Multiclass Disease Classification from Microbial Whole-Community Metagenomes using Graph Convolutional Neural Networks.

Saad Khan¹ and Libusha Kelly^{1,2†}

1) *Department of Systems & Computational Biology*, 2) *Department of Microbiology & Immunology*
Albert Einstein College of Medicine,
Bronx, NY, USA

[†]*E-mail: libusha.kelly@einstein.yu.edu*

There is a wealth of information contained within one's microbiome regarding their physiology and environment, and this is a promising avenue for developing non-invasive diagnostic tools. Here, we utilize 5643 aggregated, annotated whole-community metagenomes from 19 different diseases to implement the first multiclass microbiome disease classifier of this scale. We compared three different machine learning models: random forests, deep neural nets, and a novel graph convolutional architecture which exploits the graph structure of phylogenetic trees as its input. We show that the graph convolutional model outperforms deep neural nets in terms of accuracy (achieving 75% average test-set accuracy), receiver-operator-characteristics (92.1% average AUC), and precision-recall (50% average AUPR). Additionally, the convolutional net's performance complements that of the random forest, achieving similar accuracy but better receiver-operator-characteristics and lower area under precision-recall. Lastly, we are able to achieve over 90% average top-3 accuracy across all of our models. Together, these results indicate that there are predictive, disease specific signatures across microbiomes which could potentially be used for diagnostic purposes.

Keywords: Microbiome; Machine learning; Metagenomics

1. Introduction

There has been an immense interest in the past few years towards developing statistical methods to predict phenotypes such as disease from metagenomic sequencing of one's microbiome. One of the challenges towards achieving this goal is the problem of separating out signals for different diseases from each other. Unfortunately, many studies that have looked for signatures of individual diseases in the microbiome have produced conflicting results,¹ and there is evidence that there are general signatures of dysbiosis common to all diseases.² Thus, the standard protocol of comparing samples from a disease of interest against healthy controls, with the goal of identifying features of predictive for that disease, may instead be identifying more general features that signal a diseased or healthy microbiome. This problem can arise if the microbiome signals associated with dysbiosis are stronger than those specific to a given disease. The classifier will then have no mechanism to discriminate between general dysbiosis and the specific signatures of the disease. This lack of discrimination is problematic if we want to understand the differences between diseases on a microbial level and/or make diagnoses for specific diseases. For example, in a clinical setting, a classifier for a certain disease that

This is a preprint of an article submitted for consideration in the 2020 Proceedings of the Pacific Symposium on Biocomputing.

has not been trained against a diverse range of conditions may produce false positives when applied on a patient who has a different disease.³

We propose that approaching this problem as a multiclass classification can alleviate this issue by forcing the classifier to find features in the input that are specific for discriminating between a given class and every other class in the dataset. Additionally, this approach allows us to use a larger dataset containing samples from more conditions, potentially alleviating biases due to batch effects between studies. Unfortunately, making accurate predictions becomes much harder in the multiclass setting, because output is much more specific (random guesses are now correct with probability $\frac{1}{n}$ rather than $\frac{1}{2}$). Here, we demonstrate that the multiclass disease classification problem is tractable given the amount of publicly available metagenomic data, we quantify the performance of three machine learning approaches, and we propose graph convolutional network architectures as a powerful method for incorporating microbiome phylogeny for improving performance for deep learning based models. We posit this tractability will only improve over time as more data becomes available.

2. Previous Work

Previous studies have used feedforward and convolutional neural networks (CNN) to perform phenotype classifications on smaller datasets than the one considered here, but none have utilized a graph convolutional architecture like the one we describe. Reiman et al. implemented a CNN by embedding phylogenetic trees into $\mathbb{R}^2$ and used two-dimensional convolutional layers to construct a body-site classifier.⁴ Fioravanti et al. developed a model to diagnose inflammatory bowel disease (IBD) by projecting samples into a two-dimensional space using MultiDimensional Scaling with the patristic distance between phylogenetic trees as the distance metric.⁵ Both papers mapped phylogenetic data to a Euclidean domain to perform convolutions instead of operating in the original tree topology, as we do in this study. Additionally, these studies focused on the disease vs. healthy setting rather than disease vs. all.

More recently, there have been meta-analyses^{2,6,7} that have tried to identify disease-specific signatures which generalize beyond individual studies, but the results presented in these papers have all been for a disease vs. healthy scenario. Our major contribution here is to present three machine learning models, including a novel convolutional architecture which exploits the tree structure of bacterial phylogenies, that can make multiclass disease predictions.

