Amica Conference:
Thermodynamic Approaches to Artificial Intelligence

March 9–11 2026, Karuizawa, Japan

March 9

Time	Speaker / Session	Title
	Chair: Hamazaki	
14:00-14:05	Makoto Gonokami	Opening remark
14:05-14:15	Introduction of Daikin Industries, Ltd.	
14:15-14:30	*Self-introduction by participants*	
14:30-14:45	*Break*	
	Chair: Hamazaki	
14:45-15:15	Takahiro Sagawa	Thermodynamics Is All You Need (Probably)
15:15-15:45	Gavin Crooks	Unconventional Computing at the Edge of Thermodynamics
15:45-16:15	Massimiliano Esposito	Computing at kT: Stochastic Thermodynamics from CMOS to Machine Learning Hardware
16:15-16:45	*Break / Free Discussion*	
	Chair: Ito	
16:45-17:15	Daisuke Okanohara	Learning as a Finite-Time Non-Equilibrium Process
17:15-17:45	Max Welling	Generative AI and Stochastic Thermodynamics; a Tale of Free Energies
17:45-18:15	Taiji Suzuki	Statistical learning theory of diffusion models: Optimality and high dimensionality
18:15-19:00	*Break / Free Discussion*	
19:00-	*Dinner*	

March 10

Time	Speaker / Session	Title
	Chair: Sagawa	
9:00-9:30	Christopher Jarzynski	The virtues of virtual free energy estimation
9:30-10:00	David Wolpert	Strengthened second law for all digital computers
10:00-10:30	Hirosi Ooguri	Learning the Shape of Information
10:30-11:00	*Break / Free discussion*	
	Chair / Moderator: Sugiyama	
11:00-11:30	Jascha Sohl-Dickstein	Advice for a (young) investigator in the first and last days of the Anthropocene
11:30-12:00	Discussion	*How is generative AI changing society?*

Time	Speaker / Session	Title
12:30-13:30	*Lunch*	
13:30-14:00	*Break / Free discussion*	
	Chair: Kawaguchi	
14:00-14:30	Lenka Zdeborova	Statistical Physics of Learning in the Age of Attention
14:30-15:00	SueYeon Chung	Computing with Neural Manifolds: A Multi-Scale Framework for Understanding Biological and Artificial Neural Networks
15:00-15:30	*Break / Free discussion*	
	Chair / Moderator: Kabashima	
15:30-16:00	Discussion	*Thermodynamics for AI?*
16:00-16:30	*Break / Free discussion*	
	Chair: Ashida	
16:30-17:00	Marco Cuturi	Semidiscrete flow matching and other advances in optimal transport
17:00-17:15	Asuka Takatsu	Law of large numbers for dependent but identically distributed random variables
17:15-17:30	Sosuke Ito	Thermodynamic bounds for diffusion models based on optimal transport
17:30-18:30	*Break / Free discussion*	
18:30-	*Dinner*	

March 11

Time	Speaker / Session	Title
	Chair: Takatsu	
9:00-9:30	Masashi Sugiyama	Recent Advances in Reward Modeling for Reinforcement Learning
9:30-10:00	Yoshiyuki Kabashima	Compressed sensing based on diffusion models
10:00-10:45	*Break / Free discussion*	
	Chair: Kabashima	
10:45-11:00	Kyogo Kawaguchi	Discovering emergent structure in biological and artificial systems with machine learning
11:00-11:15	Yuto Ashida	Generative diffusion model with inverse renormalization group flows
11:15-11:30	Ryusuke Hamazaki	Quasiprobability thermodynamic uncertainty relation
11:30-11:35	Closing remark	
11:35-12:15	*Break*	
12:15-	*Lunch & Farewell*	

Abstract Book

March 9

Takahiro Sagawa (The University of Tokyo)

Title: Thermodynamics Is All You Need (Probably)

Abstract: Recent advances in generative AI have renewed interest in the links to thermodynamics and statistical physics; for example, diffusion models have been discussed in connection with stochastic thermodynamics, and the attention mechanism has been related to modern Hopfield networks. In this talk, I will provide a brief introduction to stochastic thermodynamics, including the Jarzynski equality and the Crooks fluctuation theorem, and discuss its extension to thermodynamics of information. I will emphasize that thermodynamics of information clarifies fundamental lower bounds on the energetic cost required for information processing, which would provide a quantitative perspective on the rapidly increasing power consumption for modern AI. I will conclude with our recent experimental results on realization of optimal transport using colloidal particles manipulated by optical tweezers, where we designed protocols to minimize entropy production for finite-time information erasure (https://www.nature.com/articles/s41467-025-66519-9).

