



HAL
open science

Applying Deep Machine Learning for psycho-demographic profiling of Internet users using O.C.E.A.N. model of personality

Iaroslav Omelianenko

► **To cite this version:**

Iaroslav Omelianenko. Applying Deep Machine Learning for psycho-demographic profiling of Internet users using O.C.E.A.N. model of personality. [Research Report] NewGround LLC. 2017. hal-01482770v2

HAL Id: hal-01482770

<https://hal.science/hal-01482770v2>

Submitted on 5 Jul 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC0 - Public Domain Dedication 4.0 International License

Applying Deep Machine Learning for psycho-demographic profiling of Internet users using O.C.E.A.N. model of personality

IAROSLAV OMELIANENKO

Research Director, NewGround LLC

yaric@newground.com.ua

Abstract

In the modern era, each Internet user leaves enormous amounts of auxiliary digital residuals (footprints) by using a variety of on-line services. All this data is already collected and stored for many years. In recent works, it was demonstrated that it's possible to apply simple machine learning methods to analyze collected digital footprints and to create psycho-demographic profiles of individuals. However, while these works clearly demonstrated the applicability of machine learning methods for such an analysis, created simple prediction models still lacks accuracy necessary to be successfully applied for practical needs. We have assumed that using advanced deep machine learning methods may considerably increase the accuracy of predictions. We started with simple machine learning methods to estimate basic prediction performance and moved further by applying advanced methods based on shallow and deep neural networks. Then we compared prediction power of studied models and made conclusions about its performance. Finally, we made hypotheses how prediction accuracy can be further improved. As result of this work, we provide full source code used in the experiments for all interested researchers and practitioners in corresponding GitHub repository. We believe that applying deep machine learning for psycho-demographic profiling may have an enormous impact on the society (for good or worse) and provides means for Artificial Intelligence (AI) systems to better understand humans by creating their psychological profiles. Thus AI agents may achieve the human-like ability to participate in conversation (communication) flow by anticipating human opponents' reactions, expectations, and behavior. By providing full source code of our research we hope to intensify further research in the area by the wider circle of scholars.

1. INTRODUCTION

By using various on-line services, modern Internet user leaves an enormous amount of digital tracks in the form of server logs, user-generated content, etc. All these information bits meticulously saved by on-line service providers create the vast amount of digital footprints for almost every Internet user. In recent research [Lambiotte, R., and Kosinski, M., 2014], it was demonstrated that by applying simple machine learning methods it was possible to find statistical correlations between digital footprints and psycho-demographic profile of individuals. The considered psycho-demographic profile comprise of psychometric scores based on five-factor O.C.E.A.N. model of personality [Goldberg et. al, 2006] and demographic scores such as *Age*, *Gender* and the *Political Views*. The O.C.E.A.N. is an abbreviation for *Openness* (Conservative and Traditional - Liberal and Artistic), *Conscientiousness* (Impulsive and Spontaneous - Organized and Hard Working), *Extroversion* (Contemplative - Engaged with outside world), *Agreeableness* (Competitive - Team working and Trusting), and *Neuroticism* (Laid back and Relaxed - Easily Stressed and Emotional).

In this work we decided to test whether applying advanced machine learning methods to analyze digital footprints of Internet users can outperform results of previous research conducted by M. Kosinski: *Mining Big Data to Extract Patterns and Predict Real-Life Outcomes*[Kosinski et. al, 2016]. For our experiments we used data corpus comprising of

psycho-demographic scores of individuals and their digital footprints in form of Facebook likes. The data corpus used in experiments kindly provided by M. Kosinski through corresponding web site: <http://dataminingtutorial.com>.

We started our experiments with building simple machine learning models based on linear/logistic regression methods as proposed by M. Kosinski in [Kosinski et. al, 2016]. By training and execution of simple models we estimated basic predictive performance of machine learning methods against available data set. Then we continued our experiments with advanced machine learning methods based on shallow and deep neural network architectures.

The full source code of our experiments provided in form of GitHub repository: <https://github.com/NewGround-LLC/psistats>

The source code is written in R programming language [R Core Team, 2015] which is highly optimized for statistical data processing and allows to apply advanced deep machine learning algorithms by bridging with Google Brain's TensorFlow framework [Google Brain Team, 2015].

This paper is organized as follows: In Section 2, we describe data corpus structure, and necessary data preprocessing steps to be applied. It is followed in Section 3 by details about how to build and run simple prediction models based on linear and logistic regression with results of their execution. In Section 4, we provide details how to create and execute advanced prediction models based on artificial neural networks. Finally, in Section 6 we compare the performance of different machine learning methods studied in this work and draw conclusions about the predictive power of studied machine learning models.

2. DATA CORPUS PREPARATION

In this section, we consider the creation of input data corpus from the publicly available data set, and it's preprocessing to allow further analysis by selected machine learning algorithms.

2.1. Data Set Description

The data set kindly provided by M. Kosinski and used in this work contains psycho-demographic profiles of $n_u = 110728$ Facebook users and $n_L = 1580284$ of associated Facebook likes. For simplicity and manageability, the sample is limited to U.S. users [Kosinski et. al, 2016]. The following three files can be downloaded from corresponding web site - <http://dataminingtutorial.com>:

1. *users.csv*: contains psycho-demographic user profiles. It has $n_u = 110728$ rows (excluding the row holding column names) and nine columns: anonymised user ID, gender ("0" for male and "1" for female), age, political views ("0" for Democrat and "1" for Republican), and scores of five-factor model of personality [Goldberg et. al, 2006].
2. *likes.csv*: contains anonymized IDs and names of $n_L = 1580284$ Facebook (FB) Likes. It has two columns: ID and name.
3. *users-likes.csv*: contains the associations between users and their FB Likes, stored as user-Like pairs. It has $nu_L = 10612326$ rows and two columns: user ID and Like ID. An existence of a user-Like pair implies that a given user had the corresponding Like on their profile.

2.2. Data pre-processing

The raw data preprocessing is an important step in machine learning analysis which will significantly reduce the time needed for analysis and result in better prediction power of created machine learning models.

The detailed description of data corpus preprocessing steps applied during this research given hereafter.

Construction of sparse users-likes matrix and matrix trimming

To use provided data corpus in machine learning analysis it should be transformed first into optimal format. Taking into account properties of provided data corpus (user can like specific topic only once and most users has generated small amount of likes) its natural to present it as *sparse* matrix where most of data points is zero (the resulting matrix density is about 0,006% - see Table 1). The *sparse* matrix data structure is optimized to perform numeric operations on sparse data and considerably reduce computational costs compared to the *dense* matrix used for same data set.

After users-likes *sparse* matrix creation, it was *trimmed* by removing rare data points. As a result, the significantly reduced data corpus was created, imposing even lower demands on computational resources and more useful for manual analysis to extract specific patterns. The descriptive statistics of users-likes matrix before and after trimming present in Table 1.

| Descriptive statistics | Raw Matrix | Trimmed Matrix |
|------------------------|------------|----------------|
| # of users | 110 728 | 19 742 |
| # of unique Likes | 1 580 284 | 8 523 |
| # of User-Like pairs | 10 612 326 | 3 817 840 |
| Matrix density | 0,006% | 2,269% |
| Likes per User | | |
| Mean | 96 | 193 |
| Median | 22 | 106 |
| Minimum | 1 | 50 |
| Maximum | 7 973 | 2 487 |
| Users per Like | | |
| Mean | 7 | 448 |
| Median | 1 | 290 |
| Minimum | 1 | 150 |
| Maximum | 19 998 | 8 445 |

Table 1: The descriptive statistics of raw and trimmed users-likes matrix with minimum users per like threshold set to $u_L = 150$ and minimum likes per user $L_u = 50$

The users-likes matrix can be constructed from provided three comma-separated files with the help of accompanying script written in R language: *src/preprocessing.R*. To use this script make sure that `input_data_dir` variable in the *src/config.R* points to the root directory where sample data corpus in the form of .CSV files were unpacked.

To start preprocessing and trimming, run the following command from terminal in the project's root directory:

```
$Rscript ./src/preprocessing.R \\  
-u 150 -l 50
```

where: $-u$ is the minimum number of users per like u_L , and $-l$ is the minimum number of likes per user L_u to keep in resulting matrix.

The values for the minimum number of users per like u_L and the minimum number of likes per user L_u was selected based on recommendations given in [Kosinski et. al, 2016]. We have experimented with other set of parameters as well ($u_L = 20$ and $L_u = 2$), but accuracy of trained prediction models degraded as result.

Data imputation of missed values

The raw data corpus has missed values in column with 'Political' *dependent variable* data. Before building the prediction model for this dependent variable, it is advisable to impute missed values. In this work, we applied multivariate imputation using **LDA** method with number of multiple imputations equals to $m = 5$ to fill missed values as described in [van Buuren, Groothuis-Oudshoorn, 2011].

The data imputation performed by the same *src/preprocessing.R* script, as part of users-likes matrix creation routine. The summary statistics for data imputation applied to *political* variable, presented in Table 2.

| | est | se | t | df | Pr(> t) | lo 95 | hi 95 | nmis | fmi | lambda |
|-------------|-------|------|--------|---------|-----------|-------|-------|------|------|--------|
| (Intercept) | 1.39 | 0.01 | 102.29 | 1240.53 | 0.00 | 1.36 | 1.41 | NA | 0.06 | 0.05 |
| gender | -0.02 | 0.01 | -2.80 | 25.27 | 0.01 | -0.04 | -0.01 | 0 | 0.44 | 0.40 |
| age | 0.00 | 0.00 | -0.73 | 577.61 | 0.47 | 0.00 | 0.00 | 0 | 0.09 | 0.08 |
| ope | -0.23 | 0.00 | -68.10 | 2446.16 | 0.00 | -0.23 | -0.22 | 0 | 0.04 | 0.04 |
| con | 0.05 | 0.00 | 10.92 | 20.28 | 0.00 | 0.04 | 0.06 | 0 | 0.49 | 0.44 |
| ext | 0.03 | 0.00 | 6.44 | 14.40 | 0.00 | 0.02 | 0.04 | 0 | 0.58 | 0.53 |
| agr | 0.02 | 0.00 | 6.72 | 189.32 | 0.00 | 0.02 | 0.03 | 0 | 0.15 | 0.14 |
| neu | -0.01 | 0.00 | -2.07 | 95.30 | 0.04 | -0.02 | 0.00 | 0 | 0.22 | 0.20 |

Table 2: The descriptive statistics for data imputation applied to political variable using **LDA** method with number of multiple imputations equals to $m = 5$. The plausibility of applied multivariate imputation can be confirmed by low values in column **fmi** and **lambda**. The column **fmi** contains the fraction of missing information as defined in [Rubin DB, 1987], and the column **lambda** is the proportion of the total variance that is attributable to the missing data $\lambda = \frac{B+B}{T^m}$.

Dimensionality reduction with SVD

After two previous steps, the resulting users-likes *sparse* matrix still has a considerable number of features per data sample: 8523 of feature columns. To make it even more maintainable, we considered applying singular value decomposition [Golub, G. H., and Reinsch, C. 1970], representing eigendecomposition-based methods, projecting a set of data points into a set of dimensions. As mentioned in [Kosinski et. al, 2016], reducing the dimensionality of data corpus has number of advantages:

- With reduced features space we can use fewer number of data samples, as it is required by most of the machine learning analysis algorithms that number of data samples exceeds the number of features (input variables)
- It will reduce risk of **overfitting** and increase statistical power of results
- It will remove **multicollinearity** and **redundancy** in data corpus by grouping related features (variables) in single dimension
- It will significantly reduce required computational power and memory requirements
- And finally it makes it easier to analyze data by hand over small set of dimensions as opposite to hundreds or thousands of separate features

To apply *SVD analysis* against generated users-likes matrix run the following command from project's root directory:

```
$Rscript ./src/svd_varimax.R \\  
  --svd_dimensions 50 --apply_varimax true
```

where: `--svd_dimensions` is the number of *SVD dimensions* for projection, and `--apply_varimax` is the flag to indicate whether *varimax rotation* should be applied afterwards.

Factor rotation analysis

The factor rotation analysis methodology can be used to further simplify *SVD dimensions* and increase their interpretability by mapping the original multidimensional space into a new, rotated space. Rotation approaches can be orthogonal (i.e., producing uncorrelated dimensions) or oblique (i.e., allowing for correlations between rotated dimensions).

In this work during data preprocessing we applied one of the most popular orthogonal rotations - *varimax*. It minimizes both the number of dimensions related to each variable and the number of variables related to each dimension, thus improving the interpretability of the data by human analysts.

For more details on rotation techniques, see [Abdi, H., 2003].

3. REGRESSION ANALYSIS

There is an abundance of methods developed to build prediction machine learning models suitable for analysis of large data sets. They ranging from sophisticated methods such as Deep Machine Learning [Goodfellow et al., 2016], probabilistic graphical models, or support vector machines [Cortes & Vapnik, 1995], to much simpler, such as linear and logistic regressions [Yan, Su, 2009].

Starting with simple methods is a common practice allowing the creation of good baseline prediction model with minimal computational efforts. The results obtained from these models can be used later to debug and estimate the quality of results obtained from advanced models.

3.1. Regression model description and model-specific data preprocessing

In our data corpus, we have *eight dependent variables* with psycho-demographic scores of individuals to be predicted. Among those variables, six have continuous values, and two has categorical values. To build the prediction model for variables with continuous values we applied linear regression analysis and for variables with categorical values - logistic regression analysis.

Hereafter we describe the rationale for selection of appropriate regression analysis methods as well as a description of model-specific data preprocessing needed.

Linear regression analysis

The linear regression is an approach for modeling the relationship between continuous scalar dependent variable y and one or more explanatory (or independent) variables denoted X . The case of one explanatory variable is called simple linear regression. For more than one explanatory variable, the process is called multiple linear regression

[David A. Freedman, 2009]. In linear regression, the relationships are modeled using linear predictor function $y = \Theta^T X$ whose unknown model parameters Θ estimated from the input data. Such models are called linear models [Hilary L. Seal, 1967].

We used linear regression to build prediction models for analysis of six continuous dependent variables in given data corpus: *Age, Openness, Conscientiousness, Extroversion, Agreeableness, and Neuroticism*

Logistic regression analysis

The logistic regression is a regression model where the dependent variable is categorical [David A. Freedman, 2009]. It measures the relationship between the categorical dependent variable and one or more independent variables by estimating probabilities using a logistic function $\sigma(x) = \frac{1}{1+e^{-x}}$, which is the cumulative logistic distribution [Rodriguez, G., 2007].

We considered only specialized binary logistic regression because categorical dependent variables found in our data corpus (*Gender* and *Political Views*) are binominal, i.e. have only two possible types, "0" and "1".

Cross-Validation

We applied *k-fold cross-validation* to help with avoiding model overfitting when evaluating accuracy scores of prediction models. In *k-fold cross-validation*, the original sample is randomly partitioned into k equal sized subsamples. Of the k subsamples, a single subsample is retained as the validation data for testing the model, and the remaining $k - 1$ subsamples are used as training data. The cross-validation process is then repeated k times (the folds), with each of the k subsamples used exactly once as the validation data. The k results from the folds can then be averaged to produce a single estimation. The advantage of this method is that all observations are used for both training and validation, and each observation is used for validation exactly once. The 10-fold cross-validation is commonly used, but in general, k remains an unfixed parameter [Kohavi, Ron, 1995].

Dimensionality reduction

To reduce the number of features (input variables) in the data corpus was applied singular value decomposition (SVD) with subsequent *varimax* factor rotation analysis. The number of the varimax-rotated singular value decomposition dimensions (K) has a considerable impact on the accuracy of model predictions. To find an optimal number of *SVD dimensions*, we performed analysis of relationships between K and accuracy of model predictions by creating series of regression models for different values of K . Then we plotted prediction accuracy of regression models against chosen number of K SVD dimensions. Typically the prediction accuracy grows rapidly within lower ranges of K and may start decreasing once the number of dimensions becomes large. Selecting value of K that marks the end of a rapid growth of prediction accuracy values usually offers decent interpretability of the input data topics. In general, the larger K values often results in better predictive power when preprocessed data corpus further analyzed with specific machine learning algorithm [Zhang, Marron, Shen,& Zhu, 2007]. See Figure 1 for results of our experiments.

To start SVD analysis run the following command from terminal in the project's root directory:

```
$Rscript ./src/analysis.R
```

The resulting plots will be saved as "Rplots.pdf" file in the project root and include two plots:

- the plot with relationships between the accuracy of prediction models for each dependent variable and the number of the varimax-rotated SVD dimensions used for dimensionality reduction (Figure 1). With this plot,

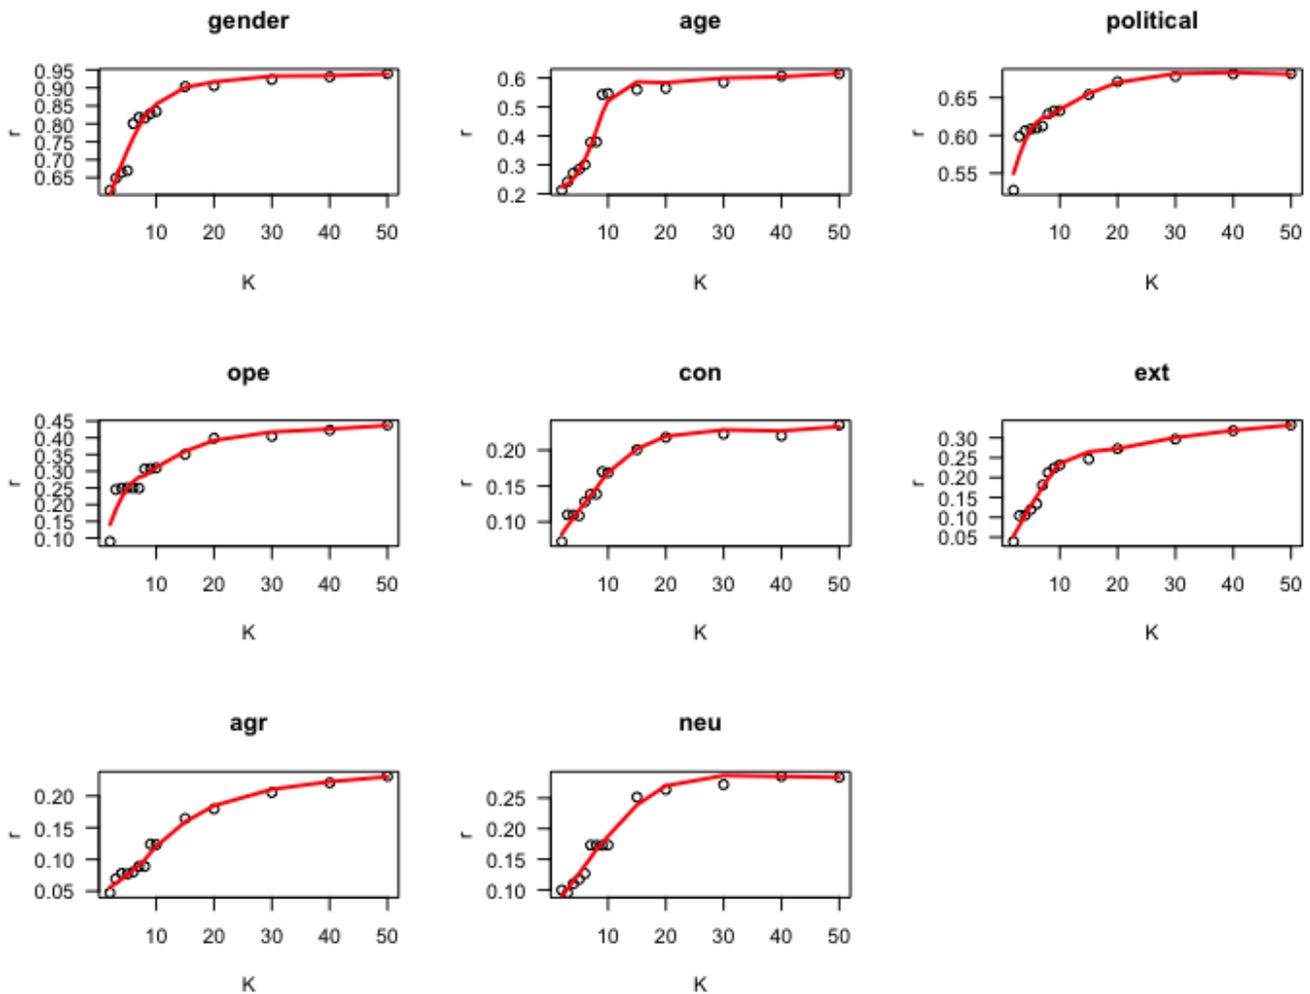


Figure 1: The relationship between the accuracy of predicting psycho-demographic traits and the number of the varimax-rotated singular value decomposition dimensions used for dimensionality reduction. The results suggest that selecting $K = 50$ SVD dimensions might be a good choice for building models predicting almost all dependent variables, as it offers accuracy that is close to what seems like the higher asymptote for this data. But for Openness, Extroversion, and Agreeableness dependent variables prediction results can be slightly improved with higher numbers of K SVD dimensions selected.

it's easy to visually find an optimal number of K SVD dimensions to maximize predicting power of regression model per particular dependent variable.

- the heat map of correlations between scores of digital footprints of individuals projected on specific number of varimax-rotated SVD dimensions and each dependent variable (Figure 2). This plot can be used to find most correlated dependent variables visually. Later, it will be shown that predictive models for dependent variables with higher correlation have better prediction accuracy.

3.2. Implementing simple prediction models and its accuracy evaluation

The given data corpus has eight dependent variables for which to build prediction models. The simple machine learning methods such as regression analysis mostly applied to estimate single dependent variable. But in case when multiple dependent variables need to be estimated the specialized methods of *multivariate* regression analysis can be used. Taking into account that our dependent variables have different types (continuous and nominal) which require different regression analysis methods to be applied, we decided to build separate regression models per each dependent variable. The metric to evaluate accuracy of prediction model is related to the regression method used in the model. In this research we have considered following metrics:

- the accuracy of prediction model applied to the continuous dependent variable will be measured as Pearson product-moment correlation [Gain, 1951]
- the accuracy of prediction model applied to the bi-nominal dependent variable will be measured as area under the receiver-operating characteristic curve coefficient (AUC) [Sing, Sander, Beerenwinkel, Lengauer, 2005]

Before executing models make sure that data corpus already preprocessed as described in Subsection: "Construction of sparse users-likes matrix and matrix trimming"

When data corpus is ready, the following command can be executed to start linear/logistic regression models building and its predictive performance evaluation (run command from terminal in the project's root directory):

```
$Rscript ./src/regression_analysis.R
```

The results will be saved into the file "*out/pred_accuracy_regr.txt*". The prediction accuracy of regression models for data corpus trimmed to contain 150 users-per-like and 50 likes-per-user and varimax-rotated against $K = 50$ SVD dimensions presented in Table 3.

| Trait | Variable | Pred. accuracy |
|-------------------|-----------|----------------|
| Gender | gender | 93.65% |
| Age | age | 61.17% |
| Political view | political | 68.36% |
| Openness | ope | 44.02% |
| Conscientiousness | con | 25.72% |
| Extroversion | ext | 30.26% |
| Agreeableness | agr | 23.97% |
| Neuroticism | neu | 29.11% |
| | Mean | 47.03% |

Table 3: The predictive accuracy of linear and logistic regression models per depended variable (for $u_L = 150$, $L_u = 50$, and $K = 50$ SVD dimensions).

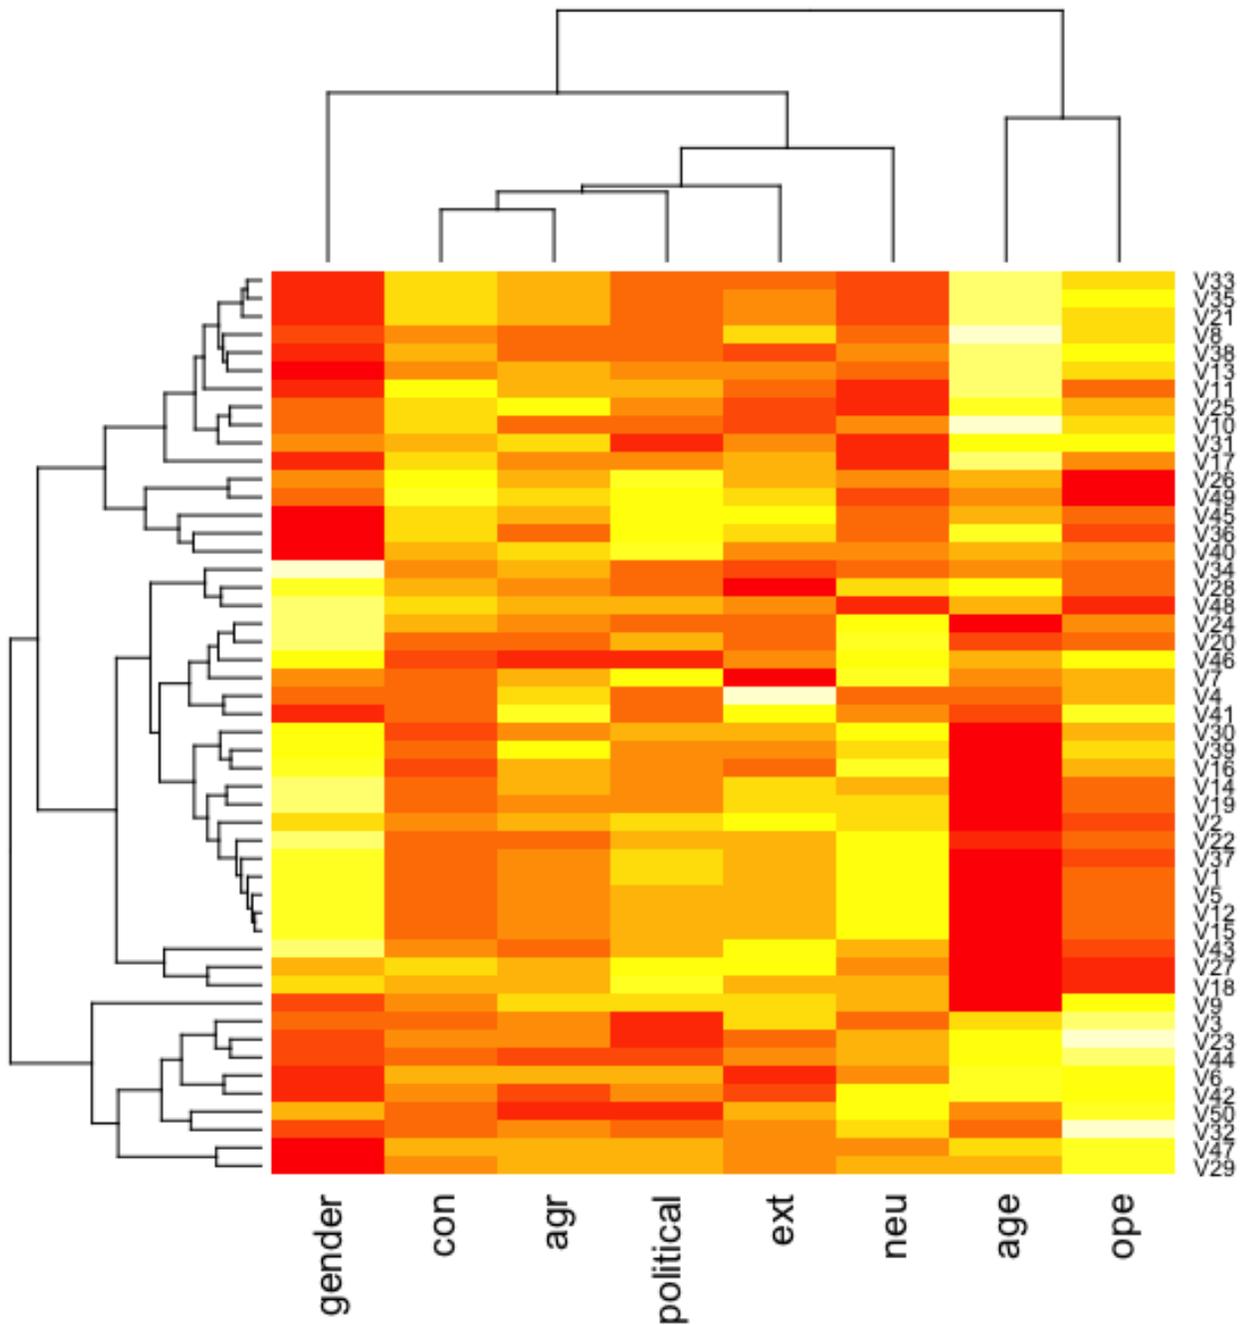


Figure 2: The heat map is presenting correlations between scores of digital footprints of individuals projected on $K = 50$ varimax-rotated singular value decomposition dimensions and scores of psycho-demographic traits of individuals. The heat map suggests that Age, Gender, and the Political view dependent variables have maximum correlation with a maximal number of SVD dimensions. The higher correlation will result in higher prediction power of regression model for particular dependent variable (which will be shown later).

From the Table 3 it can be seen that prediction accuracy of linear/logistic regression models differs per each dependent variable. Furthermore for most variables the accuracy is too low to be applied in real-life predictions. The most accurate predictions made for *Gender*, *Age*, and *Political view* with *Openness* following after. That correlates well with our previous analysis of SVD correlations heat map (see Figure 2). In general only prediction model for *Gender* is accurate enough to be useful in real-life applications. Thus, simple linear/logistic regression models can not be used to accurately estimate psycho-demographic profiles of Internet users based only on their Facebook likes.

In following sections, we will test if applying advanced deep machine learning methods can improve prediction accuracy any further.

4. FULLY CONNECTED FEED FORWARD ARTIFICIAL NEURAL NETWORKS

In this work, we considered multilayer *fully connected feed-forward* neural networks (NN) for building simple (shallow) and deep machine learning NN models. The *fully connected* NN characterized by interconnectedness of all units of one layer with all units of the layer before it in the graph. The *feed-forward* NN is not allowed to have cycles from latter layers back to the earlier.

Hereafter we will describe Artificial NN architectures evaluated and prediction accuracy results obtained.

4.1. The Shallow Feed Forward Artificial NN evaluation

A *Shallow Neural Network* (SNN) is an artificial neural network with one hidden layer of units (neurons) between input and output layers. Its hidden units (neurons) take inputs from input units (columns of input data matrix) and feeds into the output units, where linear or categorical analysis performed. To mimic biological neuron, hidden units in the neural network apply specific non-linear activation functions. One of the popular activation functions is ReLU non-linearity that we considered as activation function for the units of hidden layers in studied network architectures [Nair, Hinton, 2010]. It improves information disentangling and linear separability producing efficient variable size representation of model's data. Furthermore ReLU activation is computationally cheaper: there is no need for computing the exponential function as in case of sigmoid activation [Glorot, Bordes, Bengio, 2011].

To reduce overfitting was applied dropout regularization with drop-probability 0.5, which means that each hidden unit if randomly omitted from the network with the specified probability. This helps to break the rare dependencies that can occur in the training data [Hinton, G. et al., 2012].

The NN architecture was build using Google Brain's TensorFlow library - an open source software library for numerical computation using data flow graphs. [Google Brain Team, 2015] Nodes in the graph represent mathematical operations, while the graph edges represent the multidimensional data arrays (tensors) communicated between them. The resulting two layer (one hidden layer) ANN's architecture graph depicted in Figure 5

As the *loss function* to be optimized was selected *Mean Squared Error* (MSE) with *Adam* optimizer (Adaptive Moment Estimation) to estimate it's minimum. The Adam optimizer was selected for *it's proven advantages some of them are that the magnitudes of parameter updates are invariant to rescaling of the gradient, its step sizes are approximately bounded by the step size hyper-parameter, it does not require a stationary objective, it works with sparse gradients, and it naturally performs a form of step size annealing* [Kingma, Ba, 2014]. The batch size was selected to be 100. We also tested training with batch size 10 but found no statistically relevant prediction accuracy difference between runs with either batch size but reducing batch size considerably increased training time of NN models.

It was found that optimal number of SVD dimensions for Shallow ANN is $K = 128$, with number of units in the hidden layer 512, and learning rate $\gamma = 0.0001$. See Table 4.

| Trait | Variable | Prediction accuracy, $\gamma = 0.0001$ | | | |
|-------------------|-----------|--|---------|---------|---------|
| | | K = 50 | K = 128 | K = 256 | K = 512 |
| Gender | gender | 93.76% | 93.61% | 93.00% | 93.82% |
| Age | age | 85.00% | 85.07% | 85.03% | 83.59% |
| Political view | political | 66.65% | 66.86% | 67.28% | 65.32% |
| Openness | ope | 47.22% | 51.24% | 47.26% | 44.81% |
| Conscientiousness | con | 28.56% | 28.45% | 29.65% | 28.26% |
| Extroversion | ext | 32.53% | 32.87% | 30.44% | 30.32% |
| Agreeableness | agr | 25.21% | 27.66% | 24.58% | 24.80% |
| Neuroticism | neu | 33.53% | 36.24% | 33.09% | 31.96% |
| | Mean | 51.58% | 52.75% | 51.29% | 50.36% |

Table 4: The predictive accuracy results of SNN per K SVD dimensions with learning rate $\gamma = 0.0001$ and 512 units in the hidden layer.

The optimal learning rate was selected by comparing ratio of survived hidden units with *non zero ReLU activation* and monitoring loss function over iterations plot (Figure 3). With presented hyper-parameters it was achieved maximum ratio 0.57 of zero ReLU activations for largest learning rate value, which is acceptable taking into account tendency of ReLU to saturate at zero during gradient back propagation stage when strong gradients applied, due to high learning rates [Nair, Hinton, 2010]. The maximal number of iterations (50 000) and correspondingly number of training epochs was selected based on loss function plot (Figure 3). With series of experiments it was selected as optimal the learning rate value $\gamma = 0.0001$, which gives smooth loss function, stable acceptable number of "dead" neurons after ReLU activation (ReLU zero activations ratio is about 0.5), and best prediction scores among runs (see Table 5).

| Trait | Variable | Prediction accuracy, K = 128 SVD | | |
|-------------------|-----------|----------------------------------|-------------------|--------------------|
| | | $\gamma = 0.001$ | $\gamma = 0.0001$ | $\gamma = 0.00001$ |
| Gender | gender | 93.88% | 93.61% | 91.97% |
| Age | age | 83.68% | 85.07% | 73.81% |
| Political view | political | 66.67% | 66.86% | 66.95% |
| Openness | ope | 49.43% | 51.24% | 45.92% |
| Conscientiousness | con | 27.18% | 28.45% | 27.79% |
| Extroversion | ext | 29.93% | 32.87% | 31.93% |
| Agreeableness | agr | 25.38% | 27.66% | 24.51% |
| Neuroticism | neu | 31.15% | 36.24% | 32.24% |
| | Mean | 50.91% | 52.75% | 49.39% |

Table 5: The predictive accuracy results of SNN per learning rate with K = 128 SVD dimensions and 512 units in the hidden layer.

The accompanying launch script provided to conduct experiments under Unix:

```
./eval_mlp_1.sh ul_svd_matrix_file
```

where: **ul_svd_matrix_file** the path to the preprocessed users-likes matrix with dimension of feature columns reduced as described in Construction of sparse users-likes matrix and matrix trimming

The source code of shallow ANN implementation used for experiment can be found in *src/mlp.R* of accompanying

GitHub repository.

4.2. Feed Forward Deep Learning Networks Architecture Evaluation

A *Deep Neural Network* (DNN) is an artificial neural network with multiple hidden layers of units between the input and the output layers. The first hidden layer will take inputs from each of the input units and the subsequent hidden layer will take inputs from the outputs of previous hidden layer’s units [Christopher M. Bishop, 1995]. Similar to shallow, deep neural network can model complex non-linear relationships. But added extra layers enable composition of features from lower layers, giving the potential of modeling complex data with fewer units than a similarly performing shallow network [Bengio, Yoshua, 2009]. *Deep learning discovers intricate structure in large data sets by using the backpropagation algorithm to indicate how a machine should change its internal parameters that are used to compute the representation in each layer from the representation in the previous layer* [LeCun, Bengio, Hinton, 2015].

As with shallow ANNs, many issues can arise with training of the deep neural networks, with two most common problems - are overfitting and computation time [Tetko, Livingstone, Luik, 1995].

Hereafter we will consider DNN architectures studied in this research.

The three-layer Deep Learning Network Architecture Evaluation

Our experiments with deep learning networks we started with simple DNN architecture comprising of two hidden layers with ReLU activation and dropout after each hidden layer with keep probability of 0.5. The experimental network graph depicted in Figure 7.

| | | Prediction accuracy, $\gamma = 0.0001$ | | | | |
|-------------------|-----------|--|-----------|-----------|-----------|------------|
| | | $K = 50$ | $K = 128$ | $K = 256$ | $K = 512$ | $K = 1024$ |
| Trait | Variable | 512,256 | 512,256 | 512,256 | 1024,512 | 2048,1024 |
| Gender | gender | 92.60% | 92.14% | 92.53% | 93.61% | 95.68% |
| Age | age | 85.03% | 82.97% | 80.67% | 81.45% | 82.53% |
| Political view | political | 66.66% | 67.34% | 67.06% | 67.26% | 65.06% |
| Openness | ope | 44.76% | 47.56% | 47.68% | 46.39% | 40.61% |
| Conscientiousness | con | 24.33% | 25.79% | 25.68% | 27.74% | 24.67% |
| Extroversion | ext | 30.86% | 32.35% | 34.35% | 28.16% | 26.53% |
| Agreeableness | agr | 24.83% | 25.66% | 26.48% | 21.65% | 29.90% |
| Neuroticism | neu | 32.17% | 33.01% | 35.94% | 28.50% | 30.83% |
| Mean | | 50.15% | 50.85% | 51.30% | 49.34% | 49.48% |

Table 6: The predictive accuracy results of three layer DNN per K SVD dimensions with learning rate $\gamma = 0.0001$, with sizes of hidden layers presented in third table header row as $[hidden_1, hidden_2]$.

We started with learning rate $\gamma = 0.0001$ which resulted in best prediction performance for shallow ANN and attempted series of experiments to estimate optimal value of K SVD dimensions. The optimal prediction accuracy of DNN model was achieved with $K = 256$ SVD dimensions and two hidden layers comprising of $[512, 256]$ units correspondingly. Similar prediction accuracy can be achieved with $K = 1024$ SVD dimensions and $[2048, 1024]$ units per layer with learning rate $\gamma = 10^{-5}$. But we have not considered later set of hyper-parameters due to its extra computational overhead while giving statistically same results as former set. The results of experiments presented in Table 6.

After finding optimal values of K SVD dimensions and number of units per hidden layer, we have conducted series of experiments to determine optimal *initial* learning rate value for found hyper-parameters. It was experimentally confirmed that *initial* learning rate value $\gamma = 0.0001$ is the optimal one. See Table 7

| Trait | Variable | Prediction accuracy, $K = 256$ SVD | | |
|-------------------|-----------|------------------------------------|-------------------|--------------------|
| | | $\gamma = 0.001$ | $\gamma = 0.0001$ | $\gamma = 0.00001$ |
| Gender | gender | 90.89% | 92.53% | 82.75% |
| Age | age | 79.60% | 80.67% | 79.44% |
| Political view | political | 67.33% | 67.06% | 59.69% |
| Openness | ope | 44.73% | 47.68% | 39.48% |
| Conscientiousness | con | 21.41% | 25.68% | 22.86% |
| Extroversion | ext | 24.62% | 34.35% | 24.07% |
| Agreeableness | agr | 17.27% | 26.48% | 16.34% |
| Neuroticism | neu | 27.72% | 35.94% | 28.53% |
| | Mean | 46.70% | 51.30% | 44.14% |

Table 7: The predictive accuracy results of three layer DNN per learning rate with $K = 256$ SVD dimensions, and with sizes of hidden layers 512 and 256 correspondingly.

In our experiments, we applied exponential learning rate decay with the number of steps before decay 10000 and decay rate of 0.96. Such scheme has positive effect on network convergence speed due to the learning rate annealing effect, which gives a system the ability to escape from poor local minima to which it might have been initialized [Kirkpatrick et al., 1983]. We selected batch size 100 as optimal for this experiment.

The accompanying launch script provided to conduct experiments under Unix:

```
$. / eval_dnn.sh ul_svd_matrix_file
```

where: `ul_svd_matrix_file` is the path to the preprocessed users-likes matrix with dimension of feature columns reduced as described in Construction of sparse users-likes matrix and matrix trimming

The source code of DNN with two hidden layers implementation used for experiment can be found in `src/dnn.R` of accompanying GitHub repository.

The four-layer Deep Learning Network Architecture Evaluation

This architecture comprise of three hidden layers with ReLU activation and one output linear layer. All network layers are fully connected and network architecture is feed-forward as in all previous NN experiments. The experimental network graph depicted in Figure 8.

We have tested two dropout regularization schemes: (a) dropout applied after each hidden layer with keep-probability 0.5; (b) dropout applied after each second hidden layer with keep probability calculated by formula: $p_d = \frac{i}{2n}$, (where: n - number of dropouts, i - current dropout index).

It was found that former scheme gives better results than the last one. Thus for final evaluation run, we applied dropout regularization *after each hidden layer*.

Based on our previous experiments with more shallow networks we decided to start with following hyper-parameters: learning rate $\gamma = 0.0001$, learning rate decay step 10000 with decay rate 0.96, $K = 128$ SVD dimensions, and hidden layers configuration - [256, 128, 128].

| | | Prediction accuracy, $\gamma = 0.0001$ | | | |
|-------------------|-----------|--|-------------|--------------|----------------|
| | | $K = 128$ | $K = 256$ | $K = 512$ | $K = 1024$ |
| Trait | Variable | 256,128,128 | 512,256,256 | 1024,512,512 | 2048,1024,1024 |
| Gender | gender | 55.07% | 70.63% | 92.13% | 91.03% |
| Age | age | 83.82% | 82.78% | 83.80% | 84.96% |
| Political view | political | 57.42% | 61.23% | 64.52% | 66.52% |
| Openness | ope | 13.52% | 27.50% | 44.88% | 44.86% |
| Conscientiousness | con | 14.41% | 14.92% | 20.76% | 28.72% |
| Extroversion | ext | 4.06% | 7.30% | 22.77% | 26.85% |
| Agreeableness | agr | 10.12% | 10.90% | 17.69% | 20.02% |
| Neuroticism | neu | 8.81% | 11.98% | 27.48% | 28.97% |
| Mean | | 30.91% | 35.90% | 46.75% | 48.99% |

Table 8: The predictive accuracy results of three layer DNN per K SVD dimensions with the learning rate $\gamma = 0.0001$. The number of units in hidden layers differ per configuration and presented as $[hidden_1, hidden_2]$ in third table header row

To find the optimal number of K SDV dimensions and hidden layers configurations we have conducted series of experiments trying various combinations. The heuristic applied to select the number of units per hidden layer is rather naive and assumes that with dropout probability of 0.5 half of the units will be saturated to zero at ReLU activation. Thus we decided to have the number of units in the first hidden layer to be twice as much as the number of features in input data (K). The results of experiments present in Table 8

Despite the fact that best accuracy was achieved with $K = 1024$ SDV dimensions and the number of units in hidden layers $[2048, 1024, 1024]$ respectively, it was detected slight model overfitting during training/validation with these hyper-parameters applied. So, we have decided to stop increasing K and number of hidden units as it will give no further gain in prediction accuracy against validation data set and even may lead to worsening of validation accuracy with greater overfitting level. (See Figure 6)

The accompanying launch script provided to conduct experiments under Unix:

```
$/ eval_3dnn.sh ul_svd_matrix_file
```

where: `ul_svd_matrix_file` is the path to the preprocessed users-likes matrix with dimension of feature columns reduced as described in Construction of sparse users-likes matrix and matrix trimming

The source code of DNN with three hidden layers implementation used for experiments can be found in `src/3dnn.R` of accompanying GitHub repository.

5. FUTURE WORK

With conducted experiments, we have found that prediction accuracy differs considerably among machine learning methods studied and best results was achieved by using advanced methods based on neural networks architectures. At the same time, the prediction accuracy per individual dependent variable also differs per particular prediction model and selected set of hyper-parameters. From experimental results it can be seen that specific combination of NN architecture with given set of hyper-parameters are best suited for one dependent variable but worsened predictive power for some of the others.

In future studies, it is interesting to investigate this dependency and build separate NN models per each dependent variable as it was done in case of simple machine learning methods (see Section: 'Regression analysis')

Also, it seems promising to apply methodology described in [Ba, Caruana, 2014] which provide evidence that shallow networks are capable of learning the same functions as deep learning networks, and often with the same number of parameters as the deep learning networks. In [Ba, Caruana, 2014] it was shown that with wide shallow networks it's possible to reach the state-of-the-art performance of deep models and reduce training time by the factor of 10 using parallel computational resources (GPU).

6. CONCLUSION

From our experiments we found that weak correlation exists between most of *O.C.E.A.N.* psychometric scores of individuals and collected Facebook likes associated with them. Either simple or advanced machine learning algorithms that we have tested provided *poor prediction accuracy* for almost all *O.C.E.A.N.* personality traits. It seems not feasible yet to use machine learning models to accurately estimate *psychometric profile* of an individual based only on Facebook likes. But we believe that by complementing Facebook likes of user with additional data points, it is possible to greatly improve accuracy of machine learning prediction models for psychometric profile estimation.

At the same time, we have found a strong correlation with *demographic traits* of individuals such as *Age, Gender,* and the *Political Views* with their Facebook activity (likes). Our experiments confirmed that its possible to use advanced machine learning methods to build the correct demographic profile of an individual based only on collected Facebook likes.

| Trait | Var. | DNN | | | |
|----------------|--------|------------|---------|---------|----------------|
| | | Regression | SNN | 512,256 | 2048,1024,1024 |
| | | K = 50 | K = 128 | K = 256 | K = 1024 |
| Gender | gen. | 93.65% | 93.61% | 92.53% | 91.03% |
| Age | age | 61.17% | 85.07% | 80.67% | 84.96% |
| Political view | polit. | 68.36% | 66.86% | 67.06% | 66.52% |
| Openness | ope | 44.02% | 51.24% | 47.68% | 44.86% |
| Conscientious. | con | 25.72% | 28.45% | 25.68% | 28.72% |
| Extroversion | ext | 30.26% | 32.87% | 34.35% | 26.85% |
| Agreeableness | agr | 23.97% | 27.66% | 26.48% | 20.02% |
| Neuroticism | neu | 29.11% | 36.24% | 35.94% | 28.97% |
| | Mean | 47.03% | 52.75% | 51.30% | 48.99% |

Table 9: The comparison of prediction accuracy for best prediction models found. The best prediction accuracy demonstrated by shallow neural network (SNN) followed by three-layered deep neural network (DNN [512,256]). The further increase of K SVD dimensions and adding of extra hidden layers lead to model overfitting and degradation of accuracy over validation data set.

Among all studied machine learning prediction models the best overall accuracy was achieved with Shallow Neural Network architecture. We hypothesize that this may be the result of its ability to learn best parameters space function within an optimal number of *SVD dimensions* applied to the input data set (users-likes sparse matrix). Adding extra hidden layers either leads to model *overfitting* when the number of SVD dimensions is too high, or *underfitting* when the number of SVD dimensions is too low. Also, it's interesting to notice that performance of shallow networks and deep learning networks with two hidden layers are comparable, while with introducing of third and more hidden

layers it drops significantly. Thus we can conclude that no further improvements can be gained with extra hidden layers introduced. See Table 9.

Gathered experimental data confirms that advanced machine learning methods based on variety of studied artificial neural networks architectures outperform simple machine learning methods (linear and logistic regression) described in previous research conducted by M. Kosinski [Kosinski et. al, 2016]. We believe that further prediction accuracy improvements can be achieved by building separate advanced machine learning models per dependent variable, which is the subject of our future research activities.

A. THE SNN EVALUATION PLOTS

The following pages provide plots and diagrams related to evaluation of two layer (shallow) feed forward artificial neural network with one fully connected hidden layer and linear output layer.

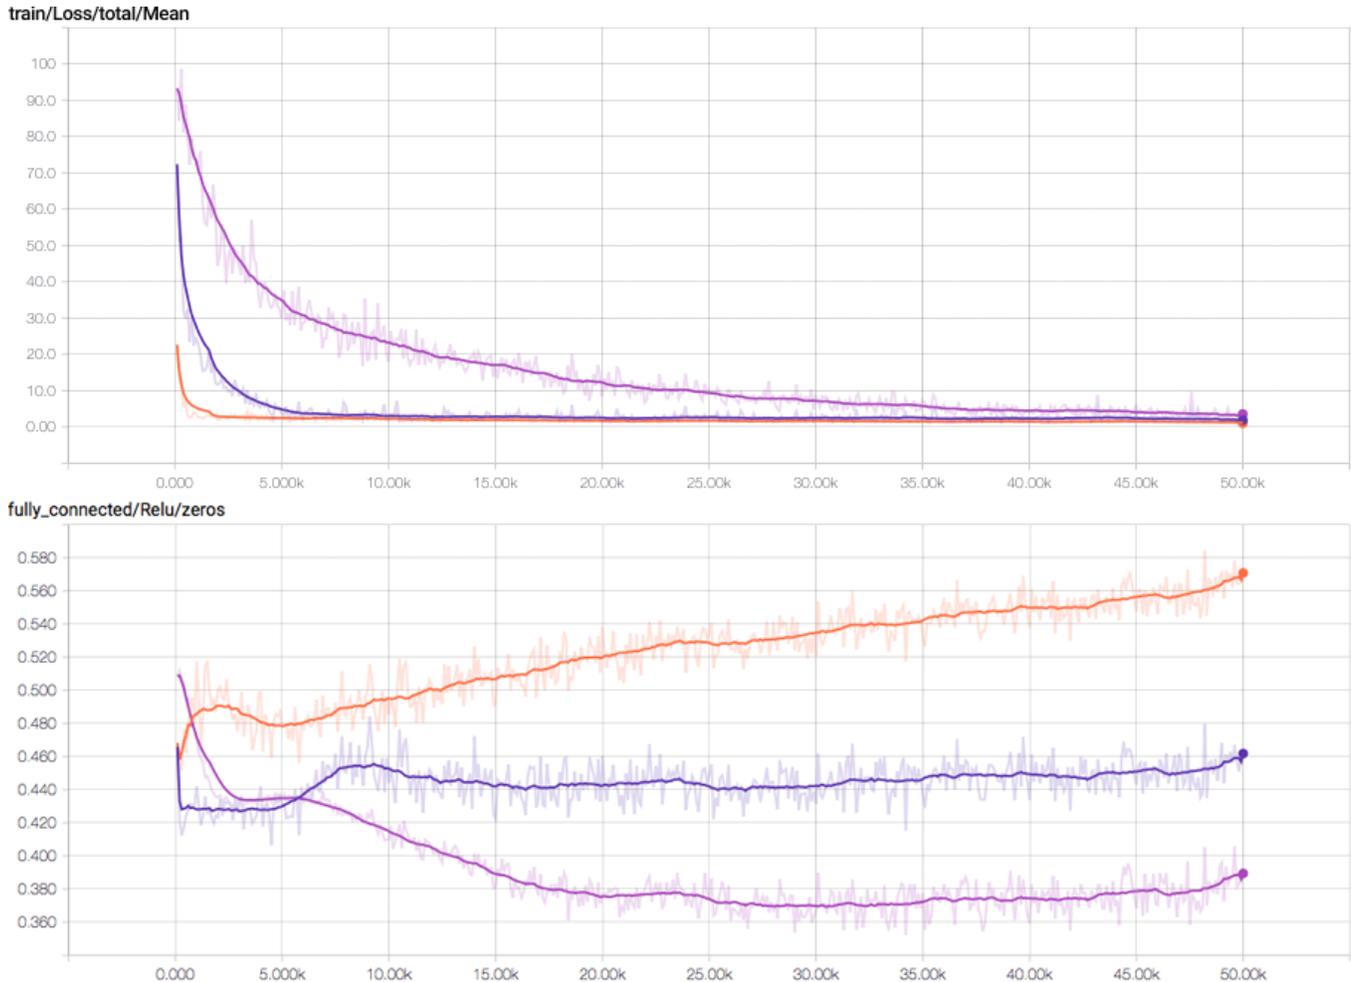


Figure 3: The training process evaluation based on loss values and ReLU zero activations per number of iterations. With higher learning rate ($\gamma = 0.001$ - orange) we have fast convergence but ratio of ReLU-zero activations is higher than 0.5 and quickly rising with relatively low evaluated prediction accuracy, which implies that optimum was missed. With medium learning rate ($\gamma = 0.0001$ - violet) we have smooth loss function plot with ratio of ReLU-zero activations bellow 0.5, giving best prediction scores among all three runs. With lowest learning rate ($\gamma = 0.00001$ - purple) we can see that learning struggled to find global minimum, reduced speed of convergence, and despite the lowest ReLU zero activations rate - worst prediction accuracy among all runs due to high loss values.

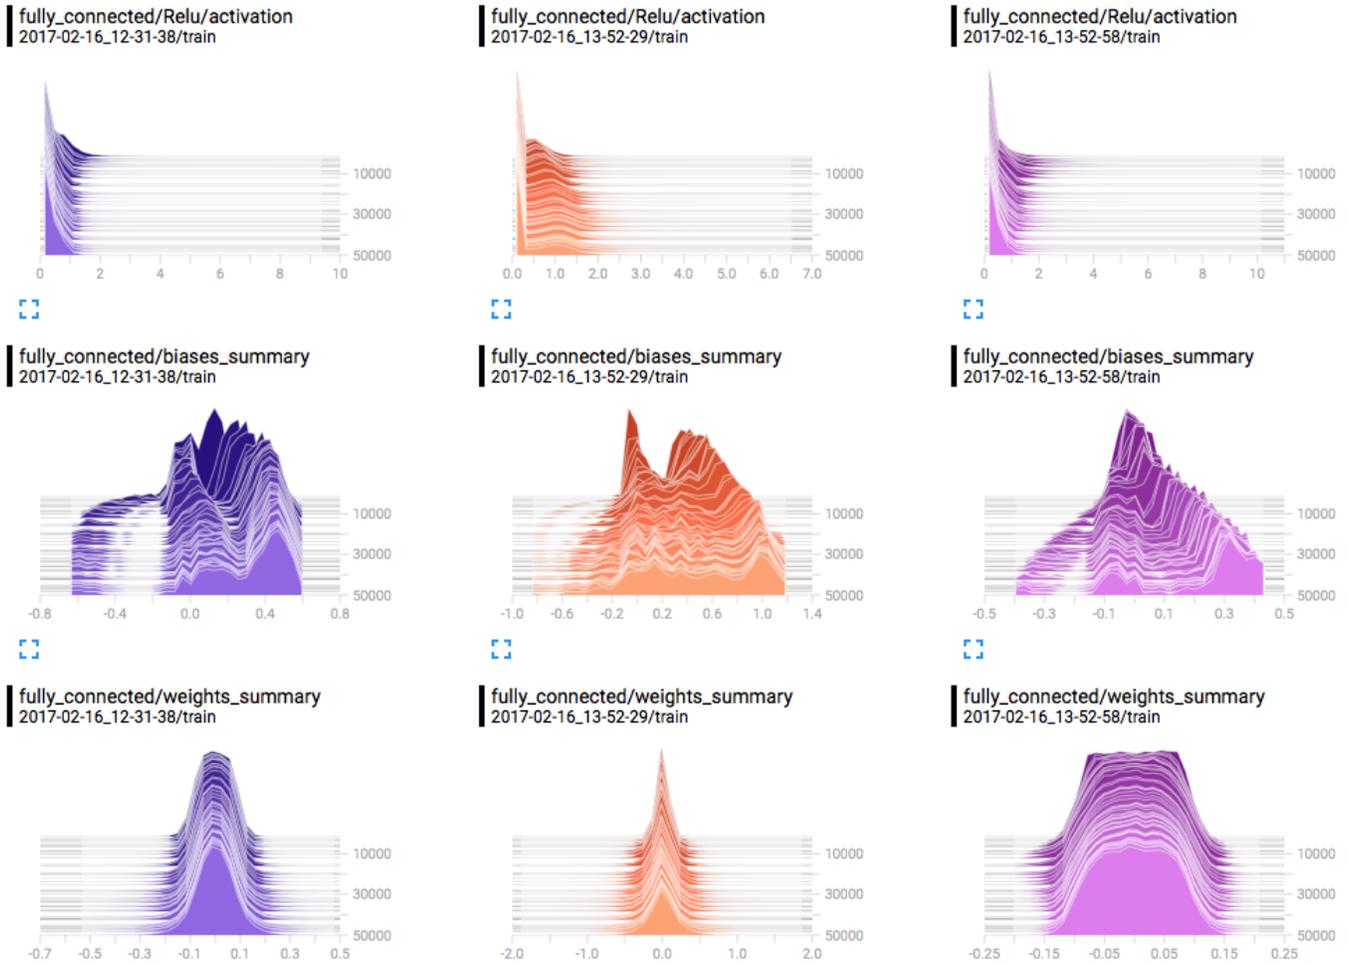


Figure 4: The histograms of various tensors collected during three runs within hidden layer (Left: $\gamma = 0.0001$, middle: $\gamma = 0.001$, right: $\gamma = 0.00001$). By examining weights histograms it may be noticed that middle has widest base with sharp peak which means that layer converged, but search space was widest among all runs. The left one has narrower base and sharp peak, which means that layer converged withing narrower search space and as result has better prediction power. The right one has narrow base but wide plateau at the top, which means that search space is narrow but algorithm still failed to converge. The left and middle histograms has sharp peaks compared to right one, which may be a signal that their learning rate values has more relevance for algorithm convergence and as result we have better predictions for those learning rates.

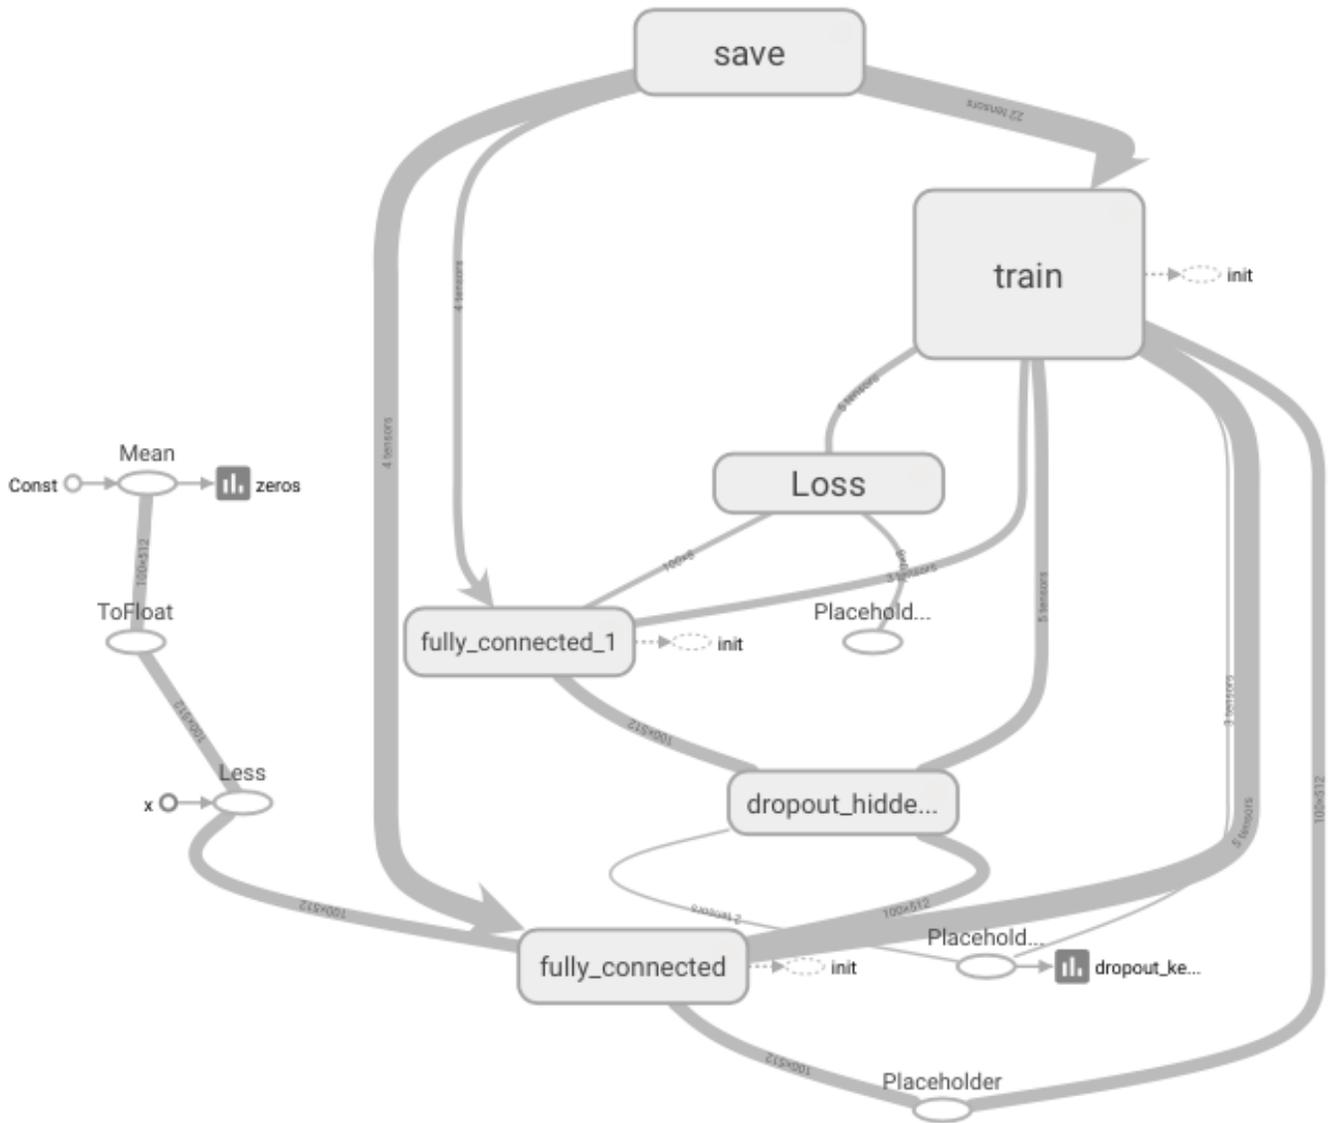


Figure 5: The tensor network graph for multi layer perceptron with one hidden layer (fully_connected), and one linear output layer (fully_connected1). The input layer presented as input tensor placeholder. The hidden layer has ReLU activation nonlinearity. The loss function is MSE (mean squared error). The train optimizer is Adam (for Adaptive Moment Estimation).

B. DEEP NN EVALUATION PLOTS

The following pages provide plots and diagrams related to evaluation of studied deep neural networks.

train/Loss/total/Mean

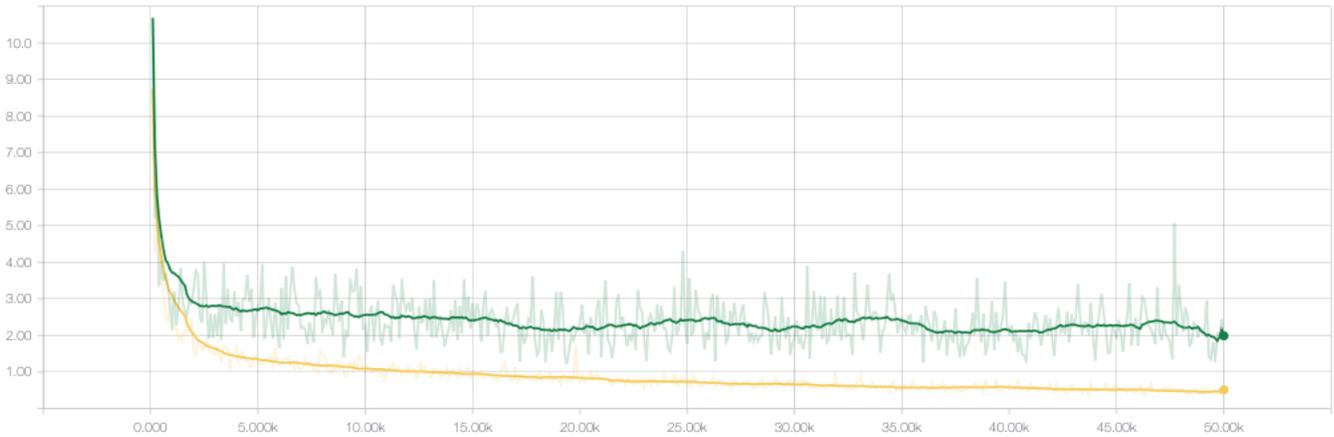


Figure 6: The loss function plot against *train* and *validation* data sets, for $K = 1024$ input features and three hidden layers - [2048,1024,1024] units correspondingly. It can be seen that model slightly overfitted against training data - the validation plot is above train plot and doesn't improve with more training steps.

