

Disentangling representations of grouped observations in adversarial autoencoders

Felix Pfeiffer

14.11.2018

- Motivation

- Motivation
- Autoencoders


- Motivation
- Autoencoders
- Disentangling Algorithm

- Motivation
- Autoencoders
- Disentangling Algorithm
- Experiments

- Motivation
- Autoencoders
- Disentangling Algorithm
- Experiments
- Conclusions, outlook, discussion

Motivation

First Impression



from Kamil Czarnogórski

int8.io/variational-autoencoder-in-tensorflow

Facial Action Classification

Facial Action Classification

- hard to get sufficient amount of labeled data

Facial Action Classification


- hard to get sufficient amount of labeled data
- finding representations unsupervised

Facial Action Classification

- hard to get sufficient amount of labeled data
- finding representations unsupervised
- profit ?

Encoder

Autoencoder



trained on one face

“Picture of a face”

trained on one face

“Picture of a face”



Autoencoder

trained on one face

“Picture of a face”



Autoencoder



Representation

=

“nearly” FA Code

trained on multiple faces

“Picture of a **X** face”



Autoencoder




Representation

=

“**X**_n**X**_{ea}**X**_r**X**_l**X**_y” F**X**A C**X**o**X**d**X**e

The solution

Input disentangling to representation of:



Autoencoders

Probabilistic Models


Probabilistic Models

- model (e.g. ANN) that acts like a certain p.d.f.


Probabilistic Models

- model (e.g. ANN) that acts like a certain p.d.f.
- naive model wont work


Generative Adversarial Network




Variational Autoencoder



Adversarial Autoencoder




AAE training




Disentangling Algorithm


Grouped observations




Batch evaluation



Batch evaluation disentangled



Hijacking the style



Experiments

Setup


Setup

- Quantitative: classifying CMNIST (4 setups)

Setup


- Quantitative: classifying CMNIST (4 setups)
- Qualitative: manipulating representations, then reconstructing

Q1: full classifier





Q2: classifier on full entangled representation

$$z \sim Q(X)$$




Q3/Q4: classifier on each representation

$$(z_c, z_s) \sim Q(X)$$




Results

Training error



Recognition rate




Reconstruction




Latent cross product

	6	6	8	6	/	/	7	0	6	8	7	2	2	5	0
6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6
6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6
8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8
6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6
/	/	/	/	/	/	/	/	/	/	/	/	/	/	/	/
/	/	/	/	/	/	/	/	/	/	/	/	/	/	/	/
7	7	7	7	7	7	7	7	7	7	7	7	7	7	7	7
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6
8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8
7	7	7	7	7	7	7	7	7	7	7	7	7	7	7	7
2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2
2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2
5	5	5	5	5	5	5	5	5	5	5	5	5	5	5	5
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Random style



Random content



Conclusions, outlook, discussion

Conclusions

Conclusions

- It lives !

Conclusions

- It lives !
- stripping the style improves classification performance

Conclusions

- It lives !
- stripping the style improves classification performance
- style holds (nearly) no information on content

Conclusions

- It lives !
- stripping the style improves classification performance
- style holds (nearly) no information on content
- no dependency on each other when reconstruction with recombination

Conclusions

- It lives !
- stripping the style improves classification performance
- style holds (nearly) no information on content
- no dependency on each other when reconstruction with recombination
- non proportional error, when randomizing style

Outlook, discussion

Outlook, discussion

- More complex dataset: e.g. YTFDB

Outlook, discussion

- More complex dataset: e.g. YTFDB
- better classifier: may outperform Q1 ?

Outlook, discussion

- More complex dataset: e.g. YTFDB
- better classifier: may outperform Q1 ?
- reconstruction imbalance: learning defect or feature complexity difference ?

Outlook, discussion

- More complex dataset: e.g. YTFDB
- better classifier: may outperform Q1 ?
- reconstruction imbalance: learning defect or feature complexity difference ?
- any difference with an explicit variance regularizer on style ?

Outlook, discussion

- More complex dataset: e.g. YTFDB
- better classifier: may outperform Q1 ?
- reconstruction imbalance: learning defect or feature complexity difference ?
- any difference with an explicit variance regularizer on style ?
- comparing to VAE that works on parameters

Outlook, discussion

- More complex dataset: e.g. YTFDB
- better classifier: may outperform Q1 ?
- reconstruction imbalance: learning defect or feature complexity difference ?
- any difference with an explicit variance regularizer on style ?
- comparing to VAE that works on parameters
- β -VAE + this.algorithm: get they along ?

A failed try on YTFDB





Notable Literature

D. Bouchacourt, R. Tomioka, and S. Nowozin:
'Multi-level variational autoencoder: Learning
disentangled representations from grouped
observations', 2017

Thank you for your attention !