Gavin Crooks (Normal Computing)

Title: Unconventional Computing at the Edge of Thermodynamics

Abstract: As AI pushes against energy, latency, and scaling limits, “more compute” is becoming an increasingly physical proposition. In this talk I will discuss the thermodynamic costs of conventional digital computation, and explore emerging unconventional computing paradigms that can circumvent these limitations, including stochastic and probabilistic computing architectures that exploit, rather than suppress, thermal fluctuations.

Massimiliano Esposito (University of Luxembourg)

Title: Computing at kT: Stochastic Thermodynamics from CMOS to Machine Learning Hardware

Abstract: As electronic devices scale toward energies comparable to kT, thermal fluctuations are no longer a perturbation but a defining feature of computation. In this regime, computation must be understood as a physical, stochastic process, and traditional deterministic abstractions of electronic circuits become insufficient. In this talk, I will show how stochastic thermodynamics provides a unified framework to describe nonlinear electronic circuits operating close to kT, with a focus on subthreshold CMOS technologies. By enforcing thermodynamic consistency at the level of stochastic dynamics, this approach captures noise, dissipation, and energy transduction on equal footing, and clarifies their role in computation. I will then discuss how this perspective naturally leads to probabilistic computing primitives such as probabilistic bits (p-bits), which exploit intrinsic device fluctuations rather than suppressing them, and briefly mention their experimental realization. Finally, I will argue that stochastic thermodynamics can serve as a selection principle for emerging AI hardware. By linking sampling, dissipation, and physical constraints, it offers guidance for identifying computational architectures that are naturally compatible with noisy, energy-efficient devices.

Daisuke Okanohara (Preferred Networks)

Title: Learning as a Finite-Time Non-Equilibrium Process

Abstract: Modern generative AI has achieved remarkable empirical success, yet we lack a principled understanding of what learning fundamentally is. Training is a finite-time, stochastic, and dissipative process far from equilibrium, and central phenomena—such as catastrophic forgetting and generalization—remain theoretically opaque. I argue that learning should be viewed as a non-equilibrium thermodynamic process in parameter space, where irreversibility constrains reachable configurations and limits future adaptability. I conclude by outlining two theoretical frameworks that formalize this perspective.

Max Welling (University of Amsterdam / CuspAI)

Title: Generative AI and Stochastic Thermodynamics; a Tale of Free Energies

Abstract: We discuss how the methods deployed in the field of Generative AI can be described by the mathematics used to describe non-equilibrium thermodynamic systems. We discuss notions such as work, heat, entropy in the context of learning systems. We show how ideas from AI can help estimate free energies for molecular systems and we show how generative models such as Normalizing Flows, Variational Autoencoders, Stochastic Normalizing Flows and Diffusion Models can all be understood as controlled thermodynamic systems that minimize free energies under external forces. We hope that shedding light on the deep relationship between these separate fields will help their cross fertilization.

Taiji Suzuki (The University of Tokyo)

Title: Statistical learning theory of diffusion models: Optimality and high dimensionality

Abstract: In this talk, I will talk about theories of diffusion models and discuss its optimality from statistical perspective. In particular, it will be discussed how the diffusion models avoid the curse of dimensionality adaptively capturing the data structure of high dimensional and infinite dimensional data. First, I introduce some relevant theories of diffusion models and show its statistical optimality. Actually, it is shown that the diffusion models can achieve the minimax-optimality under the setting where the target density has some regularity conditions. Second, I will discuss its adaptivity to high dimensional and even infinite dimensional data. Usually, generative models suffer from the curse of dimensionality. However, it will be shown that diffusion models can mitigate the issue utilizing structures of the data distribution such as (i) manifold assumption and (ii) anisotropic smoothness of the score. We show that in either setting, diffusion models can achieve a near optimal rate.

March 10

Christopher Jarzynski (University of Maryland)

Title: The virtues of virtual free energy estimation

Abstract: The convergence of numerical free energy estimation methods can be accelerated using artificial fields that "escort" simulated trajectories along near-equilibrium paths. Unfortunately, designing such fields is not easy. Taking a cue from the mathematics behind diffusion models, we introduce a method based on virtual escorting. This method adopts a post-processing approach: given a fixed set of nonequilibrium trajectories, a virtual escorting field is expressed using a neural network, and an estimate of the free energy difference is constructed from the trajectories. The neural network's parameters are then adjusted to optimize the convergence of the free energy estimate. I will describe the method, and will discuss conditions under which it produces a zero-variance estimator of the free energy difference.

David Wolpert (Santa Fe Institute)

Title: Strengthened second law for all digital computers

Abstract: All digital devices have components that implement Boolean functions, mapping that component's input to its output. Traditionally, circuit complexity theory has focused on the resource costs of a circuit's size (its number of gates) and depth (the longest path length from the circuit's input to its output). I will present recent results that extend circuit complexity theory by investigating the strictly positive lower bounds on the thermodynamic cost of running a given circuit. Going further, I will also show how to calculate strictly positive lower bounds on the thermodynamic costs of running RASP machines, and of running high-level programs. Crucially, all these results hold independent of the microscopic physics; they provide a set of strengthened second laws.

