Neural Network Disease Classification Based Feature Extraction for Imaging Genetics.

Cédric Beaulac

Simon Fraser University and University of Victoria

June the 3rd 2022

This work has been done in collaboration with colleagues at the University of Victoria: Farouk Nathoo, Michelle Miranda and Leno Rocha, and at Simon Fraser University: Sidi Wu, Jiguo Cao, Erin Gibson and Mirza Faisal Beg.

Feature Extraction for Imaging Genetics

Imaging genetics
Neural Network for feature extractions
A first experiment
Discussion

Imaging genetics

A brief introduction to imaging genetics

- Understand the genetic basis of complex neurological diseases (eg. Alzheimer's Disease (AD)).
- Each SNPs might only contribute in a small way.
- There might be additive and interactive effects
- The diseases are complex and there is a risk we lose valuable information when reducing it to a diagnosis variable.

Imaging genetics: the concept

Imaging genetics involves the use of functional neuroimaging data to study objects carrying genetic variants that may relates to neurological disorders.

- Assumption: the endophenotypes derived from the neuroimages are closer to the diseases studied.
- We want to understand the genotype to phenotype relation.
- This may help to identify important genetic variation.

Imaging genetics: the concept

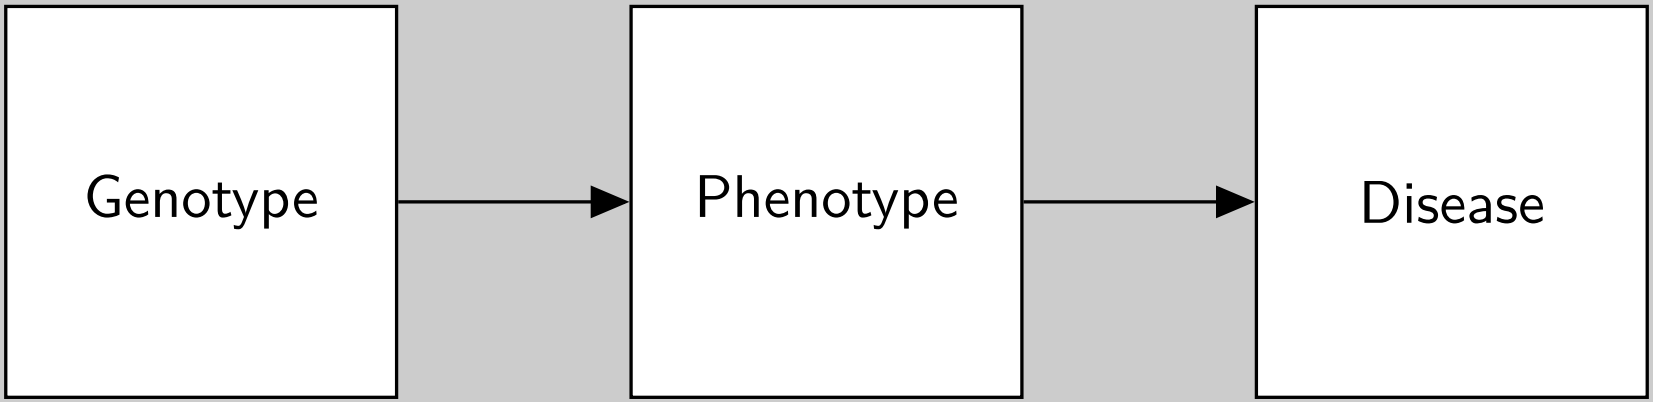

Figure: A graphical representation of the assumed generative process.

Imaging genetics: problem formulation

- $\mathbf{v}_n$ is neuroimage data and $\mathbf{g}_n$ is the genetic data.
- Problem is fitting a parametric model $f_\theta(\mathbf{g}) = \mathbf{v}$.
- and do inference with respect to the parameters θ.
- For example:
  1. f_θ is a linear model, we use confidence interval
  2. f_θ is a Bayesian model we can look at posterior distribution of θ

Imaging genetics: dimension problem

- A major problem of neuroimaging genetics is the large size of both the SNPs data and the image data (Big Data Squared).
- We need f_θ to be an *inference* model not a *blackbox* model.
- For that, we need both the predictors and the response to be of a *reasonable* size.

