

Lecture: Deep Learning and Differential Programming

3.1 Computer Vision

Semi-supervised learning

Goal: use additional unlabeled data for learning.

Example:

Dataset 1:

Industrial application: Crop growth stage
3000 images + labels
16 classes, annotated



For each image:
Label
(e.g. which class among 16)

Dataset 2:

Industrial application: Crop growth stage
100 000 images
No labels, no annotation



No manual labeling
Easy / cheap data acquisition!

Self-supervised learning

Humans can predict the future.

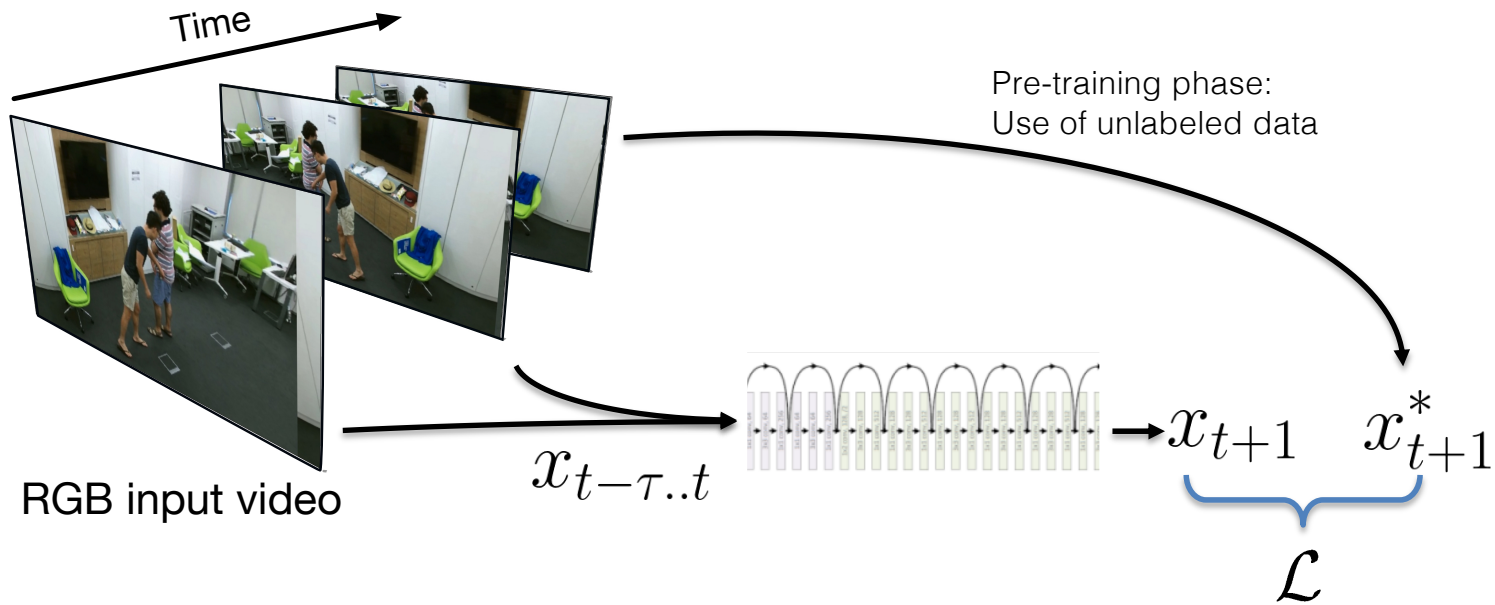
If we learn how to predict the future ... can this help us to train a model which better understands the data?

Self-supervised learning is a variant of semi-supervised learning.




[T. Han, W. Xie, A. Zisserman, Video Representation Learning by Dense Predictive Coding, ICCV 2019]

