Discussion:

Segment-based credit scoring using latent clusters in the variational autoencoder

by Mancisidor et al.

Andreas Joseph[1]

Modelling with Big Data and machine learning

Bank of England, 26 November 2018

[1]Disclaimer: The expressed views are those of the author and not necessarily those of the Bank of England (BoE) or its committees.

Summary

- Very innovative work, rich in new and general-purpose methodology
- Goal: predict customers propensity to fall into financial distress (credit scoring)
- Method: 2-step approach
  1. Clustering of loans into different risk profiles using variational autoencoder (VAE) and hierarchical clustering algorithm
  2. Predict distress for each cluster using neural nets (NN)
- Benchmark: NN without clustering
- Result: Clustered approach (slightly) better.

dataset	*n* obs. (main cluster)	fraction	H-measure	benchmark	H-diff.	*p*-value
Nor. car loans	109 969	59%	0.217	0.208	4%	0.085*
Fin. car loans	75 536	65%	0.184	0.182	1%	0.831
Kaggle	97 434	65%	0.292	0.284	3%	0.041**

Results summary for largest cluster in each dataset.

VAE principle

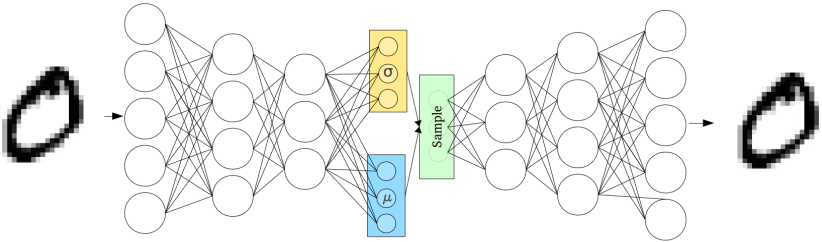

Source: Towards Data Science.

- state-of-the-art generative model: Synthetic output (training data, de-noising, inference, "fake-X", etc.)
- AE-part: Learns low-dim latent representation of data (clustering)
- V-part: Provides probabilistic model
- Trick: Hybrid loss function (reconstruction error & diff. to standard prior)

Comments

- Why V-part? Generative properties not used and different (deep) clusterings could work equally well while being simpler.
- Why clustering+NN seperately (2-steps)? NN (like VAE) learn relevant structures (1-step).
- Scoring differences could be driven by information leakage from WoE.
- Can clustering algorithm be incorporated into AE loss?
- Context & framing: previous studies' findings & logistic benchmark.
- Clarity: VAE part detailed, clustering and step 2 less so.
- More details on CV: VAE dim. $d_z = 2$ always better or only faster to converge initially?
- Look into interpretability frameworks to explain model, e.g. Shapley values / regressions or LIME.

Upshot:

State-of-the-art deep learning application paving the way for future research - THX!