Neuroimaging genetics: dimension problem

- Historically, both sets of variables are reduced based on expert advices.
- We select a few hundred SNPs suspected to be related to the disease.
- We select brain regions similarly.

Imaging genetics: our goal

- We want to reduce the size of the predictors (SNPs) and the responses (Neuroimages) in a data-driven way.
- In this project we address the latter problem.

Neural Network for feature extraction

Neural Network for feature extraction

- Even though we do not want to use a *blackbox* inference model.
- However, we can separate the feature extraction from the analysis and use a *blackbox* model for feature extraction.
- Neural Networks (NN) are commonly used for dimensionality reduction.
- AutoEncoders are common dimension-reduction models using NNs
- Let us introduce them briefly.

What is an autoencoder ?

- An AutoEncoder (AE) is an unsupervised learning model that learns how to encode (p) and decode (q) data simultaneously (input and output are of the same dimension).
- *Notations* : $\mathbf{x}$ are D-dimensional observations, $\mathbf{z}$ is the M-dimensional code, p is the encoding function for $\mathbf{x}$ ($p : \mathbb{R}^D \rightarrow \mathbb{R}^M$) and q is the decoding function ($q : \mathbb{R}^M \rightarrow \mathbb{R}^D$).
- The code (used interchangeably with latent representation) is usually of lower dimensions, say $M << D$.
- The AE compresses and decompresses high-dimensional data and is commonly used for feature extraction (dimensionality reduction).