Self-supervised learning

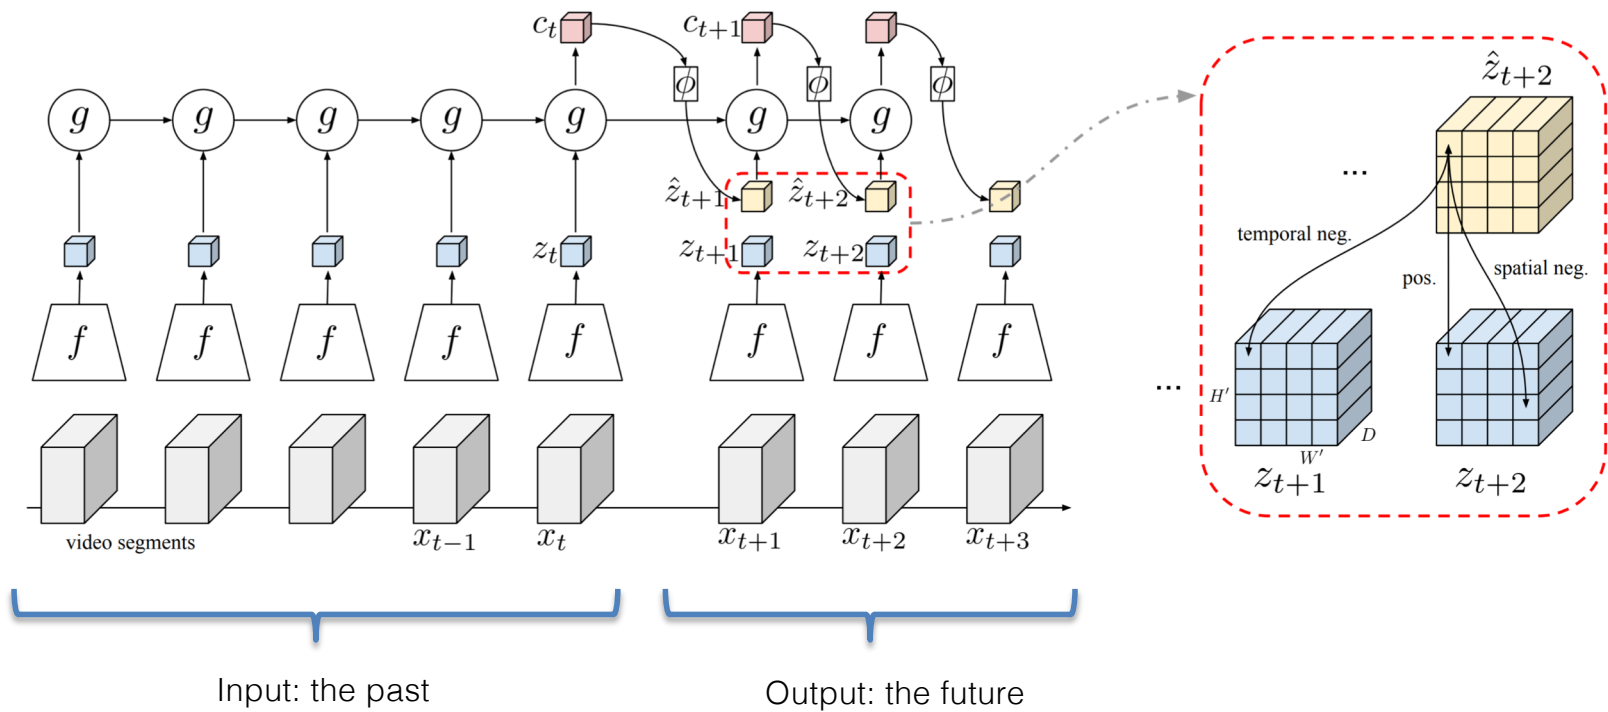


Self-supervised learning



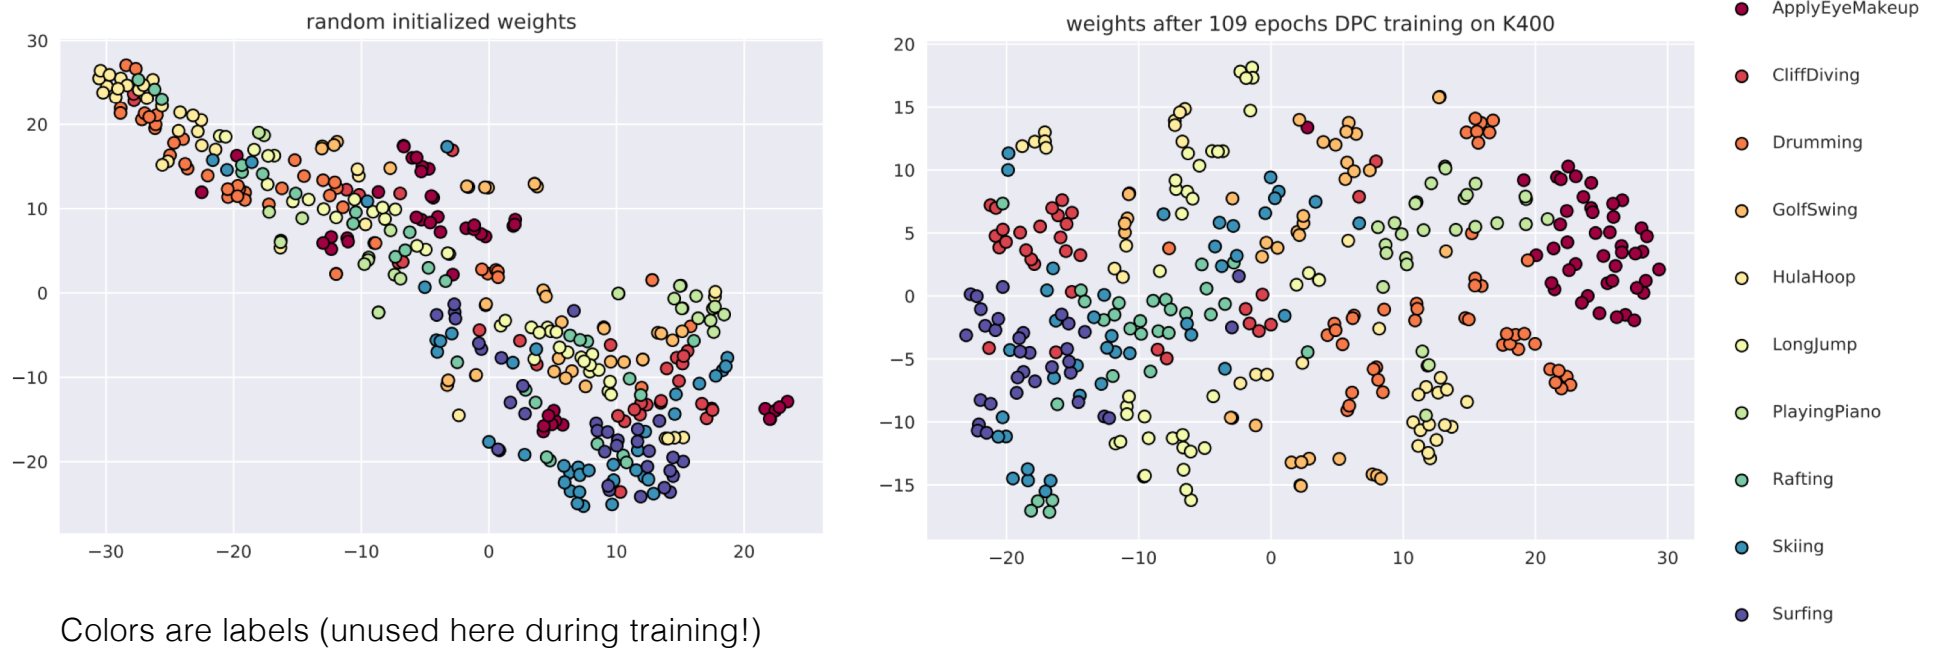
Self-supervised learning

Dense predictions



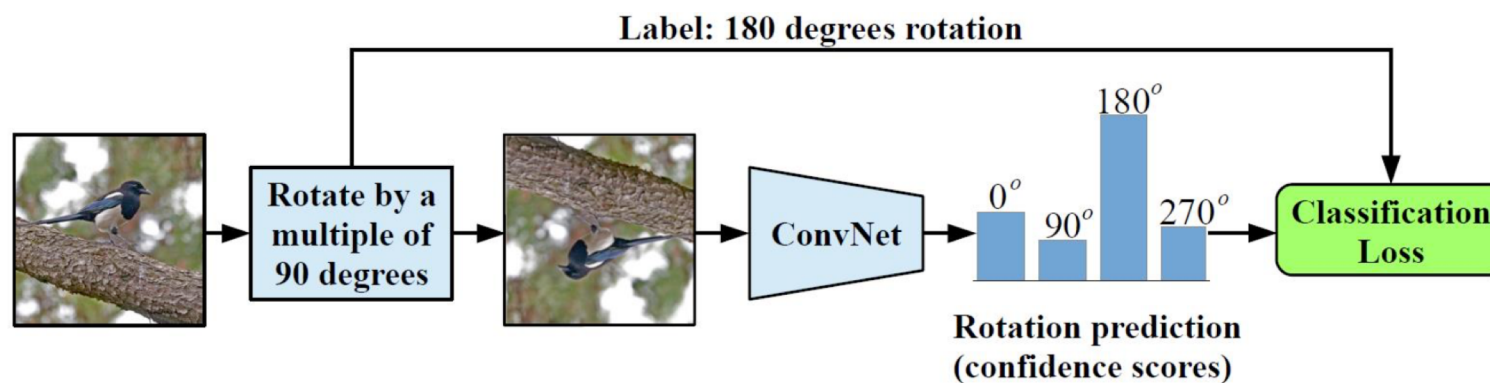
[T. Han, W. Xie, A. Zisserman, Video Representation Learning by Dense Predictive Coding, ICCV 2019]

Self-supervised learning



[T. Han, W. Xie, A. Zisserman, Video Representation Learning by Dense Predictive Coding, ICCV 2019]

How about images?



[S. Gidaris, P. Singh, N. Komodakis, Unsupervised Representation Learning by Predicting Image Rotations, CVPR 2019]

Semi-supervised learning

The industrial perspective, by Varun Nair:

<https://medium.com/@nairvarun18/from-research-to-production-with-deep-semi-supervised-learning-7caaedc39093>

Can boost your performance.

Probably can't solve your problem if performance with labeled data alone is very low.