science and Technology, Vietnam, in 2017. After that, he received the M.S. degree from Nagoya Institute of Technology, Japan, in 2020. His research interests include machine learning, data mining, and statistical data analysis.

Switzerland 11:30 / Japan 18:30 Hugo Cui (EPFL-CIS)
Title: Error rates for kernel methods under source and capacity conditions

Abstract:
We investigate the rates of decay of the prediction error for kernel methods under the Gaussian design and source/capacity assumptions. For kernel ridge regression, we derive all the observable rates, and characterize the regimes in which each hold. In particular, we show that the decay rate may transition from a fast, noiseless rate to a slow, noisy rate as the sample complexity is increased. For noiseless kernel classification, we derive the rates for two standard classifiers, margin-maximizing SVMs and ridge classifiers, and contrast the two methods. In both cases, the derived rates also describe to a good degree the learning curves of a number of real datasets. This is joint work with Bruno Loureiro, Florent Krzakala and Lenka Zeborová.

Bio:
Hugo Cui is currently a PhD student in the Statistical Physics of Computation lab in EPFL, Switzerland. He previously studied theoretical physics at ENS Paris.

Switzerland 12:00 / Japan 19:00 Yu Zhe (RIKEN-AIP)
Title: Domain Generalization via Adversarially Learned Novel Domains

Abstract: We focus on the domain generalization task, which aims to learn a model that generalizes to unseen domains by utilizing multiple training domains. More specifically, we follow the idea of adversarial data augmentation, which aims to synthesize and augment training data with ”hard” domains for improving the model’s domain generalization ability. Previous works augment training data only with samples similar to the training data, resulting in limited generalization ability. We propose a novel adversarial data augmentation method, termed GADA (Generative Adversarial Domain Augmentation), which employs an image-to-image translation model to obtain a distribution of novel domains that are semantically different from the training domains, and, at the same time, hard to classify. Evaluation and further analysis suggest that adversarial data augmentation with semantically different samples leads to better domain generalization performance.

Bio:
Yu Zhe was born in AnHui, China in 1994. He received the B.S degree in computer science from the Guangdong University, China, of Foreign Studies in 2016, and the M.S degree in computer science from the Kanazawa University, Japan, in 2019. He is currently pursuing the Ph.D. degree in computer science at University of Tsukuba, Japan.

Abstract DAY 1: Sep. 7, 2022

Switzerland 12:30 / Japan 19:30 Berfin Simsek (EPFL-CIS)
Title: Neural Network Loss Landscapes: Symmetry-Induced Saddles and the Global Minima Manifold

Abstract:
In this talk, I will present a combinatorial analysis of the number of critical manifolds in neural networks as a function of overparameterization by exploiting the permutation-symmetry between the hidden layer neurons. I will first introduce a saddle point type, the so-called symmetry-induced (SI) saddles, emerging from a redundant arrangement of neurons in deep neural networks. Then I will describe the precise geometry of the global minima manifold in a teacher-student setting by giving its number of components and flatness. Similarly, counting the possible arrangements of neuron groups inside a neural network, I will also give the number of SI saddle manifolds in terms of the student and teacher widths. Our analysis shows that overparameterization gradually smoothens the landscape due to a faster scaling of the number of global minima components than the number of SI saddle manifolds for wide networks. However, in mildly overparameterized networks, the loss landscape exhibits roughness and spurious local minima near the saddle manifolds. In this regime, we empirically show that gradient-based optimizers fail to find a zero-loss solution for a fraction of random initializations.

Bio:
I am a doctoral researcher at EPFL in Computer and Communication Sciences, supervised by Clément Hongler and Wulfram Gerstner. My research focus is the theory of deep learning, in particular, neural network landscapes, training dynamics, mild overparameterization, generalization, random features, and kernel methods. I also explored out-of-distribution generalization during my internship at Meta AI. Before my Ph.D., I studied Electrical-Electronics Engineering and Mathematics double-major at Koç University in Istanbul. Before that, I competed in International Mathematical Olympiads.

Switzerland 13:00 / Japan 20:00 Yaxiong Liu (RIKEN-AIP)
Title: Expert advice problem with noisy low rank loss

Abstract:
We consider the expert advice problem with a low rank but noisy loss sequence, where the loss vector on each round is composed by the low rank part and the noisy part. This is a generalization of the works of Hazan et al. (2016) and Barman et al. (2018), where the former one only treats noiseless loss and the latter one assumes that the low rank structure is known in advance. We propose an algorithm, where during the learning process we can re-construct the kernel of the low rank part under the assumptions, that the low rank loss is noised and there is no prior information about low rank structure. With this kernel, we obtain a satisfying regret bound. Moreover, even if in experiment, the proposed algorithm performs better than Hazan's algorithm and the Hedge algorithm.

Bio:
LIU, Yaxiong is currently a postdoctoral researcher in Computational Learning Team (leading by Prof. Hatano) AIP RIKEN. He received Dr. Sci. Degree from Kyushu University in 2022. His research interests include online learning and its applications to privacy etc.

Switzerland 13:30 / Japan 20:30 Hadrien Hendrikx (EPFL-CIS)
Title: Beyond spectral gap: the role of the topology in decentralized learning.

Abstract:
In data-parallel optimization of machine learning models, workers collaborate to improve their estimates of the model: more accurate gradients allow them to use larger learning rates and optimize faster. We consider the setting in which all workers sample from the same dataset, and communicate over a sparse graph (decentralized). In this setting, current theory fails to capture important aspects of real-world behavior. First, the 'spectral gap' of the communication graph is not predictive of its empirical performance in (deep) learning. Second, current theory does not explain that collaboration enables larger learning rates than training alone. In fact, it prescribes smaller learning rates, which further decrease as graphs become larger, failing to explain convergence in infinite graphs. This paper aims to paint an accurate picture of sparsely-connected distributed optimization when workers share the same data distribution. We quantify how the graph topology influences convergence in a quadratic toy problem and provide theoretical results for general smooth and (strongly) convex objectives. Our theory matches empirical observations in deep learning, and accurately describes the relative merits of different graph topologies.

Bio:
Hadrien Hendrikx is a post-doc in the MLO team from EPFL, working with Martin Jaggi. Before that, he completed a Ph.D. at Inria Paris, under the supervision of Francis Bach and Laurent Massoulié. His research focuses on optimization for machine learning, and on decentralized methods in particular.