Stanford University

BOĞAZİÇİ ÜNİVERSİTESİ 1863

Differentiable Feature Selection with

Concrete Autoencoders

Abubakar Abid★
Muhammed Fatih Balin★
James Zou

Poster: Thu Jun 13th 06:30 - 09:00 PM @ Pacific Ballroom #188

Unsupervised Feature Selection (UFS) is Widely Used in Machine Learning

- Identify the **subset** of most informative features in dataset
- **Simplifies** the process of training models
- Especially useful if the data is **difficult** or **expensive** to collect

Unsupervised Feature Selection (UFS) is Used Widely in Applied ML

- Example: the ***L1000 Landmark Genes*** [Lamb et al., 2006]

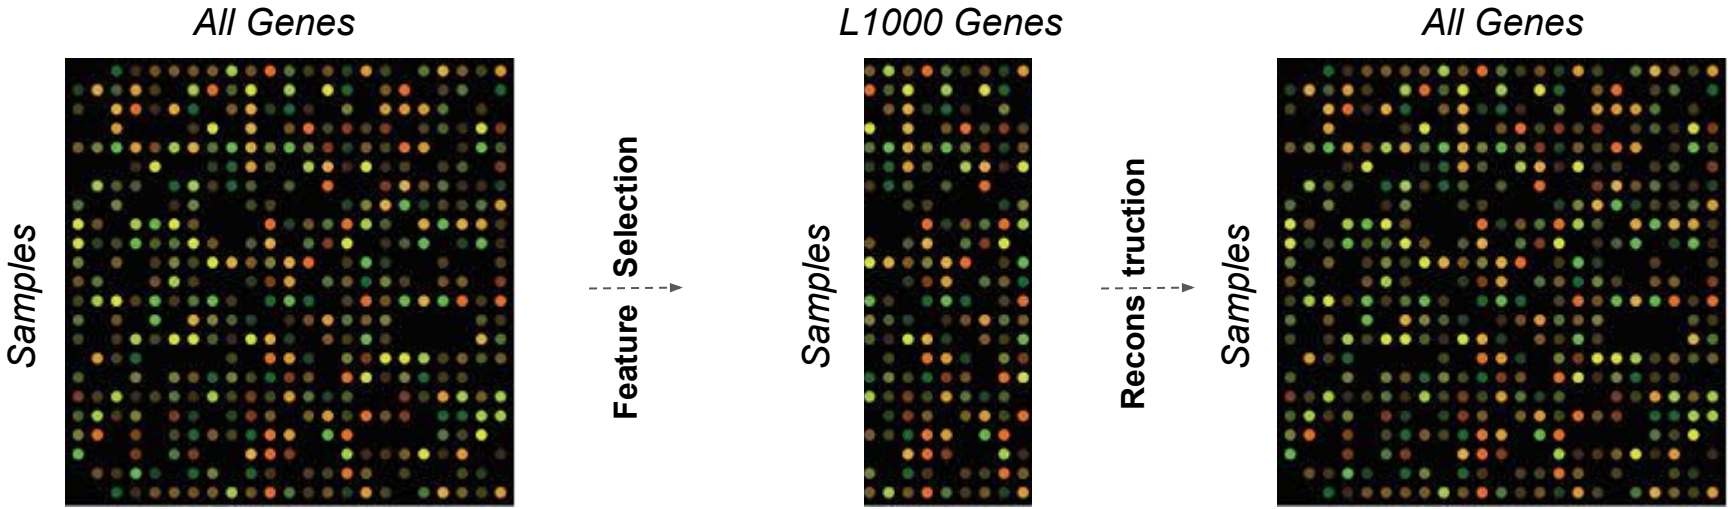

UFS Methods Typically Rely on Regularization

Unsupervised Discriminative Feature Selection (UDFS)
[Yang et al., 2011]

Multi-Cluster Feature Selection (MCFS)
[Cai et al., 2010]

Autoencoder Feature Selection (AEFS)
[Han et al., 2017]

All based on L_1 or L_{21} regularization

What about **directly backpropagating** through discrete "feature selection" nodes?

