

# Machine Learning Applications to DNA Subsequence and Restriction Site Analysis

E. Moyer<sup>1</sup> and A. Das<sup>2</sup>

1. School of Biomedical Engineering, Science and Health Systems, Drexel University, Philadelphia, Pennsylvania, USA

2. College of Engineering, Drexel University, Philadelphia, Pennsylvania, USA

{ejm374, anup.das}@drexel.edu

**Abstract**—Based on the BioBricks™ standard, restriction synthesis is a novel catabolic iterative DNA synthesis method that utilizes endonucleases to synthesize a query sequence from a reference sequence. In this work, the reference sequence is built from shorter subsequences by classifying them as applicable or inapplicable for the synthesis method using three different machine learning methods: Support Vector Machine (SVM), random forest, and Convolution Neural Network (CNN). Before applying these methods to the data, a series of feature selection, curation, and reduction steps are applied to create an accurate and representative feature space. Following these preprocessing steps, three different pipelines are proposed to classify subsequences based on their nucleotide sequence and other relevant features corresponding to the restriction sites of over 200 endonucleases. The sensitivity using SVM, random forest, and CNN are 94.9%, 92.7%, 91.4%, respectively. Moreover, each method scores lower in specificity with SVM, random forest, and CNN resulting in 77.4%, 85.7%, and 82.4%, respectively. In addition to analyzing these results, the misclassifications in SVM and CNN are investigated. Across these two models, different features with a derived nucleotide specificity visually contribute more to classification compared to other features. This observation is an important factor when considering new nucleotide sensitivity features for future studies.

**keywords:** machine learning, subsequence matching, DNA sequence classification

## I. INTRODUCTION

Originally postulated by Watson and Crick in the 1950s, the process of cellular deoxyribonucleic acid (DNA) synthesis has long been understood as a trademark to biological studies [1] [2]. Over the past few decades, understanding this process has led to the discovery and development of many synthetic manufacturing methods, namely synthetic gene synthesis [3] [4]. As a result, other areas of genomics have accumulated massive research interest. From the start of the Human Genome Project in 1990 to the discovery of the Cluster Regularly Interspaced Palindromic Repeat (CRISPR) Cas9 system in bacteria during 2005, the field of genomics has advanced at an incredible rate as the reservoir of information pertaining to DNA and gene research increases [5] [6]. Although the public adoption of these technologies is in the near future, the cost of producing novel genes remains as the field's largest determinant [7].

Stemming from the need to decrease the cost of synthetic synthesis, the rationale behind this research begins with the introduction of a hypothesized catabolic synthesis process, restriction synthesis. This procedure relies on a process known as restriction enzyme digest. This typical laboratory procedure utilizes restriction endonucleases, which are enzymes that cleave DNA at specific restriction sites, or target locations. [8]. Figure 1 shows an example restriction enzyme digest using enzyme EcoR1 on restriction site GAATTC.




Figure 1. Restriction enzyme digest of GAATTC using EcoR1. In general, DNA is represented as a permutation of the four nucleotides [adenine (A), thymine (T), cytosine (C), guanine (G)] and an ambiguous base (N).

Although restriction enzyme digest is normally reserved for a few endonucleases at a time, restriction synthesis serializes the use of many endonucleases to digest a reference sequence into fragments. These fragments are then used to synthesize a query sequence by ligating them to one another in succession as shown in Figure 2. This process is a modification to the BioBricks™ standard proposed by Dr. Thomas Knight. Anderson et al. defines the BioBricks™ standard as a process that “employs iterative restriction enzyme digestion and ligation reactions to assemble small basic parts into larger composite parts” [9]. These larger components are built using fragments flanked by two restriction enzymes, XbaI and SpeI. The key difference between restriction synthesis and the BioBricks™ standard is that the latter only digests fragments with XbaI and SpeI endonucleases. Conversely, restriction synthesis relies on a data set of approximately 200 endonucleases from New England Biolabs [10]. In flanking a reference sequence with many endonucleases, restriction synthesis builds more composite fragments homologous to those

referenced by the BioBricks™ standard [11].




