UC DAVIS
COLLEGE OF ENGINEERING

LARA Seminar Series

Understanding and Solving Dimensional Collapse in Contrastive Self-Supervised Learning

Li Jing
Facebook AI Research

Date: February 23, 2022 Time: 4-5 pm PST
Link: https://ucdavis.zoom.us/j/93750537129?pwd=alpDd0JudW80ckxwamRNOHNObHVXdz09

Abstract:

Self-supervised pretraining has become a standard technique for various domains. For visual representation learning, one of the most promising approaches is the joint-embedding method, where the models are trained to be invariant to human-designed distortions. To prevent trivial solutions, a common method is to push away representations coming from different images, known as contrastive learning. We observe that contrastive learning still suffers from a lesser collapsing problem where the embedding vectors fall into a lower-dimensional subspace, thus leading to worse solutions. This is counter-intuitive as the contrastive loss is supposed to push embedding vectors to cover the entire space. We theoretically study the underlying mechanisms. First, we show that strong augmentation will cause dimensional collapse, as the network does not have enough capacity to distinguish positive pairs and negative pairs. Second, we show that implicit regularization will also cause dimensional collapse, as the gradient descent dynamics encourage networks to find minimal rank solutions. We propose an empirical self-supervised learning method, called Barlow Twins, that solves dimensional collapse based on the redundancy reduction principle. This method encourages all the channels of the embedding vectors to be decorrelated from each other, thus maintaining more useful information. Experiments show that Barlow Twins outperforms contrastive learning methods, and is comparable to other state-of-the-art self-supervised learning approaches. In particular, Barlow Twins achieved state-of-the-art performances on semi-supervised learning benchmarks and held the advantage of being stable against training batch sizes.

References:

1. Li Jing, Pascal Vincent, Yann LeCun, Yuandong Tian. "Understanding Dimensional Collapse in Contrastive Self-supervised Learning." *International Conference on Learning Representations (ICLR)* (2022).

2. Jure Zbontar*, Li Jing*, Ishan Misra, Yann LeCun, Stephane Deny. "Barlow Twins: Self-Supervised Learning via Redundancy Reduction." *International Conference on Machine Learning (ICML)* (2021).

Li Jing is a postdoctoral researcher at Facebook AI Research (FAIR), working with Prof. Yann LeCun. Before joining FAIR, he obtained his PhD in physics from MIT, advised by Prof. Marin Soljacic. Before that, he received a BS in physics and a BA in economics from Peking University. He is a co-founder of Lightelligence, Inc. He has been awarded Forbes China 30 under 30. He has won a gold medal in the 2010 International Physics Olympiad (IPhO). His current research interest is on self-supervised learning, with a focus on computer vision.