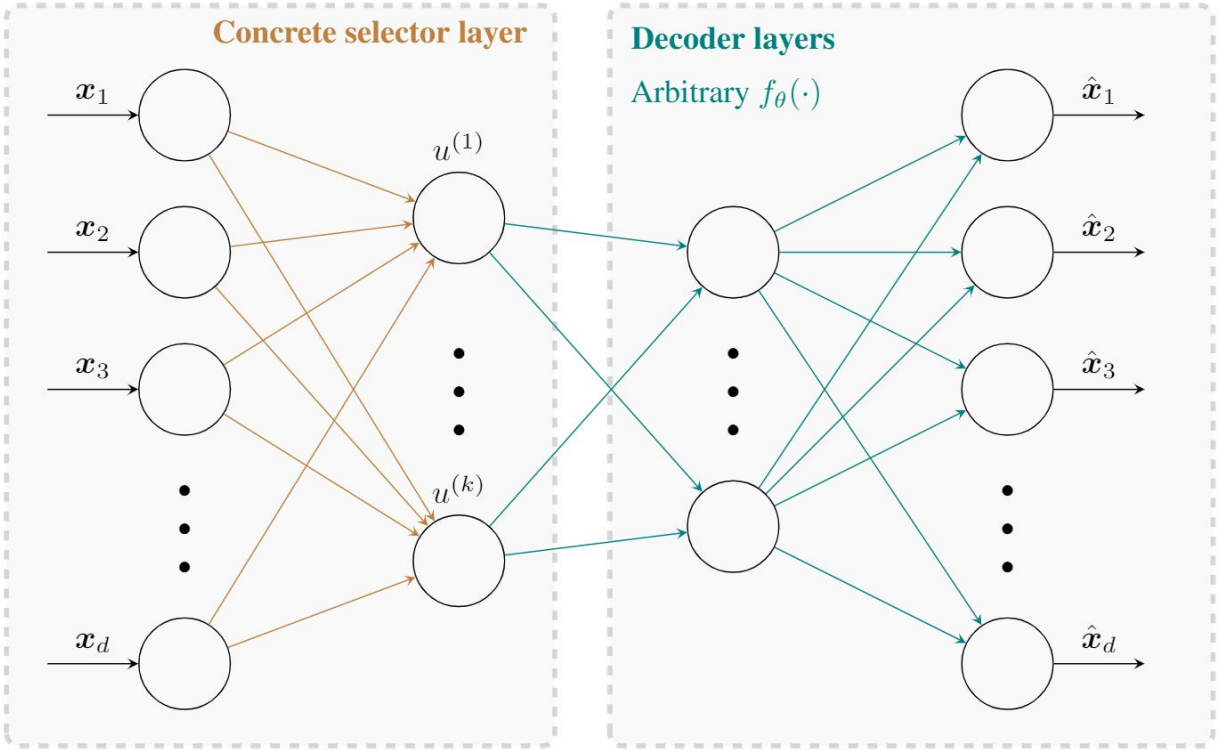

Differentiable Feature Selection and Reconstruction with Concrete Autoencoders

What about **directly backpropagating** through discrete "feature selection" nodes?