3. Problem Setup

There are many moving targets that make machine learning in bioinformatics particularly challenging, and one major problem is the paucity of standardized datasets. Data pre-processing, particularly in the relatively new microbiome field, involves numerous components that are each active areas of research, and thus being continually improved. In the case of the metagenomics data, pre-processing includes new assembly methods, decontamination algorithms, sequencing libraries, annotation methods, reference genomes and so on.⁸ Additionally, new studies are published every month, resulting in an ever-increasing catalog of potential data points to utilize for model training. The result is that different publications compile their own training datasets by aggregating a selection of studies and running a genomic annota-

tion pipeline to pre-process the raw sequence data. This non-standard pre-processing makes comparing models and results between papers difficult because the input datasets can be fundamentally different across each study. Fortunately, there are now actively-maintained standardized datasets to help with this task. For the purpose of this study, we utilized curatedMetagenomicData,⁹ which contains over 7000 whole-community sequenced (WCS) samples across dozens of diseases and multiple body sites. The metagenomic data is represented as a 12365-dimensional vector of taxonomic relative abundances that were generated using the MetaPhlAn2 software package¹⁰ (this type of sequencing is also much deeper and higher resolution than traditional 16s rRNA-based sequencing methods).

3.1. Dataset Construction

We constructed a dataset containing 5643 samples with 4885 from stool, 403 from skin, 254 from oral cavity, 93 from nares (nasal cavity), and 8 from maternal milk (healthy babies from Asnicar et al.¹¹ by including diseases that had at least 15 unique samples. One of the challenges we faced is that some samples have multiple disease labels due to the way the original studies were run. We approached this problem as a multiclass (one correct label) as opposed to a multilabel (k possible correct labels) problem and thus sought to avoid conflicts due to multiple labeling. Multiple labeling was present in four of our disease sets: Atopic Dermatitis (atopic rhinitis (28), asthma (12)); *C. difficile* (pneumonia (15), cellulitis (2), osteoarthritis (1), ureteral stone (1)); Adenoma (fatty liver (28), hypertension (19), Type II Diabetes (6), Hypercholesterolemia (2), metastases (1)); and Hepatitis B (Schistosoma (1), Hepatitis E (7), Hepatic encephalopathy (2), Hepatitis D (5), Wilson’s disease (1), Cirrhosis (97), Ascites (48)). Additionally, depending on how an individual study was annotated, some samples in the original dataset have no disease label or are instead labeled as ”control”. We chose to be conservative in our construction and only included as ”healthy” samples where the disease column in the curatedMetagenomicData database was ”healthy”.

3.2. Graph Convolutional Neural Networks

There are many architectures for generalizing convolutional neural networks to a graph setting,³² and we chose to use the method outlined by Kipf and Welling,³³ which is computationally simple while remaining robust. The method falls into the category of spectral methods, which model the convolution operation as multiplication by a filter operator g_θ in the Fourier domain against input $x \in \mathbb{R}^n$, where g_θ is a diagonal matrix with parameters $\theta \in \mathbb{R}^n$. The multiplication takes place with respect to the Fourier basis of eigenvectors U of the graph Laplacian $L = I_N - D^{-\frac{1}{2}}AD^{\frac{1}{2}}$, where D is the diagonal matrix of vertex degrees, and A is the adjacency matrix of the graph (in our case, it is the adjacency matrix of the phylogenetic tree used in our model).

$$g_\theta \star x = Ug_\theta U^\top x$$

In Euclidean domains, classical convolution operators are equivalent to this operation, and U is replaced by the Eigenbasis of the Laplacian operator in $\mathbb{R}^n$, which is the standard Fourier basis $\{e^{2\pi i k \cdot x} : k \in \mathbb{N}^n, x \in \mathbb{R}^n\}$.

