



ELLIS talk

Learning general visual representation with large-scale pre-training

Alexander Kolesnikov

Google Brain

Host: Marco Mondelli

In this talk I will demonstrate that transfer of visual representations pre-trained on large-scale data can dramatically improve sample efficiency and simplify hyperparameter tuning. In the first part of the talk I will discuss challenges that arise in large-scale pre-training and how to address them. Then I will dive into strategies for adapting pre-trained models for a target task. Finally, I will present extensive empirical evaluation of large-scale visual models and highlight many surprising findings. In particular, it turns out that huge models pre-trained on large data not only achieve state-of-the-art performance on many standard vision benchmarks, but are also very strong few-shot learners and generalize well in out-of-distribution evaluation scenarios.

Thursday, July 15, 2021 04:00pm - 05:00pm

Zoom Link: Join Zoom Meeting

<https://istaustria.zoom.us/j/91703991543?pwd=dTl3VWlWVTl1MGxxRjBKakN0ajRPZz09> Meeting ID: 917 0399 1543 Passcode: 378221



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