



General Interest Seminar

Learning strategies from fundamental physical limits

Susanne Still (University of Hawaii at Manoa)

Host: Nick Barton

Learning reveals itself to the world through actions, which might reflect an agent's underlying goals. We know that when a specific goal is given in a well defined context, then action strategies can be found as solutions to an optimization problem. But goals are largely context-dependent, and objective functions often have to be constructed by hand, and typically do not generalize well. Clearly this approach is deeply problematic if we want to understand learning and behavior: we should not have to make goals increasingly complicated by hand, to engineer complex behavior. But are there any simple, over-arching principles governing agents in the physical world? Perhaps the most basic, and invariant, agent goal" is self-continuation, and a foundation for achieving that goal is maintaining a positive free energy balance. But can agents that are driven by nothing else than pushing thermodynamic limits come up with information processing strategies comparable to those we use for scientific reasoning? Surprisingly, the answer is yes. Recent developments in far-from-equilibrium thermodynamics have allowed us to understand how optimizing for the efficient use of energy leads to predictive inference. To allow for minimal dissipation, an agent must retain predictive information [Still et al. PRL 109, 120604 (2012)]. We can find a more general version of the same principle by modeling an observer as part of an information engine. We ask: how should the observer best represent available observations to maximize the engine's overall thermodynamic bill? We find that dissipation is proportional to the irrelevant information kept by the observer. Thus, pushing to minimize dissipation leads us to two rules" for information processing: (i) retain all relevant, predictive information, and (ii) retain as little as possible beyond that. This insight allows us to derive, from a very simple physical argument, an algorithm that is widely used for lossy compression and machine learning, coined Information Bottleneck method". Curiously, Tishby et al. recently argued that this same encoding strategy might be reflected in deep neural networks. In summary, the path may now be paved for deriving complex learning strategies and behavior straight from simple fundamental physical limits. Because these limits apply also to quantum systems, this physically driven approach may offer an entirely new window into quantum machine learning.

Friday, July 27, 2018 03:30pm - 04:30pm

IST Austria Campus Meeting room 1st floor / Central Bldg. (I01.1OG - Zentralgebäude)



This invitation is valid as a ticket for the IST Shuttle from and to Heiligenstadt Station. Please find a schedule of the IST Shuttle on our webpage: <https://ist.ac.at/en/campus/how-to-get-here/> The IST Shuttle bus is marked IST Shuttle (#142) and has the Institute Logo printed on the side.