Disease	Count	Site	Studies
Atopic Dermatitis (AD)	38	Skin	Chng ¹²
Adenoma	143	Stool	Thomas, ⁶ Feng, ¹³ Hannigan, ¹⁴ Zeller ¹⁵
Bronchitis	18	Stool	Yassour ¹⁶
<i>C. difficile</i> Infection (CDI)	33	Stool	Vincent ¹⁷
Colorectal Cancer (CRC)	273	Stool	Vogtmann, ¹ Yu, ¹⁸ Feng, ¹³ Zeller, ¹⁵ Hannigan ¹⁴
Fatty Liver	94	Stool	Loomba, ¹⁹ Feng ¹³
Hepatitis B Virus (HBV)	99	Stool	Qin ²⁰ (2014)
Healthy	3808	All	25 Studies
Hypertension	169	Stool	Thomas, ⁶ Feng, ¹³ Li ²¹
Inflammatory Bowel Disease (IBD)	148	Stool	Nielsen ²²
Impaired Glucose Tolerance (IGT)	49	Stool	Karlsson ²³
Infectious Gastroenteritis	20	Stool	David, ²⁴ Yassour ¹⁶
Metabolic Syndrome	50	Stool	Vrieze ²⁵
Otitis	107	Stool	Yassour ¹⁶
Periodontitis	48	Oral	Shi ²⁶
Psoriasis	74	Skin	TettAJ ²⁷
Rheumatoid Arthritis (RA)	194	Stool	Chengping ²⁸
Type I Diabetes (T1D)	55	Stool	HeintZ-Buschart, ²⁹ Kostic ³⁰
Type II Diabetes (T2D)	223	Stool	Qin ³¹ (2012), Karlsson ²³

Since there are an infinite number of Fourier modes for any domain, we choose a cutoff K and use the first K Laplacian eigenvectors when constructing U . However, this operation is computationally expensive on graphs because there is no general analogue of Fast Fourier Transforms outside of a Euclidean domain (meaning we would have to perform direct matrix multiplications), and the Eigenbases cannot be precomputed because they are different for each graph. The method of Kipf and Welling overcomes this problem by setting $K = 1$ and reformulating the above operation to get

$$g_{\theta} \star x \approx \theta(I_N + D^{-\frac{1}{2}}AD^{\frac{1}{2}})x$$

The choice of cutoff specifies that all nodes up to K edges away from a given node contribute to the output of the convolution operator at that node. Thus, $K = 1$ implies that the output at each node is a function of that node and its immediate neighbors. By stacking layers of this form together (after applying a non-linearity at each level), we can integrate information from increasingly farther nodes. Each individual convolution layer can thus be written as

$$x' = \sigma(g_{\theta} \star x)$$

for some non-linearity σ . Lastly, this framework can be extended to incorporate more than one input / output channel per layer by adding an additional channel dimension to each of these parameters (see original paper for details).

Fig. 1. Architecture of Graph Convolutional Classifier. Purple lines indicate the flow of information from the previous to the highlighted neuron in each layer. In the convolutional layers, each neuron receives information only from its immediate neighbors in the preceding layer.

3.3. Models

We constructed our graph convolutional neural network (GCN) model by stacking two GCN layers with 64 channels each followed by a fully-connected linear layer with 350 nodes. We used exponential linear units (eLU) as our activations between each layer and a sigmoid activation at the top level for classifications.³⁴ Our standard deep neural network (DNN) model consisted of two fully connected linear layers with eLU activation with 1000 and 350 neurons respectively followed by a sigmoid classification layer. Additionally, the GCN model was initialized with the adjacency matrix of the phylogenetic tree corresponding to the taxa present in our 12365-dimensional input vector. All three classifiers took such a 12365-dimensional abundance vector as input. We implemented our neural networks using PyTorch along with the PyTorch Geometric library for the GCN components.^{35,36} Lastly, our random forest model was constructed using the default settings from the `sklearn.ensemble.RandomForestClassifier` module except for the number of trees, which was set to 1000.³⁷ We settled on these configurations after manual experimentation. We found that increasing the number of trees improved generalization performance of the random forest up to a certain point, and the size of hidden layers in our networks did not make a dramatic difference in performance.

3.4. Training

The biggest challenge in this study was dealing with the extreme class imbalance caused by more than half of the samples in our dataset coming from healthy patients and many diseases having only a few dozen samples. We found that any classifier trained directly on this dataset demonstrated extremely poor generalization performance and/or devolved to classifying every sample as the same label (healthy) regardless of input. We overcame this hurdle by employing a combination of class weighting and oversampling. During each experiment, we constructed an oversampled training set by computing the size of the largest class in the training set (healthy) and for every other class, randomly resampling from that class until each class had the same number of samples.³⁸