Simple autoencoder

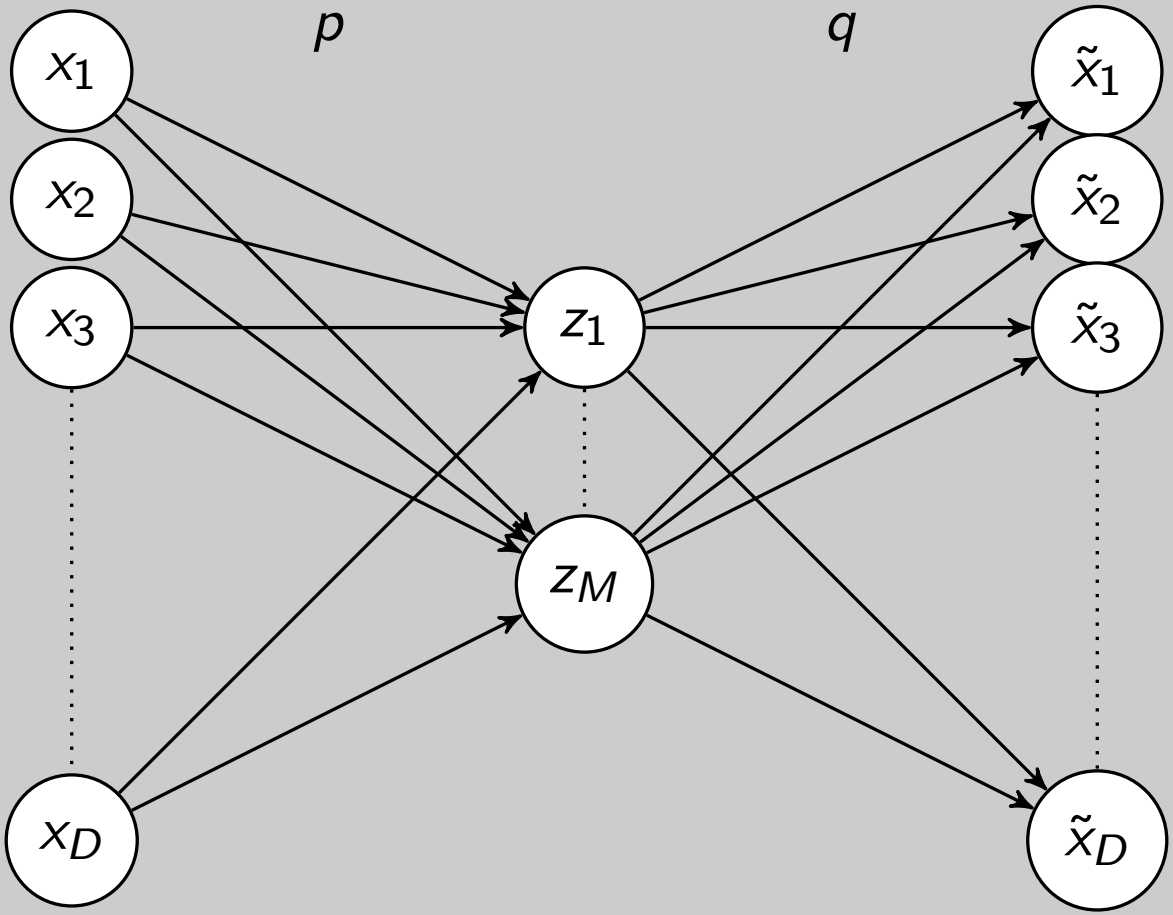

Figure: A simple AE where the code z is a linear combination of the inputs x and where the reconstruction $\tilde{x}$ is a linear combination of the code z.

Autoencoder: a generalized PCA

- There are multiple possible functions p and q and multiple ways to optimize for those.
- *Specific case*: Assume p and q are linear combinations.
- and assume we minimize the quadratic reconstruction error : $\frac{1}{n}\sum_{i=1}^{n} ||\mathbf{x_i} - \tilde{\mathbf{x}}_\mathbf{i}||^2$, where $\tilde{\mathbf{x}} = q(p(x))$.
- Then, the solutions to this problem are the principal components.
- We can think of AEs as generalized PCA (it generalizes the type of functions considered).

Autoencoder with NNs

- We want to allow for more complex p's and q's.
- A modern flexible function comes in mind: a Neural Network (NN).
- Easy to optimize with back-propagation of the gradient (chain rule of derivatives).
- We can still use the quadratic reconstruction MSE to train such a model.