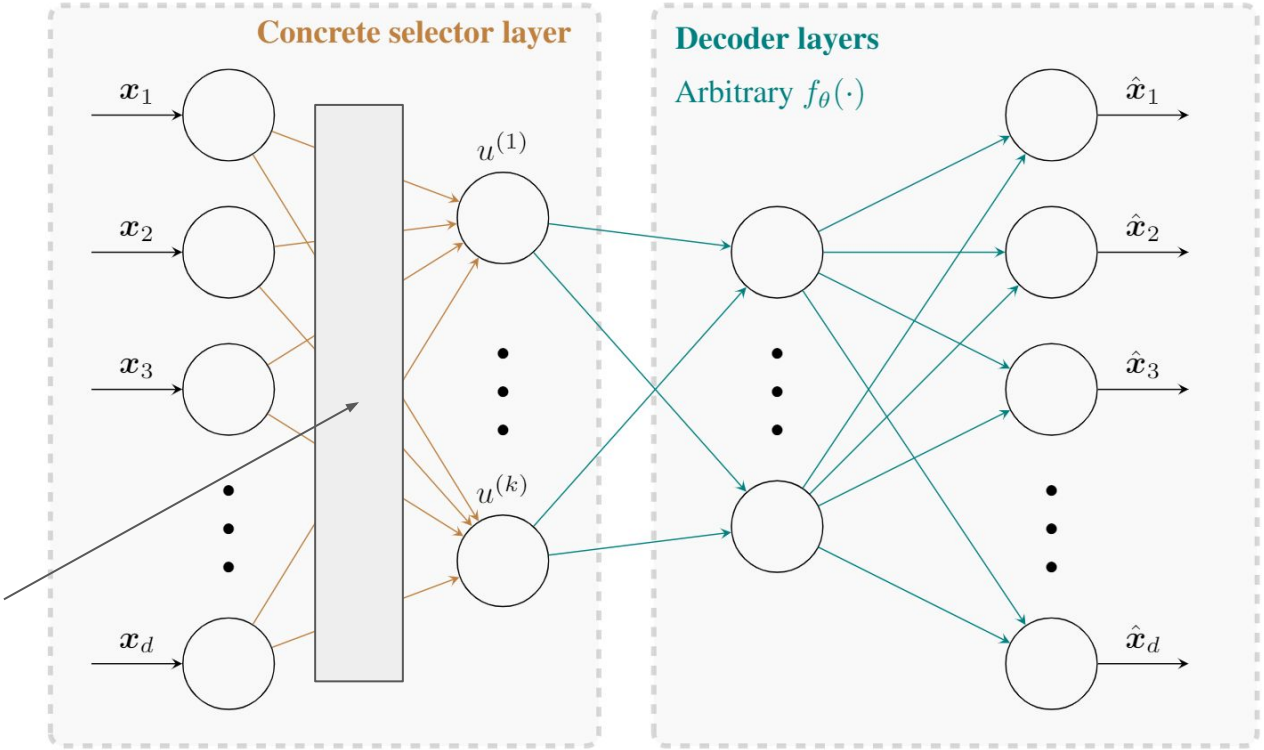

Replace the weights of the encoder with parameters of a **Concrete Random Variable** (Maddison, 2016)

Differentiable Feature Selection and Reconstruction with Concrete Autoencoders

Results on the ISOLET dataset (reconstruction error)

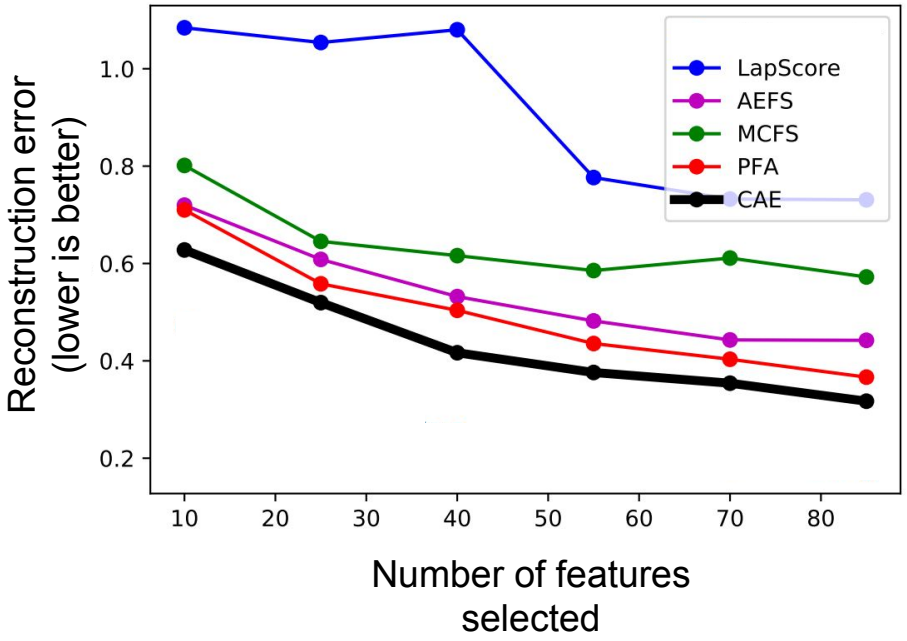

Results on the ISOLET dataset (classification accuracy)

