Amica Conference:
Thermodynamic Approaches to Artificial Intelligence

March 9–11 2026, Karuizawa, Japan

March 9

Time	Speaker / Session	Title
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*	
14:45-15:15	Takahiro Sagawa	TBA
15:15-15:45	Gavin Crooks	TBA
15:45-16:15	Massimiliano Esposito	Computing at kT: Stochastic Thermodynamics from CMOS to Machine Learning Hardware
16:15-16:45	*Break / Free Discussion*	
16:45-17:15	Daisuke Okanohara	TBA
17:15-17:45	Max Welling	Generative AI and Stochastic Thermodynamics; a Tale of Free Energies
17:45-18:15	Taiji Suzuki	TBA
18:15-19:00	*Break / Free Discussion*	
19:00-	*Dinner*	

March 10

Time	Speaker / Session	Title
9:00-9:30	Christopher Jarzynski	TBA
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*	
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?*
12:00-12:30	*Break / Free discussion*	
12:30-13:30	*Lunch*	
13:30-14:00	*Break / Free discussion*	
14:00-14:30	Lenka Zdeborova	Statistical Physics of Learning in the Age of Attention

Time	Speaker / Session	Title
14:30-15:00	SueYeon Chung	TBA
15:00-15:30	*Break / Free discussion*	
15:30-16:00	Discussion	*Thermodynamics for AI?*
16:00-16:30	*Break / Free discussion*	
16:30-17:00	Marco Cuturi	TBA
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
9:00-9:30	Masashi Sugiyama	Recent Advances in Reward Modeling for Reinforcement Learning
9:30-10:00	Yoshiyuki Kabashima	TBA
10:00-10:45	*Break / Free discussion*	
10:45-11:00	Kyogo Kawaguchi	TBA
11:00-11:15	Yuto Ashida	Generative diffusion model with inverse renormalization group flows
11:15-11:30	Ryusuke Hamazaki	Quasiprobability thermodynamic uncertainty relation
11:30-11:45	Makoto Gonokami	Closing remark
11:45-12:15	*Break*	
12:15-	*Lunch & Farewell*	

Abstract Book

March 9

Takahiro Sagawa (The University of Tokyo)

Title: TBA

Abstract: TBA

Gavin Crooks (Normal Computing)

Title: TBA

Abstract: TBA

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: TBA

Abstract: TBA

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: TBA

Abstract: TBA

March 10

Christopher Jarzynski (University of Maryland)

Title: TBA

Abstract: TBA

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 for 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. I will conclude with a brief perspective on where AI methods are starting to influence theoretical physics.

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: TBA

Abstract: TBA

Marco Cuturi (Apple ML Research)

Title: TBA

Abstract: TBA

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: TBA

Abstract: TBA

Kyogo Kawaguchi (The University of Tokyo / RIKEN)

Title: TBA

Abstract: TBA

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, arXiv: 2508.14354 (2025)