Autoencoder with NNs

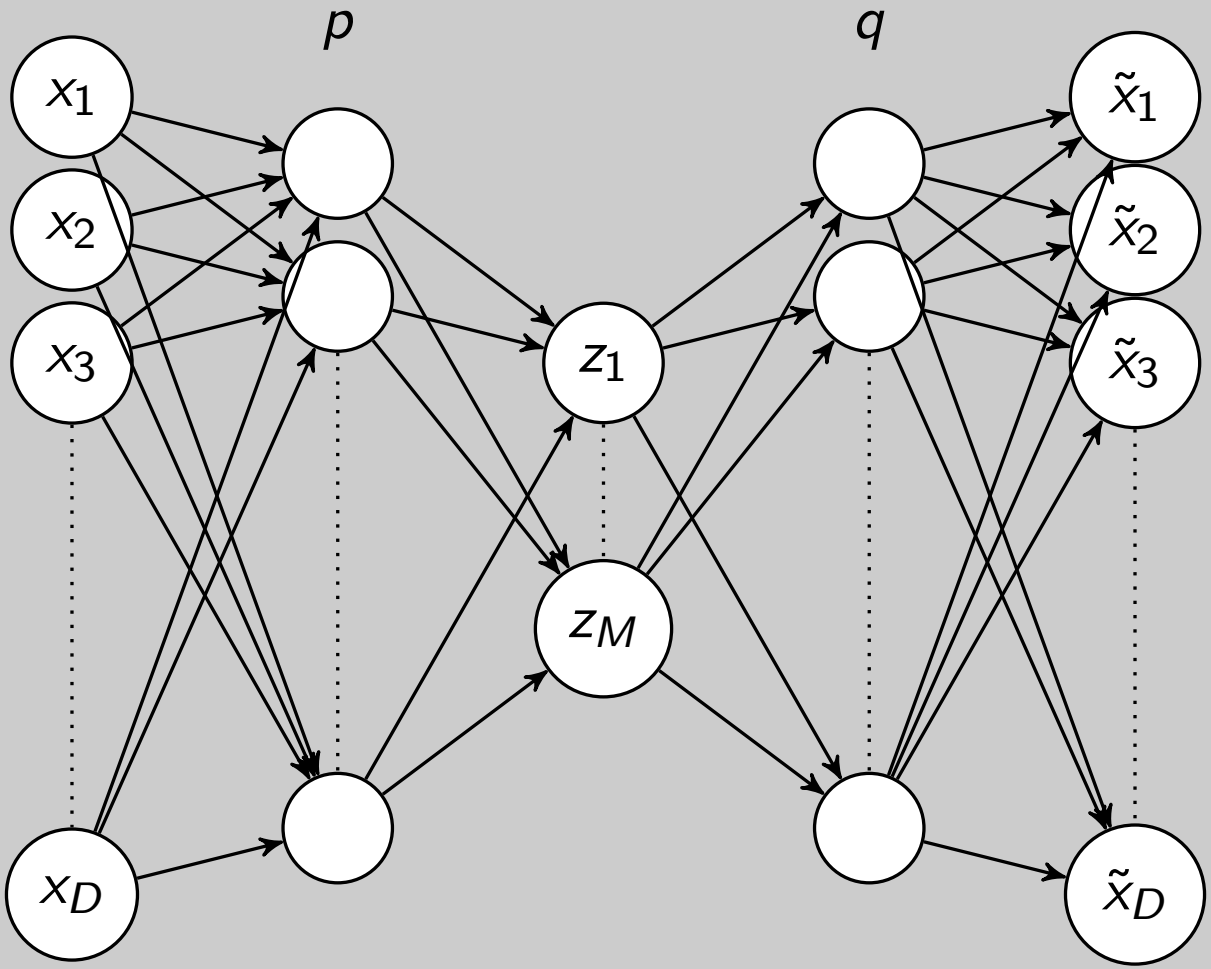

Figure: An AE where both functions p and q and NNs.

Autoencoder: extracted features

- The latent representation z is a lower dimension representation of the observation x.
- It can be perceived as features extracted from the AE (similar to principal components).
- Given the function p we can extract features for new observations.
- Given the function q we can reconstruct an observation given it's latent representation z.

Autoencoder: extracted features

- By design, those features z (and the function p and q) are optimized for reconstruction. (to maximize the variance (PCA)).
- This is certainly a desired property; those are relevant features.
- Remember that our goal is to use these feature to identify genes related to a neurological disease (say AD).

Autoencoder: new decoding function

- Can we modify the AE to extract features that are more relevant to the neurological disease ?
- If we change the function q to be a predictive function, then features z are optimized for the prediction of the disease.

