2025 KIAS Center for AI and Natural Sciences Spring Workshop

ABSTRACTS

May 28(Wed.) - 30(Fri.), 2025

Park Roche Resort, Jeongseon

No.1

Sobolev acceleration for neural networks

Hwijae Son

Konkuk University

Sobolev training, which integrates target derivatives into the loss functions, has been shown to accelerate convergence and improve generalization compared to conventional L2 training. However, the underlying mechanisms of this training method remain incompletely understood. In this work, we show that Sobolev training provably accelerates the convergence of Rectified Linear Unit (ReLU) networks and quantify such 'Sobolev acceleration' within the student-teacher framework. Our analysis builds on an analytical formula for the population gradients of ReLU networks under centered spherical Gaussian input. Extensive numerical experiments validate our theoretical findings and show that the benefits of Sobolev training extend to modern deep learning tasks, including diffusion models.

No.2

Estimation of System Parameters Including Repeated Cross-Sectional Data through Emulator-Informed Deep Generative Model

Sung Woong Cho

Stochastic Analysis and Application Research Center (SaaRC) /

Korea Advanced Institute of Science and Technology (KAIST)

[Abstract] Differential equations (DEs) are crucial for modeling the evolution of natural or engineered systems. Traditionally, the parameters in DEs are adjusted to fit data from system observations. However, in fields such as politics, economics, and biology, available data are often independently collected at distinct time points from different subjects (i.e., repeated cross-sectional (RCS) data). Conventional optimization techniques struggle to accurately estimate DE parameters when RCS data exhibit various heterogeneities, leading to a significant loss of information. To address this issue, we propose a new estimation method called the emulator-informed deep-generative model (EIDGM), designed to handle RCS data. Specifically, EIDGM integrates a physics-informed neural network-based emulator that immediately generates DE solutions and a Wasserstein generative adversarial network-based parameter generator that can effectively mimic the RCS data. We evaluated EIDGM on exponential growth, logistic population models, and the Lorenz system, demonstrating its superior ability to accurately capture parameter distributions. Additionally, we applied EIDGM to an experimental dataset of Amyloid beta 40 and beta 42, successfully capturing diverse parameter distribution shapes. This shows that EIDGM can be applied to model a wide range of systems and extended to uncover the operating principles of systems based on limited data.

No.3

Weight Initialization Methods for FFNNs with Diverse Activation Functions

Hyunwoo Lee

Korea Institute for Advanced Study

Weight initialization plays a critical role in training neural networks, as it directly affects signal propagation and optimization dynamics. Proper initialization can lead to faster convergence and improved model performance, while poor initialization may cause issues such as vanishing gradients or activation saturation, especially in deep architectures. Since such phenomena are highly dependent on the choice of activation function, various initialization schemes have been developed accordingly, for instance, He initialization for ReLU-based networks and Xavier initialization for sigmoidal activations. In this presentation, we briefly introduce novel initialization methods tailored for ReLU and tanh activations. We also outline future directions toward constructing activation-aware initialization frameworks that propose appropriate initial weights based on the mathematical properties of a given activation function.

No.4

Floating-Point Neural Networks Can Represent ``Almost All'' Floating-Point Functions

Geonho Hwang

Gwangju Institute of Science and Technology

Existing works on the expressive power of neural networks typically assume real-valued parameters and exact mathematical operations during the evaluation of networks. However, neural networks run on actual computers can take parameters only from a small subset of the reals and perform inexact mathematical operations with round-off errors and overflows. In this work, we study the expressive power of \emph{floating-point} neural networks, i.e., networks with floating-point parameters and operations. We first observe that for floating-point neural networks to represent \emph{all} functions from floating-point vectors to floating-point vectors, it is necessary to \emph{distinguish} different inputs: the first layer of a network should be able to generate different outputs for different inputs. We also prove that such distinguishability is \emph{sufficient}, along with mild conditions on activation functions. Our result shows that with practical activation functions, floating-point neural networks can represent \emph{almost all} floating-point functions.