Figure 2. Restriction Synthesis. Reference sequence is flanked by Enzyme A and Enzyme B. This fragment is then ligated to the end of the query sequence.

One assumption of restriction synthesis is that the reference sequence and the query sequence are both known prior to synthesis. If each query sequence is synthesized from an independent reference sequence, this assumption would require costly sequencing. Although this one-to-one correspondence between the query sequence and the reference sequence would result in incredibly accurate synthesis, it would massively increase the cost of this method. A solution is to build one reference sequence from which many different query sequences can be synthesized. While this solution is a less expensive alternative to the former, it requires implementing a procedure to carefully choose which subsequences should be included in the reference sequence. An applicable subsequence is defined as containing a candidate fragment that can be flanked by two endonucleases. It follows that subsequences can be classified based on whether they are applicable (true) or inapplicable (false) for inclusion in the reference sequence.

The data set for this research is obtained through the synthesis of mycobacterium structural genes during restriction synthesis. Each gene sequence ranges from 1,000 to 5,000 base pairs and is obtained from the National Center for Biotechnology Information (NCBI) [12]. Using this data, the analysis of subsequences is explored using three popular techniques in machine learning and neural network architecture: Support Vector Machine (SVM), random forest, and Convolution Neural Network (CNN) [13].

## II. FEATURE SELECTION

The feature set for each subsequence entry includes 16 unique variables. The 'SEQ' feature is responsible for capturing the nucleotide sequence of each entry formatted based on the following convention: nucleotides A, T, C, and G are assigned ordinal values 1/4, 1/2, 3/4, and 1, respectively, and the ambiguous base, N, is assigned an ordinal value of 0. This feature most directly relates whether a specific sequence can be flanked by two enzymes due to the restriction sites of each endonuclease. Another feature, 'LEN,' describes the length of each sequence, which is related to the

dimension of the 'SEQ' feature.

Eight other features are obtained using Equation 1, which describes the proportion of each nucleotide in a sequence. Four of the eight features describes the proportion of each nucleotide in the subsequence, and the other set of four describes that of the reference sequence.

The next set of four features are related to a novel sequence metric, the complexity rating, as shown by Equation 2 and Equation 3. These equations are used to accurately describe the nucleotide composition with a continuous value ranging from 0 exclusive to 1 inclusive. Equation 2 is responsible for determining the deviation of an entire sequence from an equal nucleotide composition, whereas Equation 3 segments a sequence based on preset values ranging from  $b$  to  $p$  and examines this deviation for each segment proportionally. In other words, the rationale behind Equation 3 is to sum complexity ratings from Equation 2 in proportion to the number of occurrences of segment size  $i$  in sequence size  $n$ . In short, both of these ratings determine how far a sequence deviates from an equal nucleotide composition. The four features represent both the  $r_1$  and  $r_2$  values of the subsequence and reference sequence.

$$p(x,L) = \frac{\text{occurrence of } x}{L} \quad (1)$$

$$r_1(L) = 1 - \sum_{i \in \alpha} \left( \frac{1}{4} - p(i,L) \right)^2 \quad (2)$$

$$r_2(n) = \sum_{i=b}^p \frac{\sum_{j=0}^{n-i} r_1(i) * \frac{i}{n}}{\sum_{k=b}^p \frac{k}{n}} \quad (3)$$

Figure 3. Relevant equations for features describing the nucleotide proportion and complexity ratings of both the subsequence and the query sequence. Note:  $\alpha = \{A,T,C,G\}$

Two additional auxiliary features, 'C' and 'EZY,' are included to describe the number of query sequences in a specific list that contain the subsequence and the number of available enzymes that can digest the subsequence, respectively. These are incorporated to gauge whether a particular subsequence is common in the data set.