Autoencoder with *prediction decoder*

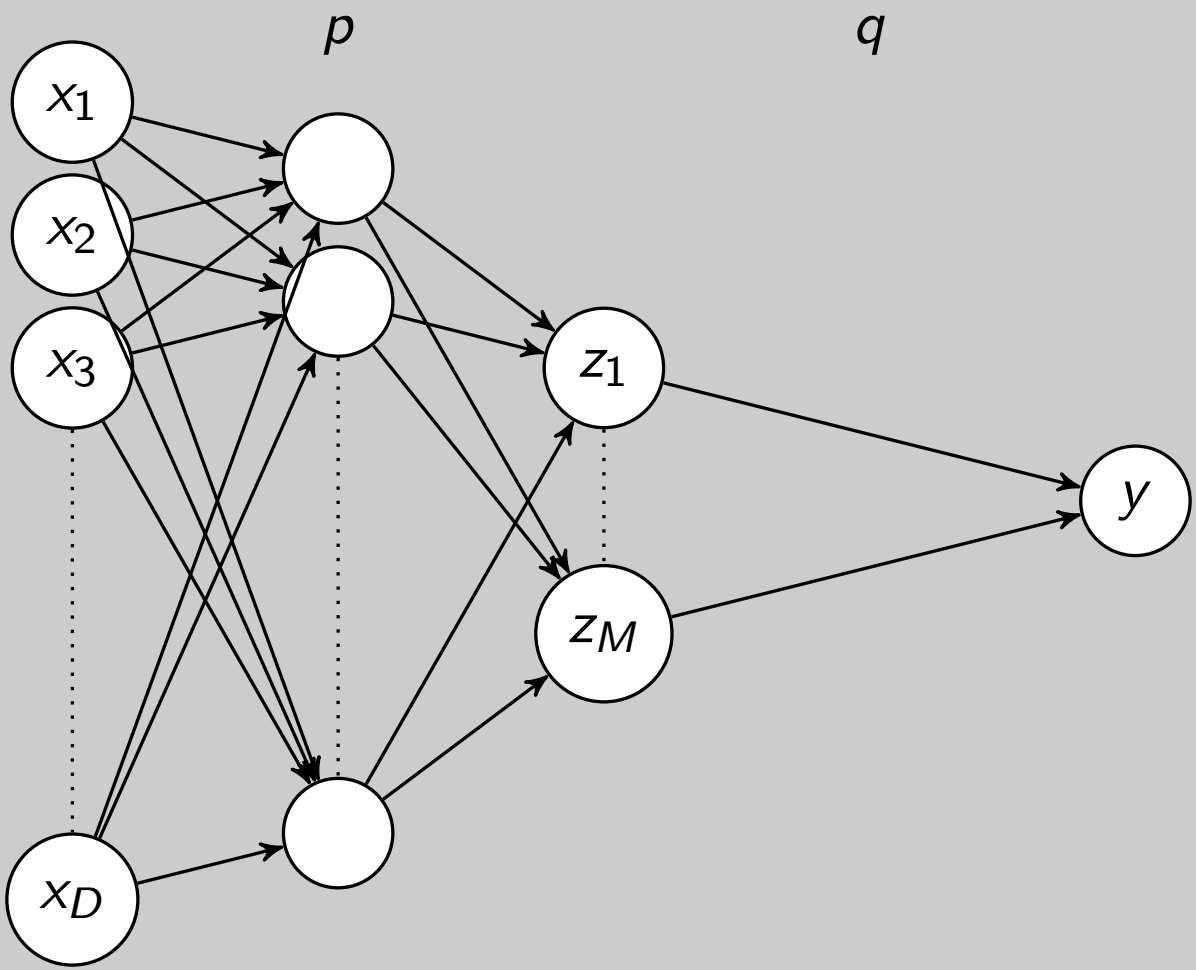

Figure: An AE where both functions p is NN and q is a classification function.

Extracting features with a NN for classification

- This turns the AE into a classification NN.
- The second to last layer (z in the previous slides) are the extracted features.
- These features are *one logistic regression away* from the neurological disease of interest.
- That way we make sure the extracted features are related with the neurological disease of interest.
- More broadly, we can use different objective functions to obtain different latent representations suited for different problems.

Extracting features with a NN for classification

- This idea stems from AEs, but simply put, we can use a predictive model to extract features.
- Any hidden layers between the input and the output is a lower-dimension representation of the input constructed to predict the response y.
- This feature selection process is automated, data-driven and construct features using all of the brain regions.
- Another key feature is that this process is done independently of every other steps. Which means we can preprocess images in any way, or conduct the genetic association study using any suited model.

