Deep Adversarial Gaussian Mixture Auto-Encoder for Clustering

Warith HARCHAOUI   Pierre-Alexandre MATTEI
Charles BOUVEYRON

Université Paris Descartes — MAP5
Oscaro.com — Research & Development

February 2017

Clustering

Clustering is grouping **similar** objects together!

Thesis

Representation Learning and Clustering operate a **symbiosis**

Gaussian Mixture Model

- Density Estimation applied to Clustering for **K modes/clusters**
- **Linear complexity** suitable for Large Scale Problems

Learning Representations

- Successful in a **supervised** context (Kernel SVM)
- Successful in an **unsupervised** context (Spectral Clustering)

Auto-Encoder

An auto-encoder is a neural network that consists of:

- an Encoder: $\mathcal{E} : \mathbb{R}^D \rightarrow \mathbb{R}^d$ (compression)
- a Decoder: $\mathcal{D} : \mathbb{R}^d \rightarrow \mathbb{R}^D$ (decompression)

$$D >> d$$

$$\mathcal{D}(\mathcal{E}(x)) \simeq x$$

Optimization Scheme

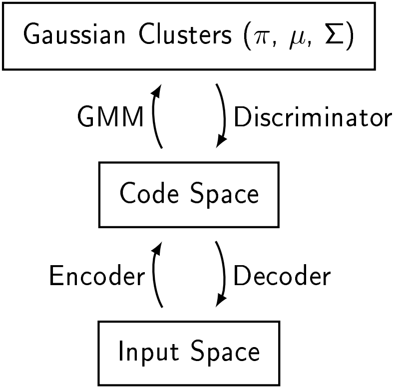

Figure: Global Optimization Scheme for DAC

Adversarial Auto-Encoder

An adversarial auto-encoder is a neural network that consists of:

- an Encoder: $\mathcal{E} : \mathbb{R}^D \to \mathbb{R}^d$ (compression)
- a Decoder: $\mathcal{D} : \mathbb{R}^d \to \mathbb{R}^D$ (decompression)
- a Prior: $\mathcal{P} : \mathbb{R}^d \to \mathbb{R}$ and $\int_{\mathbb{R}^d} \mathcal{P} = 1$ associated with a random generator of distribution $\mathcal{P}$
- a Discriminator: $\mathcal{A} : \mathbb{R}^d \to [0, 1] \subset \mathbb{R}$ that distinguishes fake data from the random generator and real data from the encoder

Optimizations

3 lines objectives:

- The encoder and decoder try to minimize the reconstruction loss
- The discriminator tries to distinguish fake codes (from the random generator associated with the prior) and real codes (from the encoder)
- The encoder also tries to fool the discriminator (opposite discriminator loss function)

Results

Datasets	MNIST-70k	Reuters-10k	HHAR
DAC EC (Ensemble Clustering)	**96.50**	**73.34**	**81.24**
DAC	94.08	72.14	80.5
GMVAE	88.54	-	-
DEC	84.30	72.17	79.86
AE + GMM (full covariances, median accuracy over 10 runs)	82.56	70.12	78.48
GMM	53.73	54.72	60.34
KM	53.47	54.04	59.98

Table: Experimental accuracy results (%, the higher, the better) based on the Hungarian method

Visualizations

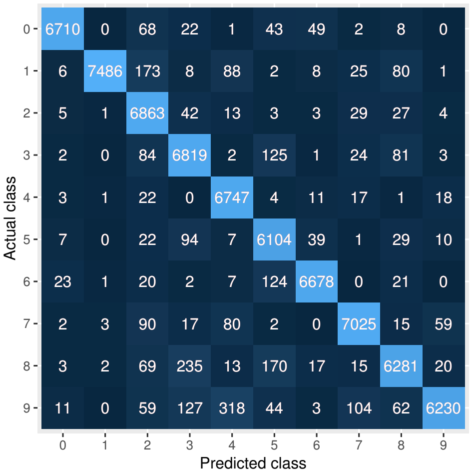

Figure: Confusion matrix for DAC on MNIST. (best seen in color)

Visualizations

μ_k

$\mu_k + 0.5\sigma$

$\mu_k + 1\sigma$

$\mu_k + 1.5\sigma$

$\mu_k + 2\sigma$

$\mu_k + 2.5\sigma$

$\mu_k + 3\sigma$

$\mu_k + 3.5\sigma$

Figure: Generated digits images. From left to right, we have the ten classes found by DAC and ordered thanks to the Hungarian algorithm. From top to bottom, we go further and further in random directions from the centroids (the first row being the decoded centroids).

Visualizations

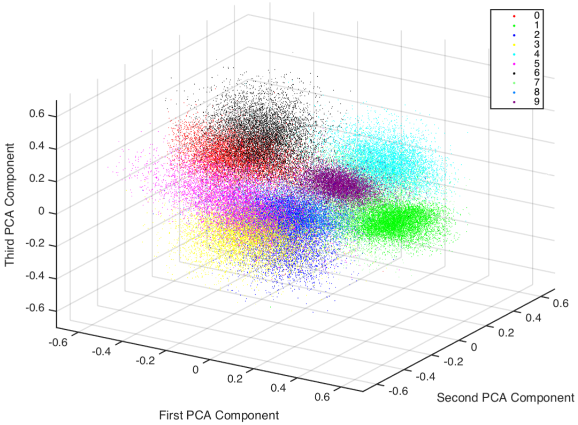

Figure: Principal Component Analysis rendering of the code space for MNIST at the end of the DAC optimization, with colors indicating the true labels. (best seen in color)

Conclusion

Representation Learning and Clustering operate a **symbiosis**

References I

Christopher M.. Bishop.
Pattern recognition and machine learning.
Springer, 2006.

Ian Goodfellow, Yoshua Bengio, and Aaron Courville.
Deep Learning.
MIT Press, 2016.
http://www.deeplearningbook.org.

Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio.
Generative adversarial nets.
In *Advances in Neural Information Processing Systems*, pages 2672–2680, 2014.

References II

Alireza Makhzani, Jonathon Shlens, Navdeep Jaitly, and Ian Goodfellow.
Adversarial autoencoders.
arXiv preprint arXiv:1511.05644, 2015.

Pascal Vincent, Hugo Larochelle, Isabelle Lajoie, Yoshua Bengio, and Pierre-Antoine Manzagol.
Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion.
Journal of Machine Learning Research, 11(Dec):3371–3408, 2010.