No.5

Thermodynamics from neural networks: Neural Density Functional Theory

Ji Woong Yu

KIAS, CAINS

Predicting a material's macroscopic thermodynamic properties from microscopic particle interactions remains a central scientific challenge addressed by Classical Density Functional Theory (cDFT). This framework uses an unknown "excess free energy functional" to derive the "one-body direct correlation function" (c_1), which characterizes inter-particle influences. Traditionally, c_1 was approximated through "integral equation theories" with various "closure" approximations. The Percus-Yevick (PY) closure works well for hard-sphere repulsive systems, while the Hypernetted-Chain (HNC) better handles softer or longer-ranged potentials like Coulombic interactions, despite struggling with highly correlated systems. The Mean Spherical Approximation (MSA) provides analytical solutions for specific potentials by linearizing correlations beyond a core. Though these closures yielded c_1 and thermodynamic data through cDFT, accuracy depended on the approximation's suitability to the specific system. Neural Density Functional Theory (NDFT) transforms this approach by training neural networks to map particle density profiles ρ(r) directly to c_1 using simulation data, eliminating reliance on system-specific closure approximations. NDFT's key advantage lies in constructing accurate c_1 functionals from data, bypassing analytical closures' limitations. This enables superior predictions across diverse systems, captures many-body effects driving emergent phenomena, and provides a thermodynamic modeling framework that improves systematically with more training data.

No.6

Importance of equivariance and message passing in machine-learning interatomic potentials

Sungwoo Kang

Computational Science Research Center, KIST

Density functional theory (DFT) is a quantum mechanics-based computational method widely employed for accurately modeling atomic interactions and material properties. Despite its high predictive accuracy, DFT calculations are computationally expensive, severely limiting their application to large-scale realistic systems. Machine-learning interatomic potentials (MLIPs) have emerged as efficient surrogate models that learn from DFT data, enabling rapid predictions of energies and atomic forces at significantly reduced computational cost. While MLIPs have traditionally been trained on specialized datasets tailored to specific material systems, recent efforts focus on developing universal MLIPs (or foundation models) through extensive training on large and diverse datasets. Such universal potentials demonstrate broad transferability and predictive accuracy across diverse material classes without requiring extensive system-specific retraining. Therefore, major global technology corporations, such as Google, Microsoft, and META, are actively conducting research in this field [1-3].

The key features enabling the successful training of models on large datasets are the implementation of equivariance and message-passing algorithms. In this talk, I will first explain these concepts and discuss how they significantly enhance both the learning efficiency and transferability of MLIPs [4]. Then, I will briefly introduce our recent work on developing coarse-grained graph models for all-atom force fields, achieved by leveraging the equivariant nature of the recently developed MLIP architectures [5].

[1] Merchant, et al. Nature 624, 80 (2023)
[2] Yang, et al. arXiv:2405.04967 (2024)
[3] Fu et al. arXiv:2502.12147 (2025)
[4] S. Kang. J. Chem. Phys. 161, 244102 (2024)
[5] S. Kang. arXiv.2505.01058 (2025)

No.7

Assessing molecular representations for Bayesian active learning with small initial data

Min Young Ha

Department of Chemical Engineering, Kyung Hee University,

1732, Deogyeong-daero, Giheung-gu, Yongin-si, Republic of Korea, 17104

Extending data-driven methods to the small-data limit is essential for accelerating the early stages of materials discovery, where only a small amount of proof-of-concept data is available. This work proposes a two-stage Bayesian active learning workflow capable of operating with small initial data, and the impact of molecular representations on their downstream performance is analyzed. After mapping molecular graphs to the embedded vector space on a hypersphere, the quality of the embeddings is assessed by alignment and uniformity of the representation space, which describes the representation robustness and information content, respectively. Different graph embedding methods are compared, and the explorative and predictive performance in Bayesian active learning iterations is attributed to the alignment and uniformity of the representation space. This work provides a broadly applicable small-data active learning framework and a decision strategy to select among numerous candidate approaches.

No.8

Overcoming Fake Solutions in Semi-Dual Neural Optimal Transport: A Smoothing Approach for Learning the Optimal Transport Plan

Jaewoong Choi

Sungkyunkwan University

We address the convergence problem in learning the Optimal Transport (OT) map, where the OT Map refers to a map from one distribution to another while minimizing the transport cost. Semi-dual Neural OT, a widely used approach for learning OT Maps with neural networks, often generates fake solutions that fail to transfer one distribution to another accurately. We identify a sufficient condition under which the max-min solution of Semi-dual Neural OT recovers the true OT Map. Moreover, to address cases when this sufficient condition is not satisfied, we propose a novel method, OTP, which learns both the OT Map and the Optimal Transport Plan, representing the optimal coupling between two distributions. Under sharp assumptions on the distributions, we prove that our model eliminates the fake solution issue and correctly solves the OT problem. Our experiments show that the OTP model recovers the optimal transport map where existing methods fail and outperforms current OT-based models in image-to-image translation tasks. Notably, the OTP model can learn stochastic transport maps when deterministic OT Maps do not exist, such as one-to-many tasks like colorization.