Model Architecture

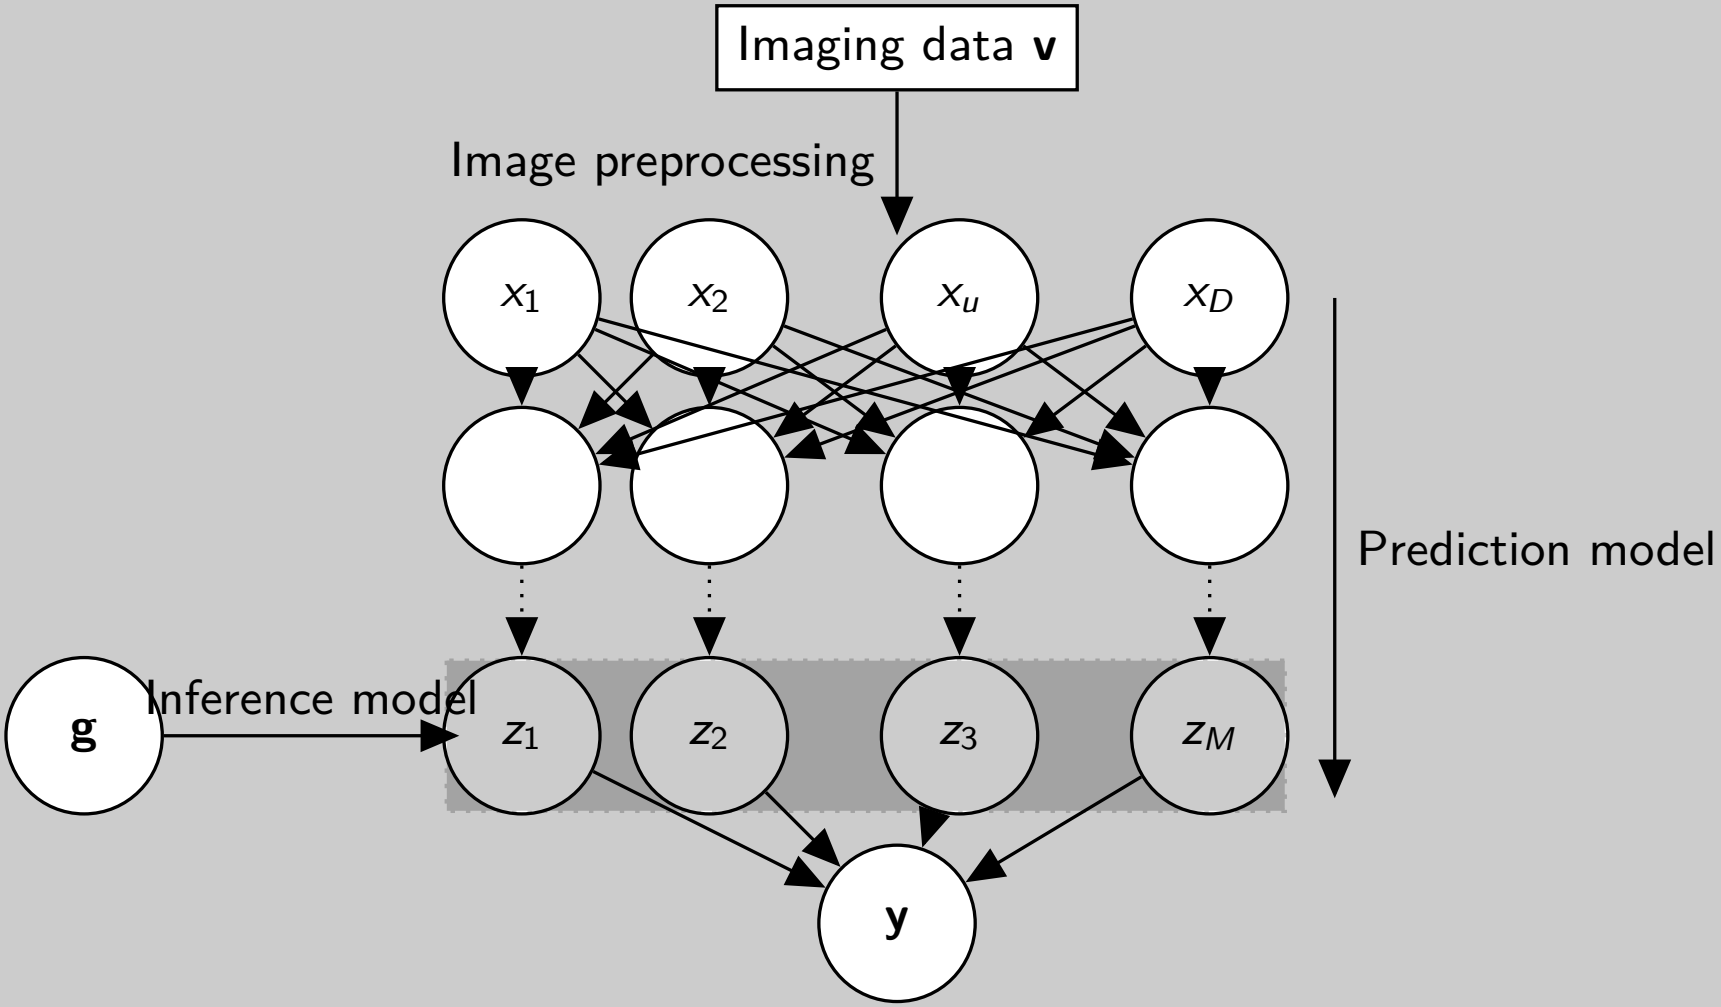

A first experiment

A first experiment on the ADNI1 data base.