Hirosi Ooguri (Caltech / The University of Tokyo)

Title: Learning the Shape of Information

Abstract: Entanglement entropy is a fundamental quantity connecting quantum information to spacetime geometry in quantum gravity. For general quantum states, classifying all linear inequalities among multipartite entanglement entropies beyond three parties remains open. However, the situation is dramatically different for systems admitting semi-classical gravity duals. In 2015, we introduced an exact combinatorial reduction that makes the classification problem finite. Using this framework, we proved that the holographic entropy cone is rational polyhedral (hence finitely generated) and obtained complete classifications for 2, 3, and 4 parties. The new families we found for 5 parties were later shown to complete the classification. However, the search becomes intractable for more than 5 parties due to a double-exponential explosion in the underlying combinatorics.

In this talk, I will present our recent progress using reinforcement learning (RL) to navigate the search space and identify candidate new inequalities. RL is well suited here because the objective is high-dimensional and non-smooth; since the entropies are computed by min-cuts, gradients are unavailable or unstable. As proof of concept, our policy model, trained to reconstruct weighted graph models from target entropy data, successfully rediscovers the holographic entropy inequality called the monogamy of mutual information. The model also provides evidence for new families of inequalities for six parties. Our work illustrates how RL can act as a discovery engine for precise, verifiable structures in quantum gravity.

Jascha Sohl-Dickstein (Anthropic)

Title: Advice for a (young) investigator in the first and last days of the Anthropocene

Abstract: Within just a few years, it is likely that we will create AI systems that outperform the best humans on all intellectual tasks. This will have implications for your research and career. I will give practical advice, and concrete criteria to consider, when choosing research projects, and making professional decisions, in these last few years before AGI.

Lenka Zdeborova (EPFL)

Title: Statistical Physics of Learning in the Age of Attention

Abstract: Over the past decades, statistical physics has provided a powerful framework for analyzing exactly solvable models of learning in high dimensions, revealing fundamental limits on generalization, phase transitions in performance, and the interplay between data, architecture, and learning algorithms. In this talk, I will present how this perspective has recently been extended beyond classical perceptron-type models toward modern architectures that process sequences of tokens through attention layers, as in transformers.

SueYeon Chung (Harvard University)

Title: Computing with Neural Manifolds: A Multi-Scale Framework for Understanding Biological and Artificial Neural Networks

Abstract: Recent breakthroughs in experimental neuroscience and machine learning have opened new frontiers in understanding the computational principles governing neural circuits and artificial neural networks (ANNs). Both biological and artificial systems exhibit an astonishing degree of orchestrated information processing capabilities across multiple scales - from the microscopic responses of individual neurons to the emergent macroscopic phenomena of cognition and task functions. At the mesoscopic scale, the structures of neuron population activities manifest themselves as neural representations. Neural computation can be viewed as a series of transformations of these representations through various processing stages of the brain. The primary focus of my lab's research is to develop theories of neural representations that describe the principles of neural coding and, importantly, capture the complex structure of real data from both biological and artificial systems.

In this talk, I will present three related approaches that leverage techniques from statistical physics, machine learning, and geometry to study the multi-scale nature of neural computation. First, I will introduce new theories based on statistical physics and convex geometry that connect complex geometric structures that arise from neural responses (i.e., neural manifolds) to the efficiency of neural representations in implementing a task. Second, I will employ these theories to analyze how these representations evolve across scales, shaped by the properties of single neurons, learning dynamics, and the transformations across distinct brain regions. Finally, I will show how these insights extend efficient coding principles beyond early sensory stages, linking representational geometry to efficient task implementations. This framework not only help interpret and compare models of brain data but also offers a principled approach to designing ANN models for higher-level vision. This perspective opens new opportunities for using neuroscience-inspired principles to guide the development of intelligent systems.

Marco Cuturi (Apple ML Research)

Title: Semidiscrete flow matching and other advances in optimal transport

Abstract: I will present works that illustrate recent advances in machine learning that are inspired by computational optimal transport approaches. I will cover how computing couplings between noise and data (either with large batches or with a semidiscrete formulation) facilitates training in flow matching, and how to apply the Brenier theorem for fast retrieval.

Asuka Takatsu (The University of Tokyo / RIKEN)

Title: Law of large numbers for dependent but identically distributed random variables

Abstract: I introduce the notion of dependent, identically distributed random variables with the help of information geometry. Then I demonstrate a kind of law of large numbers for these random variables.