We augmented resampling by assigning each class a weight of $1 - \frac{1}{n}$, where n is the number of classes, and training our neural networks in one-vs-all manner for each class (using a binary cross-entropy loss function).³⁹ This approach is equivalent to training n different classifiers with shared weights, one for each class, where the objective of each classifier is to discriminate between its class and every other sample in the dataset. To compute the top prediction of the classifier, we then ran the outputs of the network through a softmax function and returned the index of the highest class. This is a commonly used technique in multiclass classification settings, because it reduces the difficult problem of discriminating between n classes to n easier problems of discriminating between 2 classes.⁴⁰ We used a $1 - \frac{1}{n}$ weight in order to address a new class imbalance problem that arises in this setting; if there are m examples of each class in the dataset after oversampling, then each binary classifier will see m positive examples and mn negative examples. A $1 - \frac{1}{n}$ weight magnifies the importance of the positive samples for each class in order to compensate for this. We used a 70-30 split between training and test set, rather than the usual 80-20 in order to have sufficient test samples for each class to evaluate; additionally, we generated a new train-test split for each training run and ran 30 iterations of the GCN model, 30 of the DNN, and 20 of the random forest.

Next, we trained our classifiers. Data preprocessing, ingest, and analysis for our neural networks was performed using Fireworks, a PyTorch-based library that we previously developed to facilitate common machine learning tasks.⁴¹ The GCN and DNN were trained using the Adam optimizer with $2 * 10^{-5}$ and $1 * 10^{-}$ respective learning rates, 40 and 100 respective batch sizes, weight decay parameter set to 1, binary cross-entropy loss, and an early stopping condition when the loss dropped to 2. We performed an analogous procedure with our random forest classifier by assigning the same class weights and using the same over-sampled training set as with the neural network, and we trained the model using its `.fit` method.

4. Results

The average accuracy varied widely by disease and classifier, indicating the difficulty of this problem given the number of positive samples for each category. However, most of the the classifications were statistically better than random (5.3%). In general, the convolutional net performed on par with or better than the deep net in terms of accuracy, and both models had higher accuracy than the random forest for most diseases except for AD, HBV, RA, and hypertension. The random forest excelled in the healthy vs. disease task, producing an

impressive 99% accuracy. Because healthy was also the largest component of the test set, constituting 66.9% of samples on average, this was responsible for the bulk of its performance in terms of weighted accuracy. However, overall accuracy can be biased by the distribution of the test set (ie. if there were more samples for the classes that a given classifier excelled at, then that classifier would benefit). When we weighted each class equally, then the GCN had the highest average accuracy (46% vs 34.6% for random forest and 44.9% for DNN), indicating that it was more accurate across a broader range of diseases than the random forest model.

Additionally, we often want the ability to tune our models in order to have a trade off between sensitivity and specificity. The AUC, or the area under the receiver operating characteristic curve (ROC), is one such measurement which is invariant to label distributions. We generated ROCs for our neural nets by varying the bias threshold of the final layer and evaluating their true positive and false positive rate on the test set, and for the random forest by using the `sklearn.metrics.roc_curve` function on the test set predictions. We found that the GCN model had higher or statistically equivalent AUCs across all labels than both the random forest and the deep net. In particular, we achieved an average AUC of 89.5% for T2D, which previous studies have found to be a particularly challenging task⁷ (AUCs typically range in the 60s).

Disease	Accuracy		
	GCN (n=30)	DNN (n=30)	RF (n=20)
Healthy	.91 ± .01	.87 ± .01	.99 ± 0
CRC	.36 ± .02	.35 ± .02	0.19 ± .02
T2D	.39 ± .03	.4 ± .03	0.08 ± .01
RA	.73 ± .04	.79 ± .02	.80 ± .03
Hypertension	.56 ± .03	.55 ± .03	.56 ± .03
IBD	.4 ± .03	.4 ± .03	.07 ± .01
Adenoma	.15 ± .02	.13 ± .02	.04 ± .01
otitis	.12 ± .03	.13 ± .02	.00 ± 0
HBV	.62 ± .03	.59 ± .03	.80 ± .04
Fatty Liver	.32 ± .04	.22 ± .02	.02 ± .01
Psoriasis	.67 ± .04	.63 ± .04	.51 ± .05
T1D	.37 ± .07	.41 ± .07	.25 ± .06
Metabolic Syndrome	.52 ± 0.06	.5 ± .04	.08 ± .03
IGT	.12 ± .03	.18 ± .03	.00 ± 0
Periodontitis	.93 ± .02	.93 ± .03	.93 ± .03
AD	.8 ± .06	.78 ± .06	.82 ± .06
CDI	.67 ± .07	.53 ± .05	.54 ± .09
Infectious Gastroenteritis	.08 ± .04	.09 ± .05	.01 ± .02
Bronchitis	.02 ± .03	.05 ± .04	0 ± 0
Average	.46 ± .27	.45 ± .12	.35 ± .14

While problematic, class-imbalanced datasets are also the norm in reality, especially if

we are interested in the clinical context where most patients do not have the disease being screened for. While ROC curves are a useful class-distribution invariant performance metric, they can also mask the effect of a large false-positive rate. We thus computed precision-recall curves (PRC), which plot the relationship between positive-predictive-value (precision) and sensitivity (true-positive rate), in order to compare these classifiers in the context of class imbalance. We found that the random forest performed much better than the two neural network models with respect to the area under the PRC, implying that its predictions are more reliable in the context of class imbalance.