- The proposed architecture can be used on a multitude of data sets, a wide range of inference model, various image processing techniques, different imaging information and a large collection of NN classifiers.
- Here, we introduce a first application of our pipeline.
- We used a cohort of subjects from the ADNI1 data base for which we had: disease status, SNPs and baseline MRI scans.
- *Actually* getting the data set in a digestible manner is surprisingly complicated.

The data set

- We have 543 subjects categorized as either NC, MCI, or AD.
- For every subject we have 521,014 SNPs.
- MRI scans were processed with FreeSurfer for registration but also to extract volumetric and cortical thickness statistics for a total 1860 imaging variables ($\mathbf{v}^*$). (Erin Gibson and Mirza F. Beg for the Functional Anatomical Imaging Shape Analysis Lab).

The models

- The prediction model is:
  1. Fully-connected NN
  2. 1 Hidden layer
  3. Trained to minimize the negative log likelihood loss
  4. Only trained on AD and NC subjects
- The inference model is the Bayesian group sparse multi-task regression model (BGSMR). (Farouk Nathoo and Leno Rocha)
- Group (gene) sparsity.
- Both models had multiple hyper-parameters selected using cross-validation. (Sidi Wu)

Results: Prediction accuracy

- We compare the NN-extracted feature with expert-selected feature from the literature.
- For predicting the patient status, the NN-extracted features are more accurate:

Features	Mean	Standard Dev.
Expert	0.81808	0.03552
NN	0.91726	0.02340

Table: Mean and standard deviation of the accuracy of a logistic regression that separate NC from AD using 2 different sets of features: the expert-selected features and the features automatically extracted by our proposed NN classifier.