No.9

Distributional Learning for Tabular Data Synthesis

Seunghwan An

Dept. of Information and Telecommunication Engineering

Incheon National University

Tabular data is one of the most prevalent data types in real-world applications, spanning domains such as healthcare, finance, e-commerce, and public services. As the demand for privacy-preserving and shareable data continues to grow, synthetic tabular data generation has emerged as a promising solution. This talk introduces a range of frameworks for tabular data synthesis grounded in distributional learning, where the objective is to estimate the underlying data distribution and generate realistic samples from it. We begin by outlining the unique challenges posed by tabular data, particularly the presence of heterogeneous feature types, which often hinder the direct application of standard generative models. To address these challenges, we explore various nonparametric distributional learning methods, including quantile function estimation, histogram-based estimation, variational autoencoders, (any-order) autoregressive modeling, and their combinations. We further discuss three essential criteria for evaluating synthetic data: (1) statistical similarity to real data, (2) utility in downstream machine learning tasks, and (3) preservation of privacy. The talk concludes with ongoing research directions, such as constraint-aware generation and causal modeling, aimed at improving the realism and reliability of synthetic tabular data.

No.10

On the Foundations of Machine Learning

Jinsook Kim

Center for AI and Natural Sciences

We critically examine three theories of machine learning and propose a new theory, according to which machines learn a function when they successfully compute it. We argue that our theory challenges established assumptions in statistical and computational learning theories. Specifically, it implies that learning true probabilities is not equivalent either to correctly calculating those probabilities or to achieving an almost-sure convergence to them. We defend our arguments further while engaging with literature on computable learning.

No.11

Modeling anharmonic baths with reservoir computing for vibronic dynamics

Kwang Hyun Cho

Center for AI and Natural Sciences, KIAS

Vibronic dynamics, or coupled electronic-nuclear dynamics provides useful framework to understanding a wide range of chemical reactions. Often, molecular motion is approximated as a set of harmonic oscillators based on normal mode analysis, which provides a computationally efficient yet physically meaningful framework. However, this model has certain limitations. First, real potential functions are not fully harmonic, including higher-order terms lead to variations in oscillation frequencies. Second, presence of mode-mode couplings results in heat dissipation between modes, which is not captured by models that treat vibrational modes as independent oscillators. Although more realistic approaches, such as molecular dynamics simulations can account for these effects, they are computationally demanding and typically do not offer an explicit functional form of the PES. This lack of analytical expression limits their utility in developing deeper theoretical frameworks. To address these challenges, we propose using reservoir computing as a data-driven approach to process atomistic simulation outputs, enabling the modeling of anharmonic molecular motion with minimal efforts.

No.12

Automated Searches for Inverse Design in Self-Assembly of Branched Block Copolymers

Daeseong Yong

KIAS

Branched block copolymers have attracted growing interest due to their ability to stabilize complex morphologies that are often difficult to achieve with linear polymers. To accelerate the discovery and design of such morphologies, we have developed an algorithmic approach that automatically eliminates redundant computations in self-consistent field theory (SCFT) software. In addition, by identifying and removing duplicate architectures based on SMILES notation, our method also removes topological redundancies, significantly reducing the design space. Using this approach, we investigated the self-assembly behavior of AB-type branched block copolymers across various candidate phases, including gyroid, diamond, primitive, and spherical structures, and found that most network phases can be stabilized. This approach is currently being extended to more complex branched polymers using graph convolutional

No.13

Hunting and identifying coloured resonances in four top events with machine learning

Thomas Flacke

KIAS

Machine learning provides a powerful tool for identifying and distinguishing new physics signals at the Large Hadron Collider (LHC). at particle colliders. In this talk, I discuss the discovery and exclusion potential at the LHC for pair- and single produced color sextet and octet scalars states in the 4 top final state. We implement a convolutional neural network combined with a fully connected DNN to separate signal from Standard Model background events.

For LHC operated at 14 TeV and a luminosity of 3 ab^{-1} we find an expected discovery reach of $m_8=1800$~GeV and $m_6=1920$~GeV for pair produced color octets and sextets and an expected exclusion reach of $m_8=2020$~GeV and $m_6=2140$~GeV. We also show that retrained networks allow for very clear distinction between color sextet and octet signals if an excess is found.