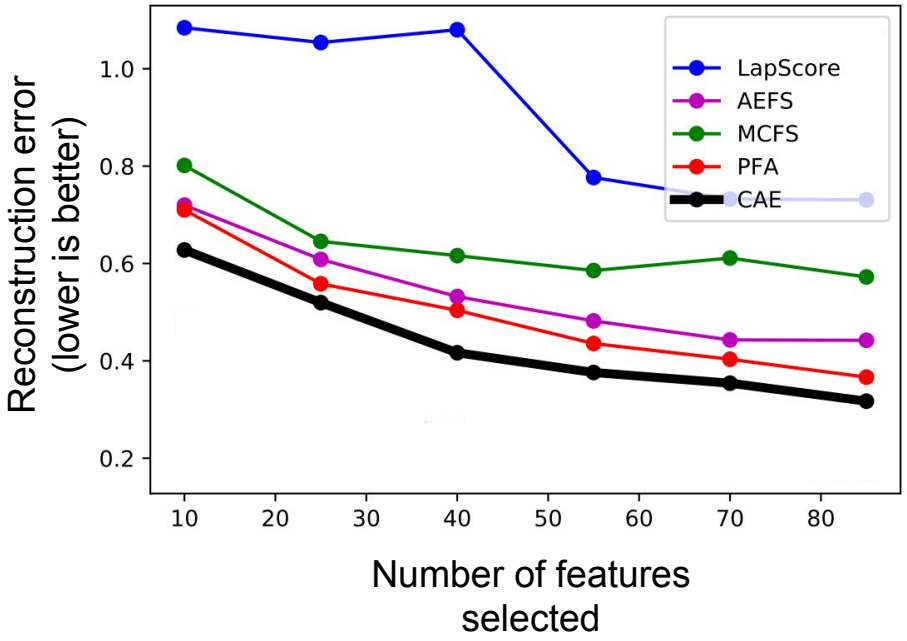

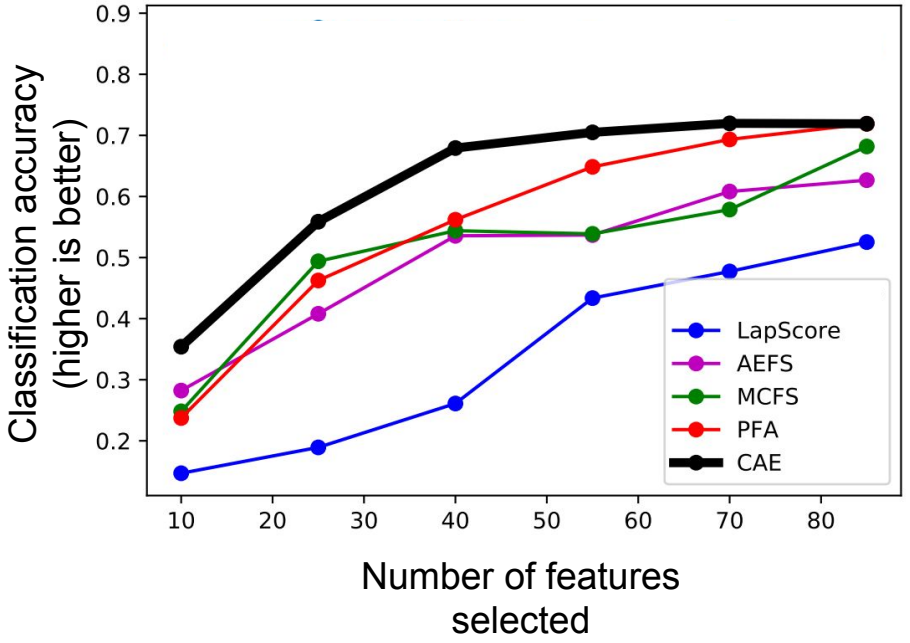

Concrete Autoencoder (CAE) Genes Outperform the L1000 Landmark Genes!