Disease	AUC			AUPR		
	GCN	DNN	RF	GCN	DNN	RF
Healthy	.84 ± .01	.82 ± .0	.83 ± 0	.73 ± .08	.88 ± .01	.93 ± 0
CRC	.84 ± .01	.81 ± .02	.5 ± 0	.41 ± .04	.35 ± .03	.53 ± 0
T2D	.9 ± .01	.86 ± .01	.5 ± 0	.4 ± .05	.31 ± .02	.52 ± 0
RA	.98 ± .01	.99 ± .0	.85 ± .02	.73 ± .04	.74 ± .03	.85 ± .02
Hypertension	.94 ± .01	.92 ± .01	.66 ± .01	.51 ± .05	.44 ± .03	.6 ± .02
IBD	.94 ± .02	.9 ± .01	.5 ± 0	.41 ± .04	.36 ± .02	.52 ± 0
Adenoma	.81 ± .02	.75 ± .02	.5 ± 0	.15 ± .03	.12 ± .01	.51 ± 0
Otitis	.86 ± .02	.74 ± .02	.5 ± 0	.18 ± .03	.1 ± .01	.51 ± 0
HBV	.97 ± .01	.96 ± .01	.69 ± .02	.64 ± .05	.58 ± .03	.7 ± .02
Fatty Liver	.89 ± .02	.8 ± .02	.5 ± 0	.35 ± .07	.2 ± .03	.51 ± 0
Psoriasis	.98 ± .01	.93 ± .02	.73 ± .03	.74 ± .03	.57 ± .04	.62 ± .03
T1D	.98 ± 0	.97 ± .01	.53 ± .03	.47 ± .07	.48 ± .07	.53 ± .02
Metabolic Syndrome	.98 ± .01	.97 ± .01	.51 ± .01	.63 ± .08	.55 ± .03	.51 ± .01
IGT	.95 ± .01	.88 ± .02	.5 ± 0	.22 ± .03	.16 ± .01	.5 ± 0
Periodontitis	.99 ± 0	.99 ± 0	.96 ± .02	.95 ± .02	.93 ± .02	.96 ± .02
AD	.99 ± 0	.99 ± 0	.88 ± .03	.89 ± .04	.78 ± .05	.85 ± .03
CDI	.99 ± 0	.91 ± .04	.64 ± .03	.75 ± .04	.64 ± .08	.62 ± .03
Infectious Gastroenteritis	.88 ± .03	.7 ± .1	.5 ± 0	.21 ± .04	.14 ± .01	.5 ± 0
Bronchitis	.78 ± .04	.74 ± .06	.5 ± 0	.13 ± .03	.13 ± .02	.5 ± .0
Average	.92 ± .03	.88 ± .12	.62 ± .07	.5 ± .11	.45 ± .12	.62 ± .07

Next, we analyzed the ranking of predictions by our models. In a clinical setting, a ranked ordering of probabilities may be more useful than a single output, because a physician can integrate that information along with the rest of their resources to generate a differential diagnosis. We examined the accuracy of the top-3 and top-5 predictions that each classifier produced. Most diseases were correctly identified within the top-3 classifications, and almost every disease was correctly present with at least 90% of the time in the top-5 classifications for each model. We measured which classes were most often predicted for a given label when an incorrect prediction was made for a subset of the labels (ie. which how often the model confused a given disease for another disease). Healthy was often misclassified as T2D or CRC, and every disease was often misclassified as healthy. These patterns were also

Fig. 2. (left) Accuracy at top-1,3, and 5 levels for (top to bottom) Random Forest, GCN, DNN. (right) Chord diagram showing (for a subset of labels) the most common classification made when an incorrect classification was made for a given class.