This talk is based on joint work with Hiroshi Matsuzoe (Nagoya Institute of Technology).

Sosuke Ito (The University of Tokyo)

Title: Thermodynamic bounds for diffusion models based on optimal transport

Abstract: The optimal transport problem can be viewed thermodynamically as a problem of minimal dissipation. Mathematically, it can be discussed using thermodynamic inequalities concerning state changes and dissipation. Similarly, analogous inequalities can be used to discuss generation accuracy in generative models, such as diffusion models. We argue that using optimal transport protocols for diffusion learning processes minimizes the upper bound on generation error, resulting in optimal learning.

March 11

Masashi Sugiyama (RIKEN / The University of Tokyo)

Title: Recent Advances in Reward Modeling for Reinforcement Learning

Abstract: Reinforcement learning (RL) enables an agent to learn actions through trial and error, using rewards as guidance, and it has achieved remarkable success across a wide range of fields, including games, robotics, and the post-training of large language models. However, a major challenge remains: how to design and provide appropriate reward signals. In this talk, I will introduce our recent research on reward modeling in RL. First, I will describe how RL can be made more robust and flexible by incorporating transfer learning and weakly supervised learning. I will then present new RL frameworks that extend how rewards are defined and aggregated by addressing situations in which rewards are not available at every time step and by accommodating diverse evaluation criteria that emphasize safety and efficiency. Finally, I will discuss how research on sequential decision-making should evolve in the era of personalized and increasingly diverse AI systems.

Yoshiyuki Kabashima (The University of Tokyo)

Title: Compressed sensing based on diffusion models

Abstract: Compressed Sensing (CS) is a framework that exploits the expected *sparsity* of many natural signals, enabling their recovery from significantly fewer measurements than traditionally required. Over the past two decades, CS has led to numerous successful applications. Nevertheless, it does not generally achieve the theoretically optimal sensing and reconstruction performance. In principle, optimal performance is achieved through Bayesian inference with the correct prior distribution. In practice, however, constructing such a prior has long been considered intractable. This situation has recently changed. Advances in generative modeling now provide a practical way to approximate these complex priors directly from data. In this talk, we present our recent work on employing diffusion-based generative models as priors for enhanced signal recovery.

This is joint work with Xiangming Meng (Zhejiang University).

Kyogo Kawaguchi (The University of Tokyo / RIKEN)

Title: Discovering emergent structure in biological and artificial systems with machine learning

Abstract: Complex biological and artificial systems exhibit rich structure that emerges from many interacting components, yet identifying the principles underlying this structure remains a central challenge. In this talk, I will present examples where machine learning combined with ideas from statistical physics helps uncover hidden organization in both neural networks and biological systems. First, using transformers trained on a controlled formal language, we study the phenomenon of emergent capabilities in neural networks and show that sudden improvements in performance coincide with the learning of underlying grammatical structure in the data. The dynamics of this transition can be understood through a percolation-based phase transition model of representation formation. Second, I will present an example of an experimental and computational framework for interpreting large ribosome profiling data using RNA language models. This approach reveals sequence-level regulatory grammar and identifies RNA-binding proteins that govern compartment-specific translation across cellular regions. Together, these results illustrate how machine learning can help uncover organizing principles such as phase transitions and bio-sequence grammar that shape complex biological and computational systems.

Yuto Ashida (The University of Tokyo)

Title: Generative diffusion model with inverse renormalization group flows

Abstract: In recent years, the rapid advancement of generative AI, particularly diffusion models, has led to the adoption of new data-driven methods in widely diverse fields—from high-quality image generation to drug discovery and material design. In this talk, I will discuss their theoretical similarities to the theory of renormalization group (RG) widely used in statistical and high-energy physics. Specifically, I will introduce an approach based on RG concepts that aim to capture the hierarchical structure of data more efficiently.

Based on K. Masuki and YA, arXiv:2501.09064

Ryusuke Hamazaki (RIKEN)

Title: Quasiprobability thermodynamic uncertainty relation

Abstract: I present a new quantum extension of thermodynamic uncertainty relations (TURs), which constrain dissipation, currents, and dynamical fluctuations and provide stronger constraints than the second law of thermodynamics. In classical TUR, dynamical fluctuations are naturally characterized by the joint probability distribution of outcomes in a two-point measurement scheme. In quantum systems, however, the same scheme is problematic because it destroys initial coherence via measurement backaction. To overcome this issue, we employ the Terletsky–Margenau–Hill quasiprobability and derive a TUR that is free from the backaction problem. As an application of the quasiprobability TUR, we discuss a new criterion for dissipationless heat currents, recently proposed as a demonstration of quantum advantage.

Based on K. Yoshimura and R. Hamazaki, Phys. Rev. Lett., in press. [arXiv: 2508.14354 (2025)]