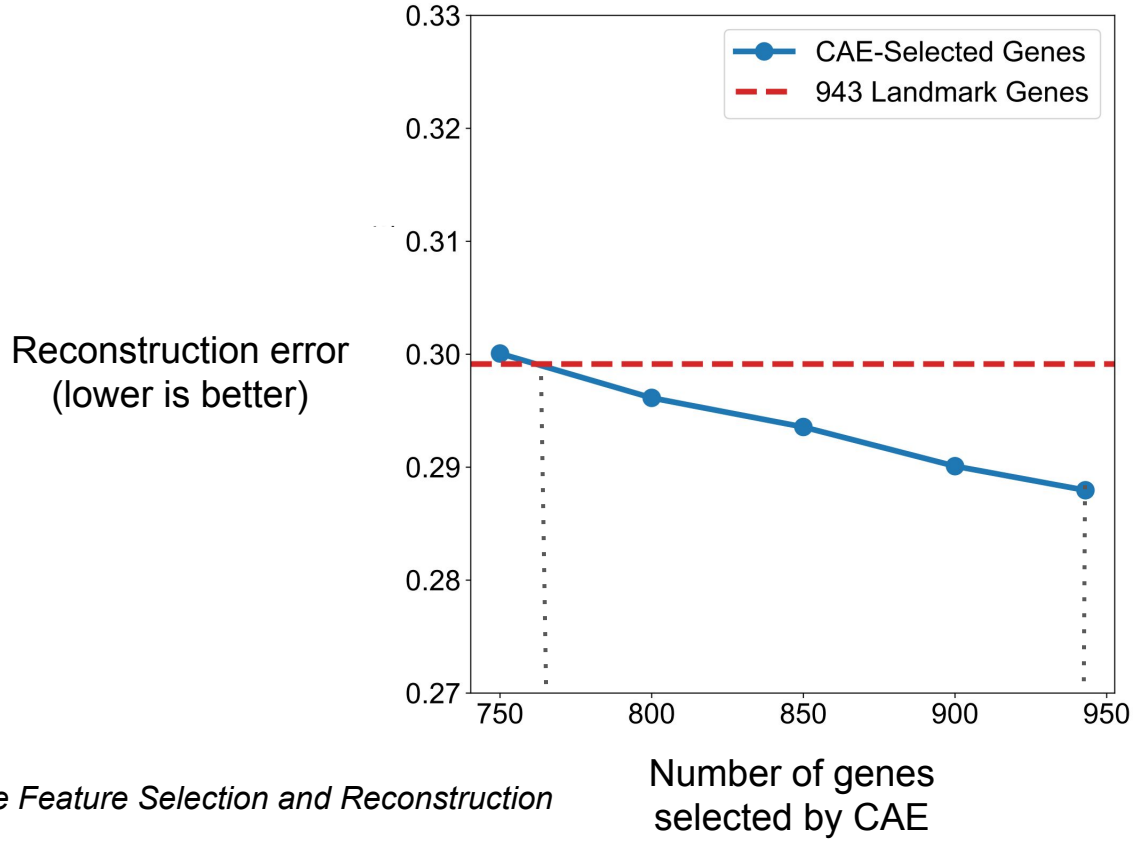

Concrete Autoencoder Takeaways

- More effective than other feature selection methods based on regularization
- Implementation is just a few lines of code from a standard autoencoder
- Training time is similar to standard autoencoder per epoch
- Can be extended to supervised/semi-supervised settings

Start using concrete autoencoders today!

Installation: `pip install concrete-autoencoder`

Code: https://github.com/mfbalin/Concrete-Autoencoders

For more details and results:

Poster: Thu Jun 13th 06:30 - 09:00 PM @ Pacific Ballroom #188

Contact: a12d@stanford.edu, fatih.balin@boun.edu.tr