## III. FEATURE ANALYSIS

Feature reduction is performed in order to simplify the feature space for each of the three machine learning techniques. A correlation test is applied across all features to determine whether any of them are redundant. A conservative threshold is set at 0.90 because of the small size of the feature set. None of the pairs of features had a correlation higher than 0.90, so all features are included in the study. A sample feature correlation test on the first six features can be observed in Table 1 below.

	N1	N2	N3	N4	N5	N6
N1	1.000					
N2	0.051	1.000				
N3	-0.013	0.006	1.000			
N4	-0.062	-0.007	-0.002	1.000		
N5	-0.015	0.088	-0.011	0.001	1.000	
N6	-0.042	0.025	0.119	-0.017	-0.028	1.000

Table 1. An example correlation test on the first six nucleotides (N1 to N6) for the 'SEQ' feature.

Surprisingly, this indicates low nucleotide-nucleotide interaction within a subsequence. One would expect some degree of interaction due to the nature of restriction enzymes cutting at specific restriction sites.

#### IV. DATA SELECTION & CURATION

A constant training set and testing set is used while evaluating each of the machine learning solutions. Based on an initial raw data set of approximately one million entries across almost one hundred gene synthesis simulations, a sample of 105,000 entries is selected with a 50/50 percent distribution of applicable and inapplicable subsequences. From that data set, a stratified sample produces both the training set and testing set with unique entries. The training set consists of 50,000 unique entries with a 60/40 percent distribution of the classes, and the testing set consists of 26,622 unique entries with a 50/50 percent distribution of the classes.

#### V. METRICS

In this work, we compare machine learning techniques using the metrics of Equation 4-7.

$$\text{sensitivity} = \frac{\text{true positive}}{\text{true positive} + \text{false negative}} \quad (4)$$

$$\text{specificity} = \frac{\text{true negative}}{\text{true negative} + \text{false positive}} \quad (5)$$

$$\text{false negative rate (FNR)} = 1 - \text{sensitivity} \quad (6)$$

$$\text{false positive rate (FPR)} = 1 - \text{specificity} \quad (7)$$

Figure 4. Sensitivity, specificity, false negative rate, and false positive rate equations.

#### VI. METHODS

##### 1. Support Vector Machines (SVM)

SVM is a supervised learning model that creates hyperplanes between groups of data [14]. Given a set of training data and corresponding labels, the algorithm attempts to best categorize data by measuring the distance of each data point to the segregating hyperplane [15]. By maximizing these distances, SVM creates highly reliable non-probabilistic classifications [16].

In addition to the feature analysis, Principle Component Analysis (PCA) is also used for feature reduction.

PCA is an unsupervised learning method that reduces a feature space while retaining relationships in the data. By transforming a high dimensional feature space into a specified number of principle components (PCs), raw data can be summarized in a few dimensions [17].

As seen in Figure 5, the classification can be easily visualized when 2 PCs are used. In this case, the most successful classifier is the 2 PCs polymetric kernel. The results from the testing set of 26,622 entries can be observed in the confusion matrix shown in Table 2.




Figure 5. Polymetric 2 PCs SVM analysis.

	Actual True	Actual False
Predicted True	12632	3009
Predicted False	678	10304

Table 2. Confusion matrix for polymetric 2 PCs SVM.

Table 2 displays that the sensitivity is 94.9% and the specificity is 77.4%. These results are seen consistently across all trials of this configuration. For this reason, the standard deviation is considered negligible. Furthermore, Table 2 indicates a 5.1% FNR and a 22.6% FPR.