Results: 2D neighbor-preserving embedding

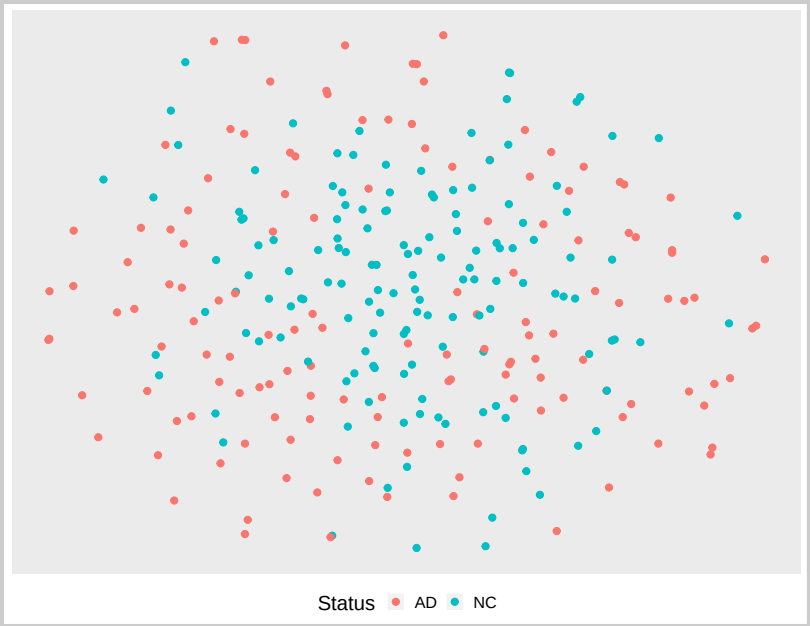

(a) Expert features

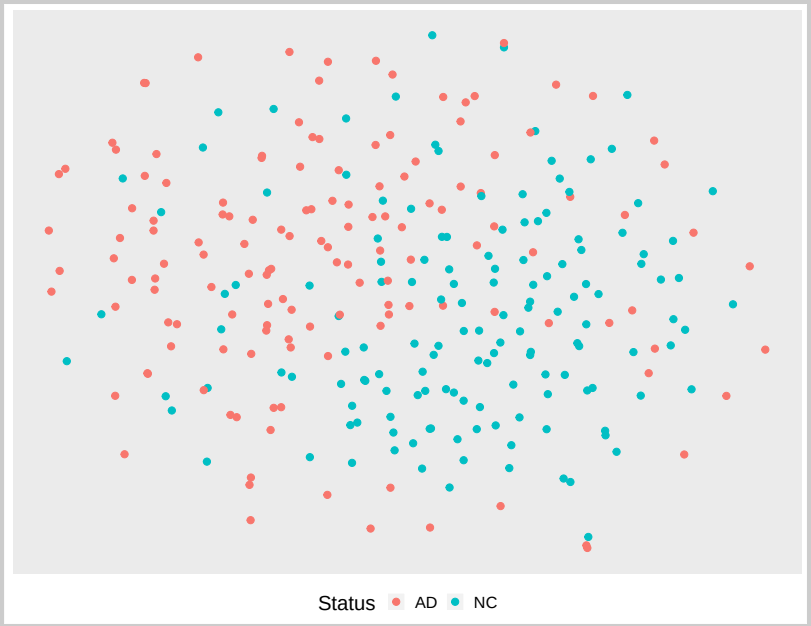

(b) NN-extracted features

Figure: 2-dimensional embedding using t-SNE. In contrast with PCA that finds a representation that captures as much variance as possible, SNEs try to identify a low-dimensional reprentation so as to optimally preserve neighborhood identity.

Results: Important genes

SNP	Gene	Chromosome	Status	No. of related features
rs2243581	SORCS1	10		1
rs1699105	SORL1	11	known	4
rs6511720	LDLR	19		15
rs6457200	NEDD9	6		8
rs11006130	TFAM	10	known	3
rs2025935	CR1	1	known	1
rs1568400	THRA	17	known	1
rs3785817	GRN	17		11
rs3026845	ECE1	1		1
rs12209631	NEDD9	6	known	5

Table: BGSMTR results: top 10 SNPs related to NN-extracted features with the highest standard score. The last column counts the number of NN-extracted features for which a 95% credible interval excluded zero. (Farouk Nathoo)

Discussion

Discussion

- There might exist multiple ways to compare the set of features extracted, the above experiment is a proof of concept analysis.
- In a future project we would like to replace FreeSurfer statistics with data-driven statistics.
- We could use a convolutional NN in place of FreeSurfer to directly consider the 3D neuroimages.

Discussion

- There are also alternative models for the genetic association study.
- We could use LASSO-NN to automatically subset SNPs to include in the Bayesian analyses.
- We could use Bayesian Model Averaging (BMA) instead of restraining ourselves to a single Bayesian model.

Conclusion

- The automatic feature extraction approach provides added value for imaging genetics studies:
  1. It requires no external expertise for feature selection.
  2. The features are built considering disease prediction through nonlinear representations of neuroimaging.
  3. It considers all of the brain regions jointly.

  A paper titled *Neural Network Disease Classification Based Feature Extraction for Imaging Genetics* is currently in its last stage of review and will be submitted in the next few weeks. (Reviewed by Michelle Miranda and Jiguo Cao).

I would love to answer your questions.

K. Greenlaw, E. Szefer, J. Graham, M. Lesperance, F. S. Nathoo, A. D. N. Initiative, A Bayesian group sparse multi-task regression model for imaging genetics, Bioinformatics 33 (16) (2017) 2513–2522

Islam, Y. Zhang, Brain MRI analysis for Alzheimer's disease diagnosis using an ensemble system of deep convolutional neural networks, Brain informatics 5 (2) (2018) 1–14.

W. Lin, T. Tong, Q. Gao, D. Guo, X. Du, Y. Yang, G. Guo, M. Xiao, M. Du, X. Qu, et al., Convolutional neural networks-based MRI image analysis for the Alzheimer's disease prediction from mild cognitive impairment, Frontiers in neuroscience 12 (2018) 777.

G. E. Hinton, S. Roweis, Stochastic neighbor embedding, Advances in neural information processing systems 15 (2002)