



e l l i s
unit

TURIN
Talk



<http://vandal.polito.it>

www.ellis.eu

Luca Moschella



Relative representations enable zero-shot latent space communication

Neural networks embed the geometric structure of a data manifold lying in a high-dimensional space into latent representations. Ideally, the distribution of the data points in the latent space should depend only on the task, the data, the loss, and other architecture-specific constraints. However, factors such as the random weights initialization, training hyperparameters, or other sources of randomness in the training phase may induce incoherent latent spaces that hinder any form of reuse. Nevertheless, we empirically observe that, under the same data and modeling choices, distinct latent spaces typically differ by an unknown quasi-isometric transformation: that is, in each space, the distances between the encodings do not change. In this talk, I will present a way to adopt pairwise similarities as an alternative data representation, that can be used to enforce the desired invariance without any additional training. Neural architectures can leverage these relative representations to guarantee, in practice, latent isometry invariance, effectively enabling latent space communication: from zero-shot model stitching to latent space comparison between diverse settings and, surprisingly, data modalities.

*December 15th 2022,
starting at 15:00*

Luca is a Ph.D. student at Sapienza University of Rome in the Gladia research group led by Professor Emanuele Rodolà, previously a research intern at NNAISENSE and currently at NVIDIA. His research focuses on geometric deep learning and representation learning, particularly in the interaction between different neural systems.