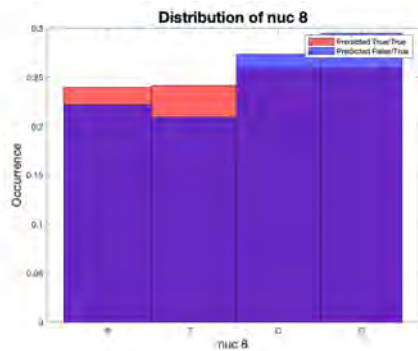
In addition to the polymetric 2 PCs classifier, five other tests are conducted: non-PCA linear kernel, 3 PCs linear kernel, 3 PCs polymetric kernel, 2 PCs RBF kernel, and 2 PCs sigmoid kernel. The results are below in Table 3.

Classifier	Sensitivity (%)	Specificity (%)	FNR (%)	FPR (%)
Non-PCA linear kernel	89.3	83.6	10.7	16.4
3 PCs linear kernel	92.5	81.3	7.5	18.7
3 PCs polymetric kernel	92.7	81.7	7.3	18.3
2 PCs RBF kernel	93.1	81.7	6.9	18.3
2 PCs sigmoid kernel	75.4	75.5	24.6	24.5

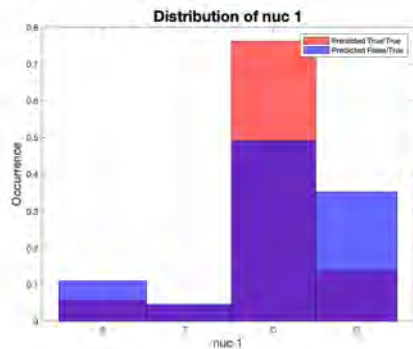
Table 3. Alternative PCA results.

The two histograms in Figure 6 display the relative proportions of each nucleotide at different features under the 'SEQ' feature. Red bars indicate correctly classified true entries, and blue bars indicate incorrectly classified true entries. Purple bars are simply the overlap of the

two. Many of the features in this data set had distributions similar to the two histograms in Figure 6a and Figure 6b. When the distribution is seemingly uniform across all four nucleotides, such as in Figure 6a, there is a displayed indifference between the four nucleotides at a specific nucleotide position. This suggests that at least across this data set, the eighth nucleotide position is insensitive to the type of base. Conversely, Figure 6b displays higher base sensitivity as it favors cytosine far more than any of the other three nucleotides on the first nucleotide position. This distribution suggests the first position is incredibly sensitive to the type of base.



(a) Distribution of the eighth nucleotide position.



(b) Distribution of the first nucleotide position.

Figure 6. Feature distribution for SVMs

Further analysis is needed in order to determine whether these differences are statistically significant. If this is the case, it would suggest a difference in nucleotide sensitivity across nucleotide positions.

Another interesting question arises on the misclassifications between these two features. For instance, Figure 6a shows both over and under classification of true and false subsequences for the eighth nucleotide position. Further statistical analysis may suggest that there is not a significant difference between misclassification rates for each nucleotide. Additionally, Figure 6b displays that the first nucleotide position misclassifies C and G in much higher proportions than A and T. Further statistical analysis may suggest this to be the case. More interestingly, these misclassifications seem to follow

the distributions of correctly matched sequences. For example, Figure 6a misclassified sequences uniformly across that feature. This phenomenon is also apparent in Figure 6b as the proportion of misclassifications follow the distribution of true classifications for that feature.

## 2. Random Forest

Random forest is an ensemble learning model that implements a set number of decision trees based on subsets of training data for classification and regression [18]. One of the features of random forest is its use of out-of-bag data [19]. At the construction of each tree, two-thirds of the data is included for training and the remaining one-third is set aside for validation after the training has completed [20]. During the training process, each decision tree is generated based on a segmented data set and a randomly selected feature set at each node of the tree [21]. While a single decision tree may overfit data, many decision trees in random forest allow for the overfitting to average out between trees. When predicting based on new data, each tree provides a prediction, or a vote, for a certain class. The final classification for random forest is based on which classification has the most votes [18].

The random forest analysis is conducted with 10, 20, 30, 40, and 50 trees. The most successful analysis is conducted when there were 30 trees, but even as this number passed 30, the additional level of accuracy in both categorizations began to level off. In other words, the classification accuracy when the number of trees is 50 is the same or less than when there were just 40 and 30 trees. The results from the test set of 26,622 entries can be observed in the confusion matrix in Table 4.

	Actual True	Actual False
Predicted True	12336	1898
Predicted False	974	11415

Table 4. Confusion matrix for random forest (n=30).

With 30 decision trees, the sensitivity is 92.7% and the specificity is 85.7% for the testing set of 26,622 entries. These results are seen consistently across all trials of random forest. For this reason, the standard deviation is considered negligible. In Table 4, the confusion matrix displays a 7.3% FNR and 14.2% FPR.

## 3. Convolution Neural Networks (CNN)

CNN is a class of deep learning models that use a linear operation called convolution in at least one of their layers [22]. The model usually begins with a convolution layer that accepts a tensor as an input with dimensions based on the size of the data [23]. The second layer, or the first hidden layer, is formed by applying a kernel or filter that is a smaller matrix of weights over a receptive field, which is a small subspace

of the inputs [25]. Kernels apply an inner product on the receptive field, effectively compressing the size of the input space [24]. As the kernel strides across the input space, the first hidden layer is computed based on the weights of the filter. As a result, the first hidden layer is a feature map formed from the kernel applied on the input space [23]. While the dimension of the kernel may be much smaller in size compared to the initial inputs of the convolution layer, the kernel must have the same depth of the input space. The inputs and convolution layer are often followed by rounds of activations, normalizations, and pooling [24]. The last layer, however, is a fully connected layer where the final outputs or categorizations are determined based on how different features fall in line with the specific classes under study [25].

Figure 7 models the proposed Convolution Neural Network (CNN) for this data set.

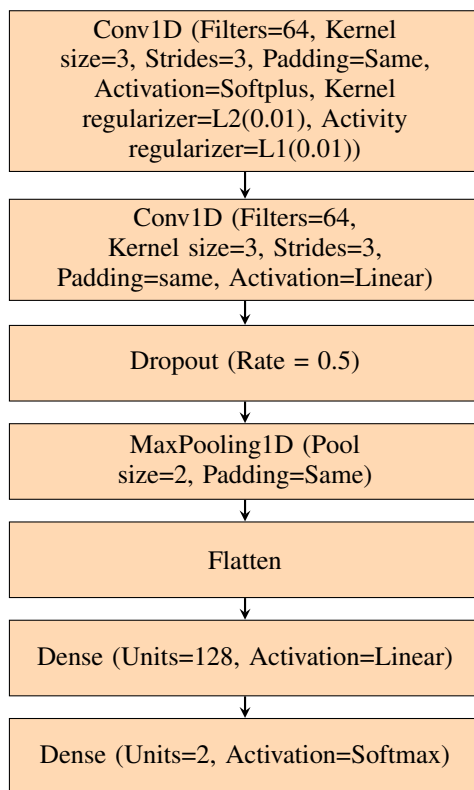
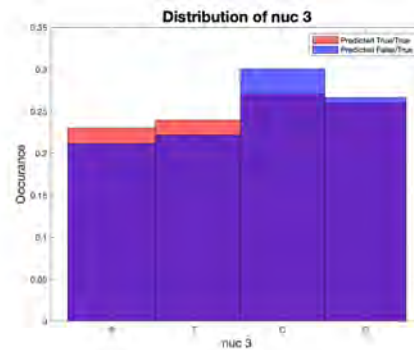


Figure 7. Currently optimal CNN

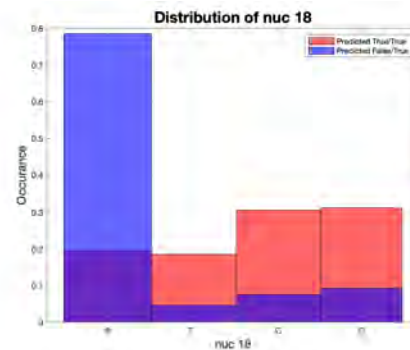
After 100 trials of classification using this model, the average sensitivity is 91.4% and the FNR is 8.6% with a standard deviation of 3.5%, and the specificity is 82.4% and the FPR is 17.6% with a standard deviation of 1.2% for the testing set of 26,622 entries.

As seen in Figure 8, uniform and nonuniform distributions with respect to feature specific misclassifications

are similar to that seen with SVM. For this reason, the nucleotide sensitivity across different positions in Figure 8a follows that observed in Figure 6a. Conversely, Figure 8b displays a new trend observed in the true misclassifications from the eighteenth nucleotide feature onward to the last nucleotide feature: the distribution of true misclassifications does not follow the distribution of correctly classified sequences. In this case, the true correctly classified true entries in red display an almost uniform distribution, while the incorrectly classified true entries in blue display a distribution almost completely favored for A. This may indicate that at least toward the end of the ‘SEQ’ feature, the model failed to recognize a uniform distribution of nucleotides.



(a) Distribution of the third nucleotide position.



(b) Distribution of the eighteenth nucleotide position.

Figure 8. Feature distribution for CNNs

## VII. DISCUSSION

When analyzing these models, one inquiry is whether a high FNR or FPR is less desirable. In the scope of this research, a false positive is defined as including a subsequence in the reference sequence when it should not have been included, and a false negative is defined as deciding not to include a subsequence in the reference sequence when it should have been included. In each optimal model, the FPR is much higher than the FNR. Initially, one would think that false positives are less desirable than false negatives. One of the arguments for this proposition is that fewer false positives allow the reference to be smaller in size. Since the size of the



reference is not determined to be a major constraint, this need not be an issue.

Instead, the inclusion of these false positives increases the genetic variation in the reference if it is used recursively. Prior to restriction synthesis, a large reference would initially be generated from a random sequence. This reference would then be subsequently used in classifying more entries and so on and so forth. Paramount to the field of genomics, genetic variation in this context would allow for more subsequences to be analyzed and classified if this recursive process is applied with references containing a slightly higher FPR and slightly lower FNR.

The performance between these three methods pose another interesting discussion. The sensitivity of SVM, random forest, and CNN are 94.9%, 92.7%, and 91.4% respectively. At least for this data set, SVM and random forest are far more consistent and accurate than CNN. Furthermore, the average specificity between all three vary greatly with random forest performing better than SVM and CNN. Of all three, SVM had the lowest specificity at 77.4%, while CNN (82.4%) trails behind random forest (85.7%) by only 3.3% and SVM behind random forest by 8.3%. In then comparing SVM to random forest based on the sensitivity, the question becomes whether a 2.2% increase in sensitivity is worth a 8.3% decrease in specificity. Because of the previous discussion on the FPR and the FNR, the answer is definitively yes. To reiterate, a higher FPR coupled with a higher sensitivity does not impact the classification if the process is implemented recursively. For this reason, SVM is ranked ahead of random forest in these respects. Because both the specificity and sensitivity of CNN are below that of SVM and random forest, random forest is ranked ahead of CNN but below SVM. It is interesting, though, that all three techniques exhibited a lower specificity. One possible interpretation is that true cases are easier to differentiate.

SVM out performs both CNN and random forest in sensitivity. One hypothesis for why this might have been the case is due to the use of PCA to build representative components from the data. Surprisingly, all 16 unique features were able to be reduced to only two components, allowing for the classification to be observed in a two-dimensional space. Although SVM are the only technique to reach near the 95% benchmark for sensitivity, the two other techniques did perform better than 90% in that category. This indicates that these techniques may be improved, but it may require more work in the realm of feature selection, data preprocessing, and architecture design.

#### ACKNOWLEDGMENT

This work is supported by the National Science Foundation Award CCF-1937419 (RTML: Small: Design of System Software to Facilitate

Real-Time Neuromorphic Computing).

#### REFERENCES

- [1] W. Tatum and R. E. Hausman. *The Cell: A Molecular Approach, Sixth Edition*. 2000.
- [2] Peter Reichard. Interactions between deoxyribonucleotide and dna synthesis. *Annual review of biochemistry*, 57(1):349–374, 1988.
- [3] R. A. Hughes and A. D. Ellington. Synthetic DNA Synthesis and Assembly: Putting the Synthetic in Synthetic Biology. *Cold Spring Harbor Perspectives in Biology*, 9, January 2017.
- [4] Marvin H Caruthers. Gene synthesis machines: Dna chemistry and its uses. *Science*, 230(4723):281–285, 1985.
- [5] National Human Genome Research Institute. The Human Genome Project, 2020.
- [6] Broad Institute. CRISPR Timeline, 2020.
- [7] Nicholas Tang, Siying Ma, and Jingdong Tian. Chapter 1 - New Tools for Cost-Effective DNA Synthesis. In Huimin Zhao, editor, *Synthetic Biology*, pages 3 – 21. 2013.
- [8] Cynthia H Bigger, Kenneth Murray, and Noreen E Murray. Recognition sequence of a restriction enzyme. *Nature New Biology*, 244(131):7–10, 1973.
- [9] et al. Anderson, JChristopher. Bglbricks: A flexible standard for biological part assembly. *Journal of Biological Engineering*, 4(1):1, 2010.
- [10] New England Biolabs. Restriction Endonucleases, 2020.
- [11] Reshma P Shetty, Drew Endy, and Thomas F Knight. Engineering biobrick vectors from biobrick parts. *Journal of biological engineering*, 2(1):5, 2008.
- [12] Gene [Internet]. Bethesda (MD): National Library of Medicine (US), National Center for Biotechnology Information, 2020.
- [13] Nucleotide [Internet]. National Library of Medicine (US), National Center for Biotechnology Information, 2020.
- [14] Johan AK Suykens and Joos Vandewalle. Least squares support vector machine classifiers. *Neural processing letters*, 9(3):293–300, 1999.
- [15] William S Noble. What is a support vector machine? *Nature biotechnology*, 24(12):1565–1567, 2006.
- [16] Baozhen Yao, Ping Hu, Mingheng Zhang, and Maoqing Jin. A support vector machine with the tabu search algorithm for freeway incident detection. *International Journal of Applied Mathematics and Computer Science*, 24(2):397–404, 2014.
- [17] J. Lever and M. Krzywinski and N. Altman. Principal component analysis. *Nature Methods*, 14(7):641–642, 2017.
- [18] Mahesh Pal. Random forest classifier for remote sensing classification. *International journal of remote sensing*, 26(1):217–222, 2005.
- [19] Andy Liaw, Matthew Wiener, et al. Classification and regression by randomforest. *R news*, 2(3):18–22, 2002.
- [20] L. Breiman and A. Cutler. Random Forests - Classification Description. Berkeley University, 2020.
- [21] Adele Cutler and John R. Stevens. Random forests for microarrays. In *DNA Microarrays, Part B: Databases and Statistics*, volume 411 of *Methods in Enzymology*, pages 422 – 432. 2006.
- [22] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016.
- [23] Yoshua Bengio, Ian Goodfellow, and Aaron Courville. *Deep learning*, volume 1. MIT press Massachusetts, USA.; 2017.
- [24] LeCun, Yann and Bengio, Yoshua and Hinton, Geoffrey. *Deep learning*, volume 521. Nature, pages 436–444, 2015.
- [25] Jürgen Schmidhuber. Deep learning in neural networks: An overview. *Neural Networks*, 61:85–117, 2015.