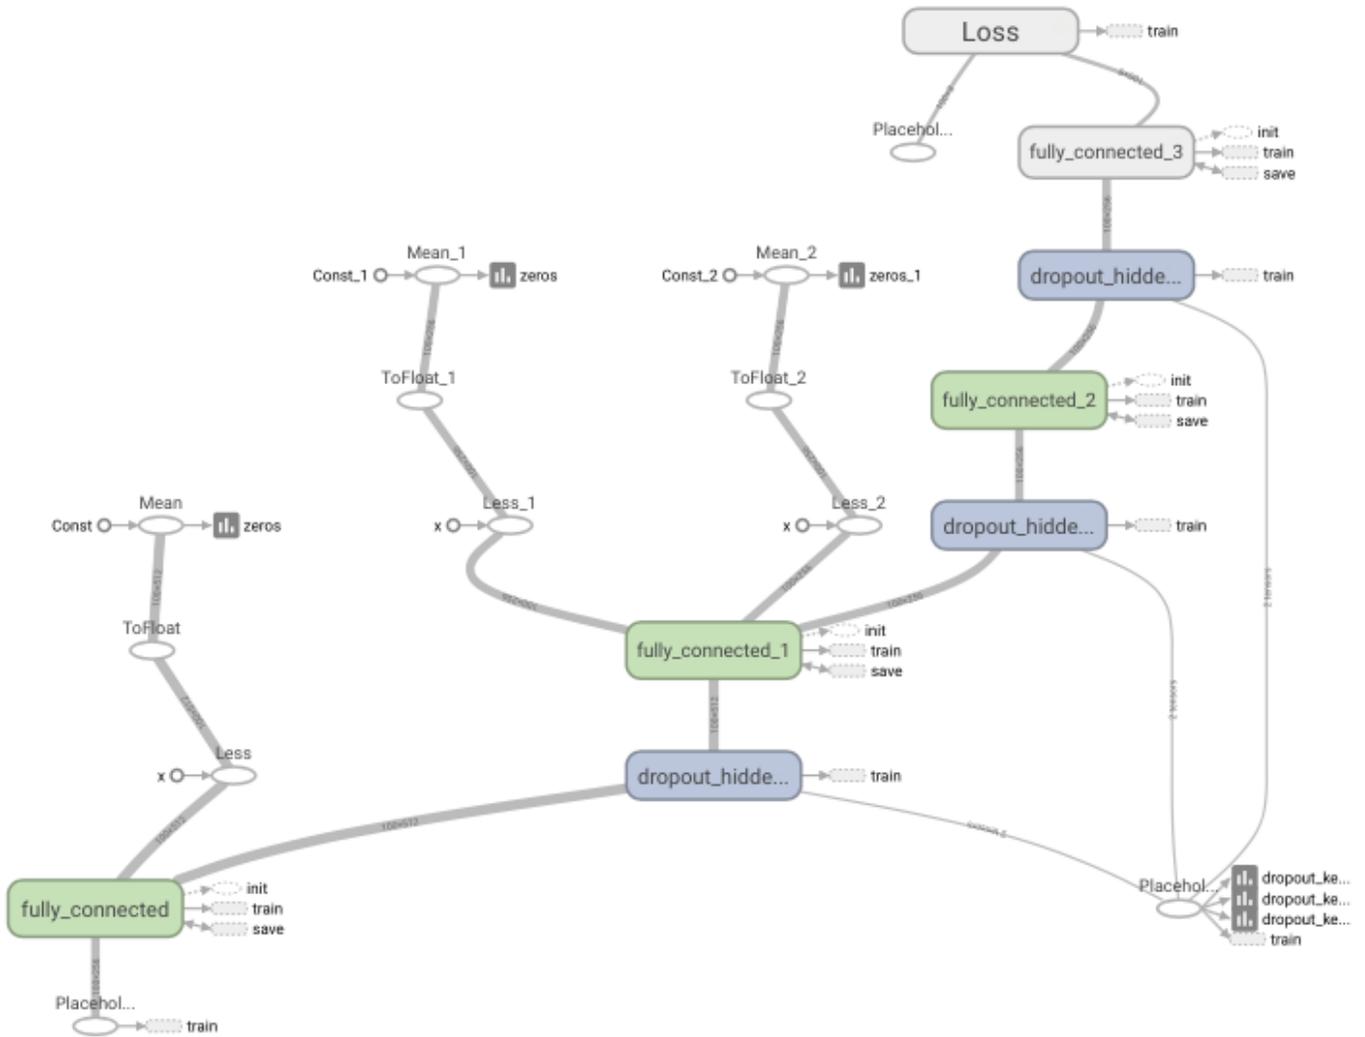


Figure 8: The tensor network graph for DNN with three hidden layers. The input layer presented as input tensor placeholder. The hidden layers has ReLU activation nonlinearity. The loss function is MSE (mean squared error). The train optimizer is Adam (for Adaptive Moment Estimation).

REFERENCES

- [Kosinski et. al, 2016] Michal Kosinski, Yilun Wang, Himabindu Lakkaraju, and Jure Leskovec, (2016). Mining Big Data to Extract Patterns and Predict Real-Life Outcomes. *Psychological Methods* 2016, Vol. 21, No. 4, 493-506. DOI: 10.1037/met0000105
- [Lambiotte, R., and Kosinski, M., 2014] Lambiotte, R., and Kosinski, M., (2014). Tracking the digital footprints of personality. *Proceedings of the Institute of Electrical and Electronics Engineers*, 102, 1934-1939. DOI: 10.1109/JPROC.2014.2359054
- [Goldberg et. al, 2006] Goldberg, L. R., Johnson, J. A., Eber, H. W., Hogan, R., Ashton, M. C., Cloninger, C. R., and Gough, H. G. (2006). The International Personality Item Pool and the future of public-domain personality measures. *Journal of Research in Personality*, 40, 84-96. DOI: 10.1016/j.jrp.2005.08.007
- [Golub, G. H., and Reinsch, C. 1970] Golub, G. H., and Reinsch, C. (1970). Singular value decomposition and least squares solutions. *Numerische Mathematik*, 14, 403-420. DOI: 10.1007/BF02163027
- [Abdi, H., 2003] Abdi, H. (2003). Factor rotations in factor analyses. In M. Lewis-Beck, A. E. Bryman, & T. F. Liao (Eds.), *The SAGE encyclopedia of social science research methods* (pp. 792-795). Thousand Oaks, CA: SAGE.
- [Goodfellow et al., 2016] Ian Goodfellow and Yoshua Bengio and Aaron Courville, (2016). *Deep learning*. Manuscript in preparation. Retrieved from <http://www.deeplearningbook.org/>
- [Daphne Koller, 2012] Daphne Koller, (2010-2012) *Probabilistic Graphical Models*. Stanford University. Retrieved from <http://openclassroom.stanford.edu/MainFolder/CoursePage.php?course=ProbabilisticGraphicalModels>
- [Cortes & Vapnik, 1995] Cortes, C. & Vapnik, V. (1995). Support-vector networks. *Machine Learning*. 20 (3):273-297. DOI: 10.1007/BF00994018
- [Yan, Su, 2009] Xin Yan, Xiao Gang Su, (2009), *Linear Regression Analysis: Theory and Computing*, World Scientific, pp. 1-2, ISBN 9789812834119
- [David A. Freedman, 2009] David A. Freedman (2009). *Statistical Models: Theory and Practice*. Cambridge University Press, p. 26.
- [Hilary L. Seal, 1967] Hilary L. Seal (1967). The historical development of the Gauss linear model. *Biometrika*. 54 (1/2): 1-24. DOI: 10.1093/biomet/54.1-2.1
- [Rodriguez, G., 2007] Rodriguez, G. (2007). *Lecture Notes on Generalized Linear Models*. pp. Chapter 3, page 45. Retrieved from <http://data.princeton.edu/wws509/notes/>
- [Sing, Sander, Beerenwinkel, Lengauer, 2005] Sing, T., Sander, O., Beerenwinkel, N., & Lengauer, T. (2005). ROCR: Visualizing classifier performance in R. *Bioinformatics*, 21, 3940-3941. DOI: 10.1093/bioinformatics/bti623
- [Gain, 1951] Gain, A. K. (1951). The frequency distribution of the product moment correlation coefficient in random samples of any size draw from non-normal universes. *Biometrika*. 38: 219-247. DOI: 10.1093/biomet/38.1-2.219
- [Kohavi, Ron, 1995] Kohavi, Ron (1995). A study of cross-validation and bootstrap for accuracy estimation and model selection. *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*. San Mateo, CA: Morgan Kaufmann. 2 (12): 1137-1143. CiteSeerX 10.1.1.48.529
- [Zhang, Marron, Shen, & Zhu, 2007] Zhang, L., Marron, J., Shen, H., & Zhu, Z. (2007). Singular value decomposition and its visualization. *Journal of Computational and Graphical Statistics*, 16, 833-854.

- [van Buuren, Groothuis-Oudshoorn, 2011] Stef van Buuren and Karin Groothuis-Oudshoorn (2011). mice: Multivariate Imputation by Chained Equations in R. *Journal of Statistical Software* 45 (3) American Statistical Association. Retrieved from <http://doc.utwente.nl/78938/>
- [Rubin DB, 1987] Rubin DB (1987). *Multiple Imputation for Nonresponse in Surveys*. John Wiley & Sons, New York.
- [Christopher M. Bishop, 1995] Christopher M. Bishop (1995). *Neural Networks for Pattern Recognition*, Oxford University Press, Inc. New York, NY, USA ©1995 ISBN:0198538642
- [Bengio, Yoshua, 2009] Bengio, Yoshua (2009). Learning Deep Architectures for AI. *Foundations and Trends in Machine Learning*. 2 (1): 1-127. DOI: 10.1561/22000000006
- [LeCun, Bengio, Hinton, 2015] LeCun, Yann; Bengio, Yoshua; Hinton, Geoffrey (2015). Deep learning. *Nature*. 521: 436-444. DOI:10.1038/nature14539
- [Tetko, Livingstone, Luik, 1995] Tetko, I. V.; Livingstone, D. J.; Luik, A. I. (1995). Neural network studies. 1. Comparison of Overfitting and Overtraining. *J. Chem. Inf. Comput. Sci.* 35 (5): 826-833. DOI: 10.1021/ci00027a006
- [Hinton, G. et al., 2012] Hinton, Geoffrey E.; Srivastava, Nitish; Krizhevsky, Alex; Sutskever, Ilya; Salakhutdinov, Ruslan R. (2012). *Improving neural networks by preventing co-adaptation of feature detectors*. arXiv preprint arXiv:1207.0580v1
- [Nair, Hinton, 2010] Vinod Nair and Geoffrey Hinton (2010). Rectified linear units improve restricted Boltzmann machines. *ICML*. Retrieved from http://machinelearning.wustl.edu/mlpapers/paper_files/icml2010_NairH10.pdf
- [Glorot, Bordes, Bengio, 2011] Xavier Glorot, Antoine Bordes, Yoshua Bengio (2011). Deep Sparse Rectifier Neural Networks. *JMLR W&CP* 15:315-323 Retrieved from <http://jmlr.org/proceedings/papers/v15/glorot11a.html>
- [Kingma, Ba, 2014] Diederik P. Kingma; Lei Jimmy Ba (2014). *Adam: A method for stochastic optimization*. arXiv preprint arXiv:1412.6980
- [Ba, Caruana, 2014] Lei Jimmy Ba, Rich Caruana (2014). *Do Deep Nets Really Need to be Deep?* arXiv preprint arXiv:1312.6184
- [Kirkpatrick et al., 1983] S. Kirkpatrick, C. D. Gelatt Jr., M. P. Vecchi (1983). Optimization by Simulated Annealing. *Science*, 13 May 1983: Vol. 220, Issue 4598, pp. 671-680 DOI: 10.1126/science.220.4598.671
- [R Core Team, 2015] R Core Team, (2015). R: A language and environment for statistical computing. *Vienna, Austria: R Foundation for Statistical Computing*. Retrieved from <http://www.R-project.org/>
- [Google Brain Team, 2015] Google Brain Team, (2015). TensorFlow is an open source software library for numerical computation using data flow graphs. Retrieved from <https://www.tensorflow.org>