Fig. 3. tSNE visualization of the final layer activations of the GCN model (excluding healthy). Inset shows the same tSNE labeled by healthy (blue) and disease (red).

consistent across the three models. These errors may simply be due to the limited size of the dataset or noise in the system. But it could also indicate that the concept of a 'healthy microbiome' is vacuous given the broad of range of microbiomes that a healthy individual can have. Additionally, diseases can manifest with a broad range of severities, which may result in different metagenomic signatures which would all get grouped under a single label.

Lastly, we visualized the hidden layer activations of our GCN model on the test set using tSNE to see how different diseases cluster. We found that the skin conditions AD and psoriasis clustered together, separate from the other diseases which were evaluated on stool samples. Periodontitis, which was represented via oral microbiome samples, also clustered independently. Some other conditions, such as hypertension, adenoma, and otitis weakly clustered in some regions of the graph. We also visualized healthy and disease using the same tSNE, and we found that healthy microbiomes were scattered throughout the plot alongside disease, consistent with the idea that there is a broad range of possible healthy microbiomes.

5. Conclusion

We have extended the results of previous work on microbiome-phenotype prediction here by demonstrating that multiclass disease prediction from whole community metagenomes, a clinically relevant task for machine learning, is a tractable problem and is improved by using the taxonomic structure of bacterial communities. We implemented multiple classifiers that were able to discriminate between 18 different diseases and healthy with greater than 70% accuracy on a dataset of over 7000 samples. Moreover, while the GCN model generally outperformed the DNN, the RF excelled on a completely different set of metrics, indicating that these two models could potentially complement one another. RF based models, achieving 99% accuracy on healthy vs. not-healthy, could be used to identify dysbiosis in general, while GCNs could potentially then be used to discriminate between individual diseases. Additionally, the success of the GCN model implies that the geometric structure entailed in microbial phylogenies contains meaningful information for disease classification. This is particularly exciting result because graph and tree data structures are ubiquitous in systems biology, so GCN architectures may be applicable to a broad range of problems. There is already a body of literature around using these techniques in computational organic chemistry,⁴² and we believe that there are many more problems involving biochemical networks, gene regulatory networks, protein structure modeling, and computational neuroscience where these graph convolutional architectures can also be applied.

While there are obvious clinical applications of a successful phenotype classifier, there are many questions that still need to be answered before these techniques are accurate enough for use in the clinic. For example, we need to understand why the model makes certain predictions, because a medical practitioner will have more confidence in a prediction if it comes with an explanation. Recently, a class of algorithms called attribution methods have emerged which can identify predictive features in the input on a per-sample basis.^{43,44} This could be useful from a personalized genomics standpoint, as attribution methods could help explain which bacteria are contributing to dysbiosis in an individual patient and potentially suggest probiotic interventions to alleviate the dysbiosis. We will perform attribution analysis for our models to consider such questions in a future study.

Lastly, one of the shortcomings of this work is that it couples multiple questions together, making it difficult to study each one in isolation. For example, in addition to comparing different models for disease classification, we also made numerous assumptions in the construction of our dataset in terms of the diseases we included and the available metadata. We have not considered how factors such as gender, age, or body site might affect or be useful for enhancing a classifier's performance. Additionally, there are numerous pieces of metadata from individual studies that could be used when analyzing these metagenomes. Although the original R package for `curatedMetagenomicData` makes it easy to explore the data, constructing training and test sets for deep learning took substantial effort. To make this exercise easier for other researchers, we have made our code available on Github, which contains not only our models and training scripts, but also our code for downloading and pre-processing the metagenome samples. We believe that the results shown here will be improved upon in the future as more studies are added to `curatedMetagenomicData` and better models and training procedures

emerge. We hope that the code we provide will help other researchers attack this problem and similar problems involving machine learning with metagenomics.

6. Acknowledgments

Saad Khan was supported by the Einstein Medical Scientist Training Program (2T32GM007288-45). Libusha Kelly is supported in part by a Peer Reviewed Cancer Research Program Career Development Award from the United States Department of Defense (CA171019).

7. External Links

Source: github.com/kellylab/Metagenomic-Disease-Classification-With-Deep-Learning

References

- Vogtmann, E. *et al.* Colorectal cancer and the human gut microbiome: Reproducibility with whole-genome shotgun sequencing. *PLoS ONE* **11**, 1–13 (2016).
- Duvallet, C., Gibbons, S. M., Gurry, T., Irizarry, R. A. & Alm, E. J. Meta-analysis of gut microbiome studies identifies disease-specific and shared responses. *Nature Communications* **8**, 1784 (2017). URL <http://www.nature.com/articles/s41467-017-01973-8>.
- Lin, H. Y. Efficient classifiers for multi-class classification problems. *Decision Support Systems* **53**, 473–481 (2012).
- Reiman, D., Metwally, A. A. & Dai, Y. PopPhy-CNN : A Phylogenetic Tree Embedded Architecture for Convolution Neural Networks for Metagenomic Data 1–9 (2018).
- Fioravanti, D. *et al.* Phylogenetic Convolutional Neural Networks in Metagenomics 1–12 (2017). URL <http://arxiv.org/abs/1709.02268>.
- Thomas, A. M. *et al.* Metagenomic analysis of colorectal cancer datasets identifies cross-cohort microbial diagnostic signatures and a link with choline degradation. *Nature Medicine* **25**, 667–678 (2019). URL <http://dx.doi.org/10.1038/s41591-019-0405-7>.
- LaPierre, N., Ju, C. J., Zhou, G. & Wang, W. MetaPheno: A critical evaluation of deep learning and machine learning in metagenome-based disease prediction. *Methods* 0–1 (2019). URL <https://doi.org/10.1016/j.ymeth.2019.03.003>.
- Quince, C., Walker, A. W., Simpson, J. T., Loman, N. J. & Segata, N. Shotgun metagenomics, from sampling to analysis. *Nature Biotechnology* **35**, 833–844 (2017).
- Pasolli, E. *et al.* Accessible, curated metagenomic data through ExperimentHub. *bioRxiv* (2017). URL <http://biorxiv.org/content/early/2017/01/27/103085.abstract>.
- Truong, D. T. *et al.* MetaPhlAn2 for enhanced metagenomic taxonomic profiling. *Nature Methods* **12**, 902–903 (2015). URL <http://www.nature.com/doi/10.1038/nmeth.3589>.
- Asnicar, F. *et al.* Studying Vertical Microbiome Transmission from Mothers to Infants by Strain-Level Metagenomic Profiling. *mSystems* **2**, e00164–16 (2017). URL <http://msystems.asm.org/lookup/doi/10.1128/mSystems.00164-16>.
- Chng, K. R. *et al.* Whole metagenome profiling reveals skin microbiome-dependent susceptibility to atopic dermatitis flare. *Nature Microbiology* **1**, 1–10 (2016). URL <http://dx.doi.org/10.1038/nmicrobiol.2016.106>.
- Feng, Q. *et al.* Gut microbiome development along the colorectal adenomacarcinoma sequence. *Nature Communications* **6**, 6528 (2015). URL <http://www.nature.com/doi/10.1038/ncomms7528>.
- Hannigan, G. D., Duhaime, M. B., Ruffin, M. T., Koumpouras, C. C. & Schloss, P. D. Diagnostic potential and interactive dynamics of the colorectal cancer virome. *mBio* **9**, 1–13 (2018).
- Zeller, G. *et al.* Potential of fecal microbiota for early-stage detection of colorectal cancer. *Molecular Systems Biology* **10**, 766–766 (2014).
- Yassour, M. *et al.* Natural history of the infant gut microbiome and impact of antibiotic treatments on strain-level diversity and stability. *Sci Trans Med* **8**, 1173–1178 (2016). URL <http://dx.doi.org/10.1126/scitranslmed.aba4066>.
- Vincent, C. *et al.* Bloom and bust: Intestinal microbiota dynamics in response to hospital exposures and *Clostridium difficile* colonization or infection. *Microbiome* **4**, 1–11 (2016). URL <http://dx.doi.org/10.1186/s40168-016-0156-3>.
- Yu, J. *et al.* Metagenomic analysis of faecal microbiome as a tool towards targeted non-invasive biomarkers for colorectal cancer. *Gut* **66**, 70–78 (2017).
- Loomba, R. *et al.* Gut Microbiome-Based Metagenomic Signature for Non-invasive Detection of Advanced Fibrosis in Human Nonalcoholic Fatty Liver Disease. *Cell Metabolism* **25**, 1054–1062.e5 (2017). URL <https://linkinghub.elsevier.com/retrieve/pii/S1550413117302061>.
- Qin, N. *et al.* Alterations of the human gut microbiome in liver cirrhosis. *Nature* **513**, 59–64 (2014). URL <http://dx.doi.org/10.1038/nature13568>.
- Li, Y., Yu, R., Shahabi, C. & Liu, Y. Diffusion Convolutional Recurrent Neural Network: Data-Driven Traffic Forecasting 1–16 (2017). URL <http://arxiv.org/abs/1707.01926>.
- Nielsen, H. B. *et al.* Identification and assembly of genomes and genetic elements in complex metagenomic samples without using reference genomes. *Nature Biotechnology* **32**, 822–828 (2014).
- Karlsson, F. H. *et al.* Gut metagenome in European women with normal, impaired and diabetic glucose control. *Nature* **498**, 99–103 (2013).
- David, L. A. *et al.* Diet rapidly and reproducibly alters the human gut microbiome. *NIH Public Access* **505**, 559–563 (2014).
- Vrieze, A. *et al.* Transfer of intestinal microbiota from lean donors increases insulin sensitivity in individuals with metabolic syndrome. *Gastroenterology* **143**, 913–916.e7 (2012). URL <http://dx.doi.org/10.1053/j.gastro.2012.06.031>.
- Shi, B. *et al.* Dynamic changes in the subgingival microbiome and their potential for diagnosis and prognosis of periodontitis. *mBio* **6**, 1–11 (2015).
- Tett, A. *et al.* Unexplored diversity and strain-level structure of the skin microbiome associated with psoriasis. *npj Biofilms and Microbiomes* **3**, 1–11 (2017). URL <http://dx.doi.org/10.1038/s41522-017-0022-5>.
- Wen, C. *et al.* Quantitative metagenomics reveals unique gut microbiome biomarkers in ankylosing spondylitis. *Genome Biology* **18**, 1–13 (2017).
- Heintz-Buschart, A. *et al.* Integrated multi-omics of the human gut microbiome in a case study of familial type 1 diabetes. *Nature Microbiology* **2**, 16180 (2016). URL <http://www.nature.com/articles/nmicrobiol2016180>.
- Kostic, A. D. *et al.* The dynamics of the human infant gut microbiome in development and in progression towards type 1 diabetes. *Cell Host and Microbe* **17**, 260–73 (2015).
- Qin, J. *et al.* A metagenome-wide association study of gut microbiota in type 2 diabetes. *Nature* **490**, 55–60 (2012). URL <http://www.nature.com/doi/10.1038/nature11450>.
- Bronstein, M. M., Bruna, J., Lecun, Y., Szlam, A. & Vanderghenst, P. Geometric Deep Learning: Going beyond Euclidean data. *IEEE Signal Processing Magazine* **34**, 18–42 (2017). URL <http://dx.doi.org/10.1109/SPM.2017.799097>.
- Kipf, T. N. & Welling, M. SEML-SUPERVISED CLASSIFICATION WITH GRAPH CONVOLUTIONAL NETWORKS 1–14 (2017). URL <https://arxiv.org/abs/1609.02907>.
- Ng, A. *Machine Learning Yearning* (deeplearning.ai).
- Pey, M. & Lenssen, J. E. Fast Graph Representation Learning with PyTorch Geometric 1–9 (2019). URL <http://arxiv.org/abs/1903.02428>.
- Paszke, A. *et al.* Automatic differentiation in PyTorch. In *NIPS Autodiff Workshop* (2017).
- Pedregosa, F. *et al.* Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* **12**, 2825–2830 (2011).
- Malla, M. A. *et al.* Exploring the human microbiome: The potential future role of next-generation sequencing in disease diagnosis and treatment. *Frontiers in Immunology* **10**, 1–23 (2019).
- Johnson, J. M. & Khoshgoftar, T. M. Survey on deep learning with class imbalance. *Journal of Big Data* **6** (2019). URL <https://doi.org/10.1186/s40537-019-0192-5>.
- Publication, K. S. & E. R. In Defense of One-Vs-All Classification. *Journal Of Machine Learning Research* **5**, 2–6 (2004).
- Khan, S. & Kelly, L. Fireworks: Reproducible Machine Learning and Preprocessing with PyTorch. *Journal of Open Source Software* **4**, 1478 (2019). URL <https://joss.theoj.org/papers/10.21105/joss.01478>.
- Coley, C. W. *et al.* A graph-convolutional neural network model for the prediction of chemical reactivity. *Chemical Science* **10**, 370–377 (2019).
- Shrikumar, A., Greenside, P. & Kundaje, A. Learning Important Features Through Propagating Activation Differences (2017). URL <http://arxiv.org/abs/1704.02685>.
- Lundberg, S. & Lee, S.-I. An unexpected unity among methods for interpreting model predictions 1–6 (2016). URL <http://arxiv.org/abs/1611.07478>.