



HAL
open science

Disease Classification in Metagenomics with 2D Embeddings and Deep Learning

Thanh Hai Nguyen, Edi Prifti, Yann Chevaleyre, Nataliya Sokolovska,
Jean-Daniel Zucker

► **To cite this version:**

Thanh Hai Nguyen, Edi Prifti, Yann Chevaleyre, Nataliya Sokolovska, Jean-Daniel Zucker. Disease Classification in Metagenomics with 2D Embeddings and Deep Learning. La Conférence sur l'Apprentissage automatique (CAp), Jun 2018, Rouen, France. hal-01819205

HAL Id: hal-01819205

<https://hal.sorbonne-universite.fr/hal-01819205>

Submitted on 20 Jun 2018

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.

Disease Classification in Metagenomics with 2D Embeddings and Deep Learning

Thanh Hai Nguyen ^{*1}, Edi Prifti¹, Yann Chevalere³, Nataliya Sokolovska⁴, et Jean-Daniel Zucker²

¹Integromics, Institute of Cardiometabolism and Nutrition. Paris, France

²Ummisco, IRD, Sorbonne Université, Bondy, France

³LAMSADE, Université Paris Dauphine, Paris, France

⁴Nutriomics, Sorbonne Université, INSERM, Hôpital la Pitié Salpêtrière, Paris, France

June 8, 2018

Abstract

Deep learning (DL) techniques have shown unprecedented success when applied to images, waveforms, and text. Generally, when the sample size (N) is much bigger than the number of features (d), DL often outperforms other machine learning (ML) techniques, often through the use of Convolutional Neural Networks (CNNs). However, in many bioinformatics fields (including metagenomics), we encounter the opposite situation where d is significantly greater than N . In these situations, applying DL techniques would lead to severe over-fitting.

Here we aim to improve classification of various diseases with metagenomic data through the use of CNNs. For this we proposed to represent metagenomic data as images. The proposed **Met2Img** approach relies on taxonomic and t-SNE embeddings to transform abundance data into "synthetic images".

We applied our approach to twelve benchmark data sets including more than 1400 metagenomic samples. Our results show significant improvements over the state-of-the-art algorithms (Random Forest (RF), Support Vector Machine (SVM)). We observe that the integration of phylogenetic information alongside abundance data improves classification. The proposed approach is not only important in classification setting but also allows to visualize complex metagenomic data. The Met2Img is implemented in Python.

Keywords: Machine Learning for Health, classification, metagenomics, deep learning, visualization.

*nthai@cit.ctu.edu.vn

1 Introduction

High throughput data acquisition in the biomedical field has revolutionized research and applications in medicine and biotechnology. Also known as "omics" data, they reflect different aspects of systems biology (genomics, transcriptomics, metabolomics, proteomics, etc.) but also whole biological ecosystems acquired through metagenomics. There is an increasing number of datasets which are publicly available. Different statistical methods have been applied to classify patients from controls [GW09] and some have also performed meta-analyses on multiple datasets [PTM⁺16]. However, exploring omics data is challenging, since the number of features d is very large, and the number of observations N is small. Up to now, the most successful techniques applied to omics datasets have been mainly Random Forest (RF), and sparse regression.

In this paper, we applied DL on six metagenomic datasets, which reflect bacterial species abundance and presence in the gut of diseased patients and healthy controls. We also evaluate our method on additional datasets with *genus* abundance. Since DL performs particularly well for image classification, we focused in this work in the use of CNNs applied to images. For this purpose, we first searched for ways to represent metagenomic data as "synthetic" images and second applied CNN techniques to learn representations of the data that could be used next for classification purposes.

There are numerous studies where machine learning (ML) is applied to analyze large metagenomic datasets. [PTM⁺16] proposed a unified methodology, MetAML, to compare different state-of-the-art methods (SVM,

RF, etc.) on various metagenomic datasets. These datasets were processed using the same bioinformatics pipeline for comparative purposes. A general overview of ML tools for metagenomics is provided by [SN17]. In [DPR15], a typical pipeline for metagenomic analysis was suggested. It includes data preprocessing, feature extraction, followed by classification or clustering. In order to extract relevant feature, raw data can be arbitrarily transformed. A new shape of data is supposed to help predictors to predict more accurately. In some cases, if domain knowledge is available, it is possible to introduce this prior knowledge into the learning procedure. It can be a heuristic and/or hand picking to select a subset of most relevant features. [DPR15] proposed to use deep belief networks and recursive neural networks for metagenomics.

A metagenomic dataset usually contains bacterial DNA sequences obtained from an ecosystem such as the intestinal microbiome for instance. These sequences are usually transformed using bioinformatics pipelines into a vector representing the abundance of a set of features (species or other taxonomic groups). One approach of conversion from sequences to vectors removed short and low quality score reads from the whole collection of sequences. Then, the remainder is clustered with CD-HIT [LG06], for example into non redundant representative sequences - usually called a catalog. Reads are typically aligned to the catalog and counted. Annotation algorithms are applied to the catalog allowing to infer abundance on different taxonomic levels. Other groups have also tackled the use of DL in classifying metagenomic data. For instance, some authors introduced approaches to design a tree like structure for data sets of metagenomics learned by neural networks. A study in [DCN+18] identified a group of bacteria that is consistently related to Colorectal Cancer (COL) and revealed potential diagnosis of this disease across multiple populations and classified COL cases with a SVM model using the seven COL-enriched bacterial species. The authors in [FGM+18] introduced a novel DL approach, Ph-CNN, for classification of metagenomics using hierarchical structure of Operational Taxonomic Unit (OTU). The performance of Ph-CNN is promising compared to SVM, RF and a fully connected neural network.

Image classification based on ML has obtained great achievements and there are numerous studies attempting to propose architectures to improve the performance. Alex et al. in [PBBW12] introduced a deep and big CNN (including 60 million parameters) which performed on a very large dataset with 1.2 million color images of 224x224 in ImageNet LSVRC-2010.

The architectures consisted of 5 convolutional and 3 fully connected layers. The network achieved top-1 and top-5 error rates of 37.5% and 17% , respectively. The authors in [ZF14] presented a novel technique (ZFNet) for visualizing feature maps through convolutional layers, and investigated how to improve the performance of Convolutional Neural Networks (convnet). A team at Google in 2014 proposed GoogLeNet [SLaPS+14], a very deep architecture of CNN with 22 layers that won the ImageNet Large-Scale Visual Recognition Challenge 2014 (ILSVRC14). In [SZ14], the authors presented how depth affects the performance of CNNs (VGGNet). The authors designed deep architectures of CNNs with very small convolutions (3x3) and achieved top-1 validation error and top-5 validation error were 23.7% and 6.8% , respectively. The authors in [HZRS15] introduced a residual learning framework (ResNet) that was up to 152 layers but being less complex. ResNet won the 1st places on the ILSVRC 2015 and COCO 2015 competitions.

Here, we present the **Met2Img** framework based on *Fill-up* and *t-SNE* to visualize features into images. Our objectives were to propose efficient representations which produce compact and informative images, and to prove DL techniques as efficient tools for prediction tasks in the context of metagenomics. Our contribution is multi-fold:

- We propose a *visualization approach for metagenomic data* which shows significant results on 4 out of 6 datasets (Cirrhosis, Inflammatory Bowel Disease (IBD), Obesity, Type 2 Diabetes) compared to MetAML [PTM+16].
- CNNs outperform standard shallow learning algorithms such as RF and SVM. This proves that deep learning is a promising ML technique for metagenomics.
- We illustrate that the proposed method not only performs competitively on species abundance data but also shows significant improvements compared to the state-of-the-art, Ph-CNN, on genus-level taxonomy abundance with six various classification tasks on IBD dataset published in [FGM+18].

The paper is organized as follows. In Section 2 and 3, we describe the benchmarks and methods to visualize feature abundance with the Fill-up and t-SNE approaches. We show our results in Section 4 for different deep learning architectures, for 1D and 2D representations. At the end of this section, we illustrate the performance of our approach on another additional

dataset consisting of six classification tasks on IBD with genus-level abundance. Conclusion is presented in the Section 5.

2 Metagenomic data benchmarks

We work with metagenomic abundance data, i.e. data indicating how present (or absent) is an OTU (Operational taxonomic unit) in human gut. We evaluated our method with each of the different visual representations on twelve different datasets (see Table 1) with two groups (A and B). Group A consists of datasets including bacterial species related to various diseases, namely: liver cirrhosis (CIR), colorectal cancer (COL), obesity (OBE), inflammatory bowel disease (IBD) and Type 2 Diabetes (T2D) [PTM⁺16], [QYL⁺14], [ZTV⁺14], [LCNQ⁺13], [QLR⁺10], [QLC⁺12], [KTN⁺13]], with CIR (n=232 samples with 118 patients), COL (n=48 patients and n=73 healthy individuals), OBE (n=89 non-obese and n=164 obese individuals), IBD (n=110 samples of which 25 were affected by the disease) and T2D (n=344 individuals of which n=170 are T2D patients). In addition, one dataset, namely WT2, includes 96 European women with n=53 T2D patients and n=43 healthy individuals. The abundance datasets are transformed to obtain another representation based on feature presence when the abundance is greater than zero. These data were obtained using the default parameters of MetaPhlan2 [TFT⁺15] as detailed in Pasolli et al. [PTM⁺16]. Group B includes Sokol’s lab data [SLA⁺16] consisting of microbiome information of 38 healthy subjects (HS) and 222 IBD patients. The bacterial abundance includes 306 OTUs with genus level abundance. Patients in this data are classified into two categories according to the disease phenotype Ulcerative colitis (UC) and Crohn’s disease (CD). Each category is divided into two conditions (flare (f), if patients got worse or reappeared the symptoms and remission (r), if patients’ symptoms are decreased or disappear). The dataset was partitioned into subset with ileal Crohn’s disease (iCD) and colon Crohn’s disease (cCD). The detail description of the data was presented in [FGM⁺18].

For each sample, species/genus abundance is a relative proportion and is represented as a real number - the total abundance of all species/genus sums to 1.

3 Methodology

Our approach consists of the following steps: First, a set of colors is chosen and applied to different binning approaches (see 3.1). The binning can be performed on a logarithmic scale or a transformation. Then, the features are visualized into images by one of two different ways (phylogenetic-sorting (using **Fill-up**) or visualized based on t-Distributed Stochastic Neighbor Embedding (t-SNE) [MH08] (see 3.2). t-SNE technique is useful to find faithful representations for high-dimensional points visualized in a more compact space, typically the 2D plane. For phylogenetic-sorting, the features which are bacterial species are arranged based on their taxonomic annotation ordered alphabetically by concatenating the strings of their taxonomy (i.e. phylum, class, order, family, genus and species). This ordering of the variables embeds into the image external biological knowledge, which reflects the evolutionary relationship between these species. Each visualization method will be used to either represent abundance or presence data. The last representation, which serves as control is the 1D of the raw data (with the species also sorted phylogenetically). For the t-SNE representation, we use only training sets to generate global t-SNE maps, images of training and test set are created from these global maps.

3.1 Abundance Bins for metagenomic synthetic images

In order to discretize abundances as colors in the images with, we use different methods of binning. Each bin is illustrated by a distinct color extracted from color strip of heatmap colormaps in Python library such as *jet*, *rainbow*, *viridis*. In [LH18], authors stated that *viridis* showed a good performance in terms of time and error. The binning method we used in the project is Unsupervised Binning which does not use the target (class) information. In this part, we use **Equal Width binning** (EQW) with ranging [**Min**, **Max**]. We test with $k = 10$ bins (for color distinct images, and gray images), width of intervals is $w = 0.1$, if $\text{Min}=0$ and $\text{Max} = 1$, for example.

3.1.1 Binning based on abundance distribution

In Fig. 1, *left* shows the histogram of the original data. We notice the Zero-Inflated distribution as typically metagenomic data are very sparse. On the *right* we notice the log-transformed distribution of the data

Group A						
Dataset name	CIR	COL	IBD	OBE	T2D	WT2
#features	542	503	443	465	572	381
#samples	232	121	110	253	344	96
#patients	118	48	25	164	170	53
#controls	114	73	85	89	174	43
ratio of patients	0.51	0.40	0.23	0.65	0.49	0.55
ratio of controls	0.49	0.60	0.77	0.35	0.51	0.45
Autofit size for images	24×24	23×23	22×22	22×22	24×24	20×20
Group B						
Dataset name	CDf	CDr	iCDf	iCDr	UCf	UCr
#features	259	237	247	257	250	237
#samples	98	115	82	97	79	82
#patients	60	77	44	59	41	44
#controls	38	38	38	38	38	38
ratio of patients	0.61	0.67	0.54	0.61	0.52	0.54
ratio of controls	0.39	0.33	0.46	0.39	0.48	0.46
Autofit size for images	17×17	16×16	16×16	17×17	16×16	16×16

Table 1: Information on datasets




Figure 1: Histogram of whole six datasets of group A (Left: original data, Right: Log-histogram (base 4))

(eliminated values of 0) computed by Logarithm (base 4) which is more normally-distributed. In logarithmic scale, the width of each break is 1 being equivalent to a 4-fold increase from the previous bin. As observed from Fig. 2, histograms of six datasets of group A with Logarithmic scale (base 4) share the same distributions. From our observations, we propose a hypothesis that the models will perform better with such breaks owning values of breaks from $0, 10^{-7}, 4 * 10^{-7}, 1.6 * 10^{-6}, \dots, 0.0065536, 1$. The first break is from 0 to 10^{-7} which is the minimum value of species abundance known in 6 datasets of group A, each multiplies four times preceding one. We called this binning “SPecies Bins” (SPB) in experiments.

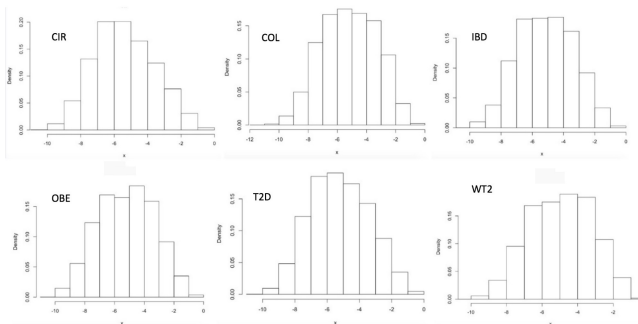


Figure 2: Log histogram of each dataset in group A

3.1.2 Binning based on Quantile Transformation (QTF)

We proposed another approach to bin the data, based on a scaling factor which is learned in the training set and then applied to the test set. With different distributions of data, standardization is a commonly-used technique for numerous ML algorithms. Quantile TransFormation (QTF), a Non-Linear transformation, is considered as a strong preprocessing technique because of reducing the outliers effect. Values in new/unseen data (for example, test/validation set) which are lower or higher the fitted range will be set to the bounds of the output distribution. In the experiments, we use this transformation to fit the fea-




Figure 3: Examples of Fill-up images, Left-Right (types of colors): heatmap, gray scale and black/white




Figure 4: Examples of t-SNE representation, Left-Right (types of images): color images (using heatmap), gray images, black/white images; Top-Down: Global t-SNE maps and images of samples created from the global t-SNE map.

tures’ signal to a uniform distribution. These implementations are provided from the scikit-learn library in Python [GM13].

3.1.3 Binary Bins

Beside the previous two methods described above, we also use the binary bins to indicate PResence/absence (PR), respectively corresponding to black/white (BW) values in subsequent images. Note that this method is equivalent to one-hot encoding.

3.2 Generation of artificial metagenomic images: Fill-up and t-SNE

Fill-up: images are created by arranging abundance/presence values into a matrix in a right-to-left order by row top-to-bottom. The image is square and empty bottom-left of the image are set to zero (white). As an example for a dataset containing 542 features (i.e. bacterial species) in the cirrhosis dataset, we need a matrix of 24×24 to fill up 542 values of species into

this square. The first row of pixel is arranged from the first species to the 24th species, the second row includes from the 25th to the 48th and so on till the end. We use distinct colors in binning scale with SPB, QTF and PR to illustrate abundance values of species and black and white for presence/ absence, where white represents absent values (see examples at Fig. 3)

T-SNE maps: are built based on training sets from raw data with perplexity $p=5$, learning rate = 200 (default), and the number of iterations set to 300. These maps are used to generate images for training and testing sets. Each species is considered as a point in the map and only species that are present are showed either in abundance or presence using the same color schemes as above (Fig. 4).

3.3 Convolutional Neural Network architectures and models used in the experiments

One-dimensional case

In order to predict disease using 1D data, we use a Fully Connected neural network (FC) and one 1D convolutional neural network (CNN1D). FC model includes one fully-connected layer and gives one output. This a very simple model but the performance is relatively high. The structure of this network contains a fully connected layer with sigmoid function. CNN1D includes one 1D convolutional layer with 64 filters and one max pooling of 2. We also use classical learning algorithms such RF (50 trees in the forest) and SVM (kernels of Sigmoid, Radial, Linear) for 1D data.

Two-dimensional case

Images in Fill-up vary in size from 16×16 to 24×24 depending on the number of features while, in t-SNE, we use 24×24 images for all datasets. The network receives either color images with three channels or gray images with one channel, then passing through a stack of one convolutional layers with 64 kernels of 3×3 (stride 1), followed by a max pooling 2×2 (stride 2). ReLU is used after each convolution. Wide convolution is implemented to preserve the dimension of input after passing through the convolutional layer. There are two approaches for binary classification including using either two output neurons (two-node technique) or one output neuron (one-node technique). We use the latter with a sigmoid activation function at the final layer (see the CNN architecture for images of 24×24 in Fig. 5).




Figure 5: The CNN architecture includes a stack of one convolutional layer with 64 filters of 3x3 and a max pooling of 2x2, followed by one fully connected layer.

Experimental Setup

All networks use the same optimization function (Adam [KB14]), learning rate = 0.001, batch size = 16, loss function: binary cross entropy, epoch = 500 (using Early Stopping to avoid over-fitting with the number of epoch patience of 5). **Met2Img** is implemented in Keras 2.1.3, Tensorflow 1.5 using Python/Scikit-Learn[GM13] and can run either in CPU or GPU architectures.

4 Results

4.1 Comparing to the-state-of-the-art (MetAML) [PTM⁺16]

We run a one-tailed t-test to find significant improvements compared to the-state-of-the-art. The p -values are computed based on function `tsum.test` (**PASWR**-package for Probability and Statistics with R) to compare results from MetAML. The results which have p -values < 0.05 are significant improvements. We use accuracy (ACC) to measure model performances. Classification accuracies were assessed by 10-fold cross validation, repeated and averaged on 10 independent runs using stratified cross validation approach.

4.1.1 The results of 1D data

As shown in the Table 2, RF both in MetAML and our framework outperforms other models. In 1D data, CNN1D gives better FC while results for the SVM models are the worst. In SVM models, Linear kernel shows the best performance. From these results, standard and shallow learning algorithms (such as RF) are more robust than deep learning for 1D data.

4.1.2 The results of 2D data

The performances of MetAML is shown in Table 3 are results of RF model, the best model in MetAML. The results with QTF bins share the same pattern with “Species bins” (SPB), but the performance of QTF is slightly lower than **SPB**. Noteworthy, t-SNE using QTF gives significant results for T2D while Fill-up shows poor performance on this dataset. The performance of Fill-up is better compared to t-SNE. In addition, t-SNE performs worse due to many overlapped points, while in contrast every features in Fill-up is visible. As shown in Fig. 6, the CNN model achieves either significant results with CIR, IBD, OBE (p -value < 0.0005) or comparative performance to the others including COL, T2D, WT2 compared to the best model (RF) in MetAML. This proves that CNN is a promising approach for metagenomic data.




Figure 6: Performance comparison in ACC between CNN model (using Fill-up with SPB and gray images) and shallow learning algorithm (RF) of MetAML and our framework. The standard deviations are shown in error bars.

4.2 Applying Met2Img on Sokol’s lab data [SLA⁺16]

In this section, we evaluate the performance of Met2Img on Sokol’s lab data [SLA⁺16], which used to evaluate Ph-CNN in [FGM⁺18]. The authors performed the six classification tasks on this data to classify HS versus the six partitions Cdf, CDr, iCDF, iCDr, UCf, UCr. They proposed Ph-CNN considered as a

Framework	Model	CIR	COL	IBD	OBE	T2D	WT2	AVG
MetAML	RF	0.877	0.805	0.809	0.644	0.664	0.703	0.750
	SVM	0.834	0.743	0.809	0.636	0.613	0.596	0.705
Met2Img	RF	0.877	0.812	0.808	0.645	0.672	0.703	0.753
	SVM- Sigmoid	0.509	0.603	0.775	0.648	0.515	0.553	0.600
	SVM- Radial	0.529	0.603	0.775	0.648	0.593	0.553	0.617
	SVM- Linear	0.766	0.666	0.792	0.612	0.634	0.676	0.691
	FC	0.776	0.685	0.775	0.656	0.665	0.607	0.694
	CNN1D	0.775	0.722	0.842	0.663	0.668	0.618	0.715

Table 2: Performance (in ACC) comparison on 1D data. The significant results are reported in **bold**. The average performance of six datasets is computed and shown in the last column (AVG)

novel DL architecture using phylogenetic structure of metagenomics for classification tasks. The structure of the model includes a stack of two 1D convolutional layers of 4x16 and a max pooling of 2x1, followed by a fully connected layer with 64 neurons and a dropout of 0.25. In order to compare to Ph-CNN, we use the same experimental procedure with 5-fold stratified cross validation repeated 10 times (internal validation) and compute Matthews Correlation Coefficient (MCC) as a performance measurement in [FGM⁺18], then applying the best model to the external validation datasets. Internal and external validation sets evaluated in our experiments are the same as in [FGM⁺18].

The performance in Tab. 4 and Fig. 7 using Fill-up with heatmap color images shows significant results compared to Ph-CNN. Especially, Fill-up using QTF outperforms Ph-CNN in internal validation and get better performance in 4 and 5 out of 6 external validation sets. SPB is conducted from species abundance distribution also shows promising results with 4-5 significant results on internal validation sets and achieves greater performances in 4 external validation sets. Noteworthy, for QTF, although CNN model shows worse performance compared to FC, this model reveals encouraging results with better performance on 5/6 external validation sets.

5 Conclusion

In this work, we introduced a novel approach **Met2Img** to learn and classify complex metagenomic data using visual representations constructed with Fill-up and t-SNE. We explore three methods of binning, namely **SPB**, **QTF** and **PR** with various colormaps using heatmap, gray, and black/white images. The bins and images approach are constructed and produced using training data only, thus avoiding overfitting issues.




Figure 7: Performance comparison in MCC between our approach and Ph-CNN on external validation sets.

Our results illustrate that the Fill-up approach outperforms the t-SNE. This may be due to several factors. First, the features in the Fill-up are all visible while features in t-SNE are often overlapping. Second, the Fill-up approach integrates prior knowledge on the phylogenetic classification of the features. In addition, the t-SNE images are more complex than the Fill-up images. Noteworthy, with the Fill-up we show significant improvements three data sets, while the t-SNE reveals significant improvement on one data set, the T2D. The FC model outperforms CNN model in color images while CNN model achieves better performance than FC for gray and black/white images. Besides, the representations based on 2D images yield better results compared to 1D data. In general, the proposed Met2Img method outperforms the state-of-the-art both on species and genus abundance data.

Currently we are investigating various deep learning architectures, and also explore integration of other heterogeneous omics data.

	Bins	Model	Color	CIR	COL	IBD	OBE	T2D	WT2	AVG
	MetAML - RF			0.877	0.805	0.809	0.644	0.664	0.703	0.750
	Met2Img - RF			0.877	0.812	0.808	0.645	0.672	0.703	0.753
Fill-up	PR	CNN	BW	0.880	0.763	0.841	0.669	0.666	0.667	0.748
Fill-up	QTF	CNN	color	0.897	0.781	0.837	0.659	0.664	0.690	0.755
Fill-up	SPB	CNN	gray	0.905	0.793	0.868	0.680	0.651	0.705	0.767
Fill-up	SPB	CNN	color	0.903	0.798	0.863	0.681	0.649	0.713	0.768
Fill-up	PR	FC	BW	0.863	0.735	0.842	0.672	0.656	0.712	0.747
Fill-up	QTF	FC	color	0.887	0.765	0.853	0.657	0.640	0.711	0.752
Fill-up	SPB	FC	grays	0.888	0.772	0.847	0.686	0.652	0.716	0.760
Fill-up	SPB	FC	color	0.905	0.794	0.837	0.679	0.659	0.713	0.764
tsne	PR	CNN	BW	0.862	0.712	0.815	0.664	0.660	0.672	0.731
tsne	QTF	CNN	color	0.878	0.746	0.811	0.658	0.684	0.648	0.737
tsne	QTF	CNN	grays	0.875	0.747	0.809	0.664	0.672	0.651	0.736
tsne	SPB	CNN	grays	0.877	0.761	0.809	0.679	0.664	0.686	0.746
tsne	SPB	CNN	color	0.886	0.769	0.802	0.682	0.658	0.685	0.747
tsne	PR	FC	BW	0.857	0.710	0.813	0.662	0.640	0.688	0.728
tsne	QTF	FC	grays	0.873	0.764	0.787	0.664	0.647	0.686	0.737
tsne	QTF	FC	color	0.889	0.770	0.825	0.660	0.638	0.650	0.739
tsne	SPB	FC	grays	0.871	0.770	0.790	0.678	0.633	0.702	0.741
tsne	SPB	FC	color	0.890	0.778	0.820	0.689	0.615	0.705	0.749

Table 3: Performance (in ACC) comparison on 2D data. The significant results are reported in **bold**. The average performance of six datasets is computed and shown in the last column (AVG)

References

- [DCN⁺18] Zhenwei Dai, Olabisi Oluwabukola Coker, Geicho Nakatsu, William K. K. Wu, Liuyang Zhao, Zigui Chen, Francis K. L. Chan, Karsten Kristiansen, Joseph J. Y. Sung, Sunny Hei Wong, and Jun Yu. Multi-cohort analysis of colorectal cancer metagenome identified altered bacteria across populations and universal bacterial markers. 6(1):70, 2018.
- [DPR15] Gregory Ditzler, Robi Polikar, and Gail Rosen. Multi-Layer and recursive neural networks for metagenomic classification. *IEEE Trans. Nanobioscience*, 14(6):608–616, September 2015.
- [FGM⁺18] Diego Fioravanti, Ylenia Giarratano, Valerio Maggio, Claudio Agostinelli, Marco Chierici, Giuseppe Jurman, and Cesare Furlanello. Phylogenetic convolutional neural networks in metagenomics. 19(2):49, 2018.
- [GM13] Raul Garreta and Guillermo Moncecchi. *Learning scikit-learn: Machine Learning in Python*. Packt Publishing Ltd, November 2013.
- [GW09] Geoffrey S Ginsburg and Huntington F Willard. Genomic and personalized medicine: foundations and applications. *Transl. Res.*, 154(6):277–287, December 2009.
- [HZRS15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *CoRR*, abs/1512.03385, 2015.
- [KB14] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014.
- [KTN⁺13] Fredrik H Karlsson, Valentina Tremaroli, Intawat Nookaew, Göran Bergström, Carl Johan Behre, Björn Fagerberg, Jens Nielsen, and Fredrik Bäckhed. Gut metagenome in european women with normal, impaired and diabetic glucose control. *Nature*, 498(7452):99–103, June 2013.

Model	Bins	Cdf	CDr	iCdf	iCDr	UCf	UCr	AVG
Internal Validation								
Ph-CNN		0.630	0.241	0.704	0.556	0.668	0.464	0.544
CNN	QTF	0.694	0.357	0.746	0.618	0.811	0.467	0.616
CNN	SPB	0.808	0.369	0.862	0.508	0.790	0.610	0.658
FC	QTF	0.743	0.362	0.815	0.642	0.796	0.519	0.646
FC	SPB	0.791	0.292	0.829	0.531	0.795	0.632	0.645
External Validation								
Ph-CNN		0.858	0.853	0.842	0.628	0.741	0.583	0.751
CNN	QTF	0.930	0.868	0.919	0.580	0.826	0.777	0.817
CNN	SPB	1.000	0.270	1.000	0.664	0.840	0.580	0.726
FC	QTF	1.000	0.868	0.842	0.514	0.916	0.713	0.809
FC	SPB	0.854	0.654	1.000	0.669	0.916	0.585	0.780

Table 4: Classification performance (in MCC) compared to Ph-CNN on six classification tasks on IBD. The better results on external validation sets are formatted ***bold-Italic***. The significant results are reported in **bold**. The average performances are revealed in the last column (AVG)

- [LCNQ⁺13] Emmanuelle Le Chatelier, Trine Nielsen, Junjie Qin, Edi Prifti, Falk Hildebrand, Gwen Falony, Mathieu Almeida, Manimozhiyan Arumugam, Jean-Michel Batto, Sean Kennedy, Pierre Leonard, Junhua Li, Kristoffer Burgdorf, Niels Grarup, Torben Jørgensen, Ivan Brandslund, Henrik Bjørn Nielsen, Agnieszka S Juncker, Marcelo Bertalan, Florence Levenez, Nicolas Pons, Simon Rasmussen, Shinichi Sunagawa, Julien Tap, Sebastian Tims, Erwin G Zoetendal, Søren Brunak, Karine Clément, Joël Doré, Michiel Kleerebezem, Karsten Kristiansen, Pierre Renault, Thomas Sicheritz-Ponten, Willem M de Vos, Jean-Daniel Zucker, Jeroen Raes, Torben Hansen, MetaHIT consortium, Peer Bork, Jun Wang, S Dusko Ehrlich, and Oluf Pedersen. Richness of human gut microbiome correlates with metabolic markers. *Nature*, 500(7464):541–546, August 2013.
- [LG06] W Li and A Godzik. Cd-hit: a fast program for clustering and comparing large sets of protein or nucleotide sequences. *Bioinformatics*, 22(13):1658–1659, 2006.
- [LH18] Yang Liu and Jeffrey Heer. Somewhere over the rainbow: An empirical assessment of quantitative colormaps. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI ’18, pages 598:1–598:12. ACM, 2018.
- [MH08] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-SNE. *J. Mach. Learn. Res.*, 9(Nov):2579–2605, 2008.
- [PBBW12] F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors. *ImageNet Classification with Deep Convolutional Neural Networks*. Curran Associates, Inc., 2012.
- [PTM⁺16] Edoardo Pasolli, Duy Tin Truong, Faizan Malik, Levi Waldron, and Nicola Segata. Machine learning meta-analysis of large metagenomic datasets: Tools and biological insights. *PLoS Comput. Biol.*, 12(7):e1004977, July 2016.
- [QLC⁺12] Junjie Qin, Yingrui Li, Zhiming Cai, Shenghui Li, Jianfeng Zhu, Fan Zhang, Suisha Liang, Wenwei Zhang, Yuanlin Guan, Dongqian Shen, Yangqing Peng, Dongya Zhang, Zhuye Jie, Wenxian Wu, Youwen Qin, Wenbin Xue, Junhua Li, Lingchuan Han, Donghui Lu, Peixian Wu, Yali Dai, Xiaojuan Sun, Zesong Li, Aifa Tang, Shilong Zhong, Xiaoping Li, Weineng Chen, Ran Xu, Mingbang Wang, Qiang Feng, Meihua Gong, Jing Yu, Yanyan Zhang, Ming Zhang, Torben Hansen, Gaston Sanchez, Jeroen Raes, Gwen Falony, Shujiro Okuda, Mathieu Almeida, Emmanuelle LeChatelier, Pierre Renault, Nicolas Pons, Jean-

- Michel Batto, Zhaoxi Zhang, Hua Chen, Ruifu Yang, Weimou Zheng, Songgang Li, Huanming Yang, Jian Wang, S Dusko Ehrlich, Rasmus Nielsen, Oluf Pedersen, Karsten Kristiansen, and Jun Wang. A metagenome-wide association study of gut microbiota in type 2 diabetes. *Nature*, 490(7418):55–60, October 2012.
- [QLR⁺10] Junjie Qin, Ruiqiang Li, Jeroen Raes, Manimozhiyan Arumugam, Kristofer Solvsten Burgdorf, Chaysavanh Manichanh, Trine Nielsen, Nicolas Pons, Florence Levenez, Takuji Yamada, Daniel R Mende, Junhua Li, Junming Xu, Shaochuan Li, Dongfang Li, Jianjun Cao, Bo Wang, Huiqing Liang, Huisong Zheng, Yinlong Xie, Julien Tap, Patricia Lepage, Marcelo Bertalan, Jean-Michel Batto, Torben Hansen, Denis Le Paslier, Allan Linneberg, H Bjørn Nielsen, Eric Pelletier, Pierre Renault, Thomas Sicheritz-Ponten, Keith Turner, Hongmei Zhu, Chang Yu, Shengting Li, Min Jian, Yan Zhou, Yingrui Li, Xiuqing Zhang, Songgang Li, Nan Qin, Huanming Yang, Jian Wang, Søren Brunak, Joel Doré, Francisco Guarner, Karsten Kristiansen, Oluf Pedersen, Julian Parkhill, Jean Weissenbach, MetaHIT Consortium, Peer Bork, S Dusko Ehrlich, and Jun Wang. A human gut microbial gene catalogue established by metagenomic sequencing. *Nature*, 464(7285):59–65, March 2010.
- [QYL⁺14] Nan Qin, Fengling Yang, Ang Li, Edi Prifti, Yanfei Chen, Li Shao, Jing Guo, Emmanuelle Le Chatelier, Jian Yao, Lingjiao Wu, Jiawei Zhou, Shujun Ni, Lin Liu, Nicolas Pons, Jean Michel Batto, Sean P Kennedy, Pierre Leonard, Chunhui Yuan, Wenchao Ding, Yuanting Chen, Xinjun Hu, Beiwen Zheng, Guirong Qian, Wei Xu, S Dusko Ehrlich, Shusen Zheng, and Lanjuan Li. Alterations of the human gut microbiome in liver cirrhosis. *Nature*, 513(7516):59–64, September 2014.
- [SLA⁺16] Harry Sokol, Valentin Leducq, Hugues Aschard, Hang-Phuong Pham, Sarah Jegou, Cecilia Landman, David Cohen, Giuseppina Liguori, Anne Bourrier, Isabelle Nion-Larmurier, Jacques Cosnes, Philippe Seksik, Philippe Langella, David Skurnik, Mathias L Richard, and Laurent Beaugerie. Fungal microbiota dysbiosis in IBD. *Gut*, 66(6):1039–1048, 2016.
- [SLaPS⁺14] Christian Szegedy, Wei Liu, Yangqing Jia and VGGNet Pierre Sermanet, Scott E. Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. *CoRR*, abs/1409.4842, 2014.
- [SN17] Hayssam Soueidan and Macha Nikol'ski. Machine learning for metagenomics: methods and tools. *Metagenomics*, 1(1), 2017.
- [SZ14] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. volume abs/1409.1556. 2014.
- [TFT⁺15] Duy Tin Truong, Eric A Franzosa, Timothy L Tickle, Matthias Scholz, George Weingart, Edoardo Pasoli, Adrian Tett, Curtis Huttenhower, and Nicola Segata. MetaPhlan2 for enhanced metagenomic taxonomic profiling. *Nat. Methods*, 12(10):902–903, 2015.
- [ZF14] Matthew D Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In *Lecture Notes in Computer Science*, pages 818–833. 2014.
- [ZTV⁺14] Georg Zeller, Julien Tap, Anita Y Voigt, Shinichi Sunagawa, Jens Roat Kultima, Paul I Costea, Aurélien Amiot, Jürgen Böhm, Francesco Brunetti, Nina Habermann, Rajna Hercog, Moritz Koch, Alain Luciani, Daniel R Mende, Martin A Schneider, Petra Schrotz-King, Christophe Tournigand, Jeanne Tran Van Nhieu, Takuji Yamada, Jürgen Zimmermann, Vladimir Benes, Matthias Kloor, Cornelia M Ulrich, Magnus von Knebel Doeberitz, Iradj Sobhani, and Peer Bork. Potential of fecal microbiota for early-stage detection of colorectal cancer. *Mol. Syst. Biol.*, 10:766, November 2014.