No.14

Expressive Power and Universal Approximation Property of Quantized Neural Networks under Fixed-Point Arithmetic

Yeachan Park

Department of Mathematics and Statistics, Sejong University

Traditional studies on the expressive power of neural networks typically assume real-valued parameters and exact arithmetic, neglecting the effects of quantization and rounding. In this work, we investigate the universal approximation property of quantized neural networks operating under discrete fixed-point parameters and fixed-point arithmetic, which may introduce rounding errors.

We begin by establishing both necessary and sufficient conditions on the arithmetic system and activation functions to ensure universal approximation. We then demonstrate that many commonly used activation functions, such as Sigmoid, ReLU, ELU, SoftPlus, SiLU, Mish, and GELU, satisfy our sufficient condition. This implies that quantized networks employing these activations retain universal approximation capabilities.

Furthermore, we prove that under a mild assumption on the activation function -specifically, the existence of a fixed-point input mapped to zero - our necessary and sufficient conditions coincide. This yields a complete characterization of the universal approximation for a broad class of activation functions.

Finally, we extend our results to show that even quantized networks constrained to binary weights in$\{-1,1\}$ can achieve a universal approximation when paired with practical activation functions.

No.15

Teaching Neural Operators the Local Physics of Hyperbolic Conservation Laws

Taeyoung Kim

KIAS

Recently, a variety of neural operators—such as the Fourier Neural Operator (FNO) and the Graph Neural Operator (GNO)—have been proposed to approximate the solution operators of partial differential equations (PDEs). These models typically learn a global mapping from a fixed distribution of initial conditions to the solution at a single terminal time, which limits their ability to generalize to out-of-distribution (OOD) samples and to perform inference at arbitrary time instances.

To address these shortcomings, we introduce the Flux Neural Operator (Flux NO), which approximates inter-cell fluxes within a finite-volume framework. By explicitly capturing local physics, Flux NO markedly improves accuracy on OOD scenarios—such as shock-dominated regimes—and enables stable long-time predictions.

No.16

Statistical Physics-Inspired Training to Mitigate Latent Gradient Bias of SGD

Yeongwoo Song

Department of Physics, KAIST

Giving up and starting over may seem wasteful in many situations such as searching for a target or training deep neural networks (DNNs). Our study, though, reveals that stochastically resetting to a checkpoint can significantly enhance generalization of DNNs when trained with noisy labels. In the presence of noisy labels, DNNs tend to first capture the general patterns of the data but eventually begin to memorize the corrupted labeled data, leading to overfitting. By deconstructing the dynamics of stochastic gradient descent (SGD), we identify the behavior of a latent gradient bias whichnduced by noisy labels, whichh harms generalization. To counteract this negative effect, we introduce stochastic resetting method to SGD, inspired by recent developments in the field of statistical physics achieving efficient target searches. We first theoretically analyze the conditions where resetting becomes beneficial, and then we empirically validate our theory, confirming the significant improvements achieved by resetting in diverse training scenarios. Our method is both straightforward to implement and readily integrates with other techniques for learning under labels noise. Taken together, this work offers insights into the learning dynamics of DNNs from an interpretability perspective, expanding the potential to analyze training methods through the lens of statistical physics.

No.17

Neighborhood-aware training of Neural ODEs to accurately predict dynamical invariants of chaotic systems

Joon-Hyuk Ko

KIAS

For training surrogate dynamics models on chaotic time series data, it is important that the trained models are able to reproduce the dynamical invariants of the ground truth system. However, training models using a simple trajectory-based loss can lead to unstable training and inaccurate invariant predictions, raising the need for alternative training methods.

In this talk, I will first introduce neural ordinary differential equations (Neural ODEs), a neural network architecture that is suited for modeling continuous dynamical systems due to its differential equation structure. Afterwards, I will present neighborhood-aware training for these models, a method that trains Neural ODEs using the time evolution of neighboring points in state space. After discussing the rationale and the preliminary results obtained with this new method, I will end the talk with future directions to explore.

No.18

Towards trustworthy generative models

Jinseong Park

KIAS

Modern deep generative models, including diffusion models, show a tremendous development for generation tasks. However, at the same time, generative AI models suffer from trustworthy issues, including safety, fairness, or privacy. Thus, in this talk, we will address a bias issue in sampling, and introduces fairness-aware sampling method, called attribute switching, which improves fairness in diffusion models without retraining or external classifiers. In the end, we will discuss about the ongoing topics concerning safety issues with harmful data generation. We conclude with a discussion on promising future directions in the field.