



UiT The Arctic University of Norway

On the Effects of Self-supervision and Contrastive Alignment in Deep Multi-view Clustering

CVPR 2023

THU-PM-322

Daniel J. Trosten, Sigurd Løkse, Robert Jenssen, Michael C. Kampffmeyer

UiT Machine Learning Group

Department of Physics and Technology, UiT The Arctic University of Norway

Motivation

- Currently: many new approaches are emerging for self-supervision-based components in deep MVC.
 - E.g. contrastive alignment
- However, lack of direction and consistent advancements makes it difficult to conduct further research.

Contributions

- Prove that contrastive alignment is harmful to cluster separability as the number of views increases.
- DeepMVC framework.
 - Open-source implementation!
 - 6 new instances (models).
- Extensive experimental evaluation.

Main findings

- New instances with state-of-the-art performance.
- Contrastive alignment works well with few views, but not with many views
 - In-line with theoretical findings.

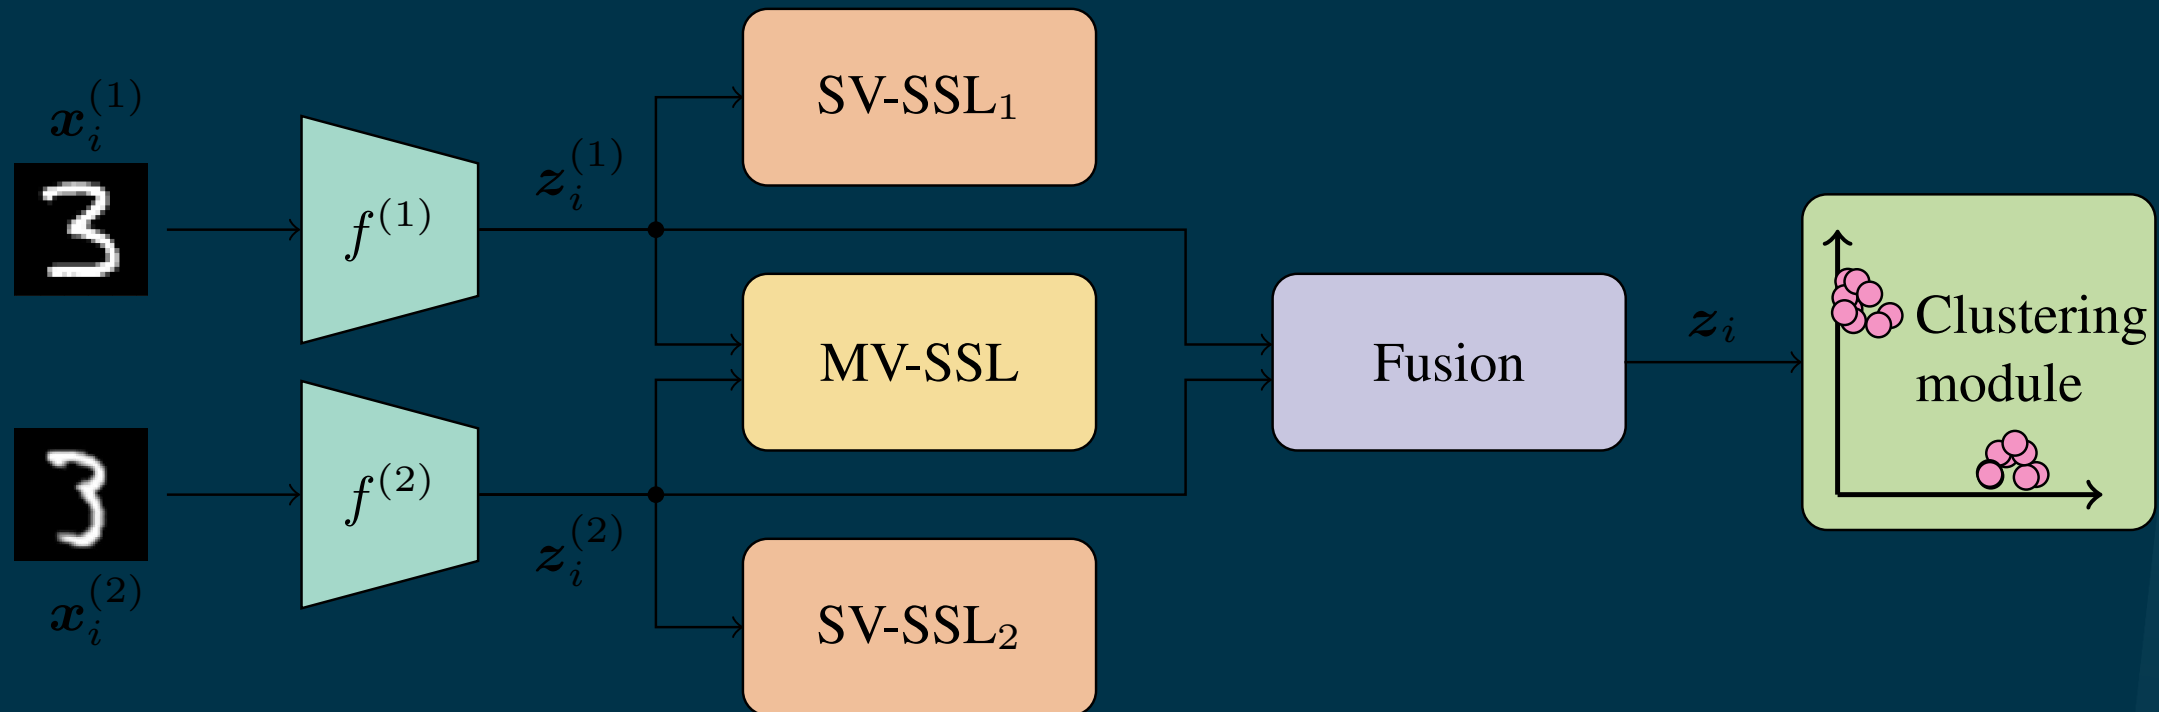
Outline

- ~~Introduction~~
- Drawbacks of alignment
- DeepMVC framework
- New instances
- Experimental findings

Drawbacks of alignment

- *Cluster separability in the representation space depends on cluster separability in the worst view.*
- Encoders can make clusters overlap but cannot separate inseparable clusters.
- Probability of bad views existing increases with the view count.

DeepMVC framework



DeepMVC framework

- Can be extended to arbitrary number of views.
- Includes many recent methods as instances.
- Open-source implementation available on GitHub:



New instances

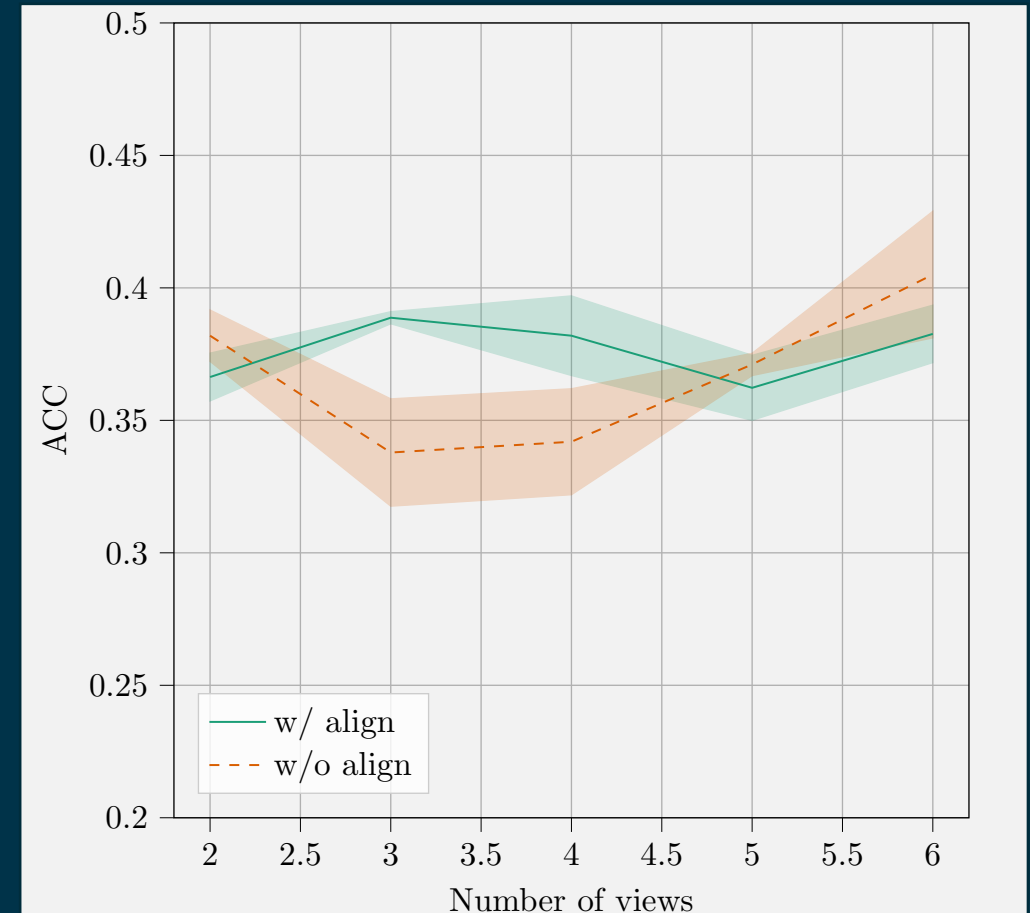
- Objective: Investigate effects of alignment and address drawbacks
- Self-supervision:
 - Contrastive alignment
 - Mutual information maximization
 - Autoencoder
- Clustering modules:
 - Deep divergence-based clustering (DDC)
 - K-means
- Simple baselines!

Experimental findings

- Contrastive alignment works well with few views
- Maximization of mutual information works better with many views
- Dataset properties significantly impact performance.
- New instances out-perform previous methods.
- All methods benefit from some form of self-supervision

Alignment vs. number of views

- Performance with alignment decays after 3 views
- Performance without alignment is best with all views



Visit our poster!

Check out our GitHub repo:

