

Spike2Vec: An Efficient and Scalable Embedding Approach for COVID-19 Spike Sequences

1st Sarwan Ali

Department of Computer Science
Georgia State University
Atlanta, GA, USA
sali85@student.gsu.edu

2nd Murray Patterson

Department of Computer Science
Georgia State University
Atlanta, GA, USA
mpatterson30@gsu.edu

Abstract—With the rapid global spread of COVID-19, more and more data related to this virus is becoming available, including genomic sequence data. The total number of genomic sequences that are publicly available on platforms such as GISAID is currently several million, and is increasing with every day. The availability of such *Big Data* creates a new opportunity for researchers to study this virus in detail. This is particularly important with all of the dynamics of the COVID-19 variants which emerge and circulate. This rich data source will give us insights on the best ways to perform genomic surveillance for this and future pandemic threats, with the ultimate goal of mitigating or eliminating such threats. Analyzing and processing the several million genomic sequences is a challenging task. Although traditional methods for sequence classification are proven to be effective, they are not designed to deal with these specific types of genomic sequences. Moreover, most of the existing methods also face the issue of scalability. Previous studies which were tailored to coronavirus genomic data proposed to use spike sequences (corresponding to a subsequence of the genome), rather than using the complete genomic sequence, to perform different machine learning (ML) tasks such as classification and clustering. However, those methods suffer from scalability issues.

In this paper, we propose an approach called Spike2Vec, an efficient and scalable feature vector representation for each spike sequence that can be used for downstream ML tasks. Through experiments, we show that Spike2Vec is not only scalable on several million spike sequences, but also outperforms the baseline models in terms of prediction accuracy, F1 score, etc. Since this type of study on such huge numbers of spike sequences has not been done before (to the best of our knowledge), we believe that it will open new doors for researchers to use this data and perform different tasks to unfold new information that was not available before. We also use information gain (IG) to compute the importance of each amino acid in the spike sequence. The amino acids with higher IG values tend to be the same as many reported by the USA based Centers for Disease Control and Prevention (CDC) for different variants.

Index Terms—COVID-19 Spike Sequences, Feature Vector Representation, k-mers, Classification, Clustering

I. INTRODUCTION

Very few fields of study remain untouched in the big data era, as massive amounts of data are collected in every domain from finance [1], [2] to astronomy [3], [4]. The field of biomedical and health informatics is no exception; one which has had a recent and rather rapid growth spurt in the amount of available data, due to the COVID-19 pandemic [5]–[8]. One facet of this increase is the amount of genomic

data becoming available for COVID-19 in databases such as GISAID [9], where several million viral genome (virome) sequences of COVID-19 — or more precisely, SARS-CoV-2¹ — are available.

Such data has a high *volume*, as the SARS-CoV-2 virome has $\approx 30K$ nucleotide base-pairs, and there are more than 2.5 million such sequences available in GISAID alone. While the number COVID-19 patients being sequenced is a fraction of the actual number of cases, the sheer number of infections (both now and in the past) means that the *velocity* in which SARS-CoV-2 virome sequences are appearing is very high. For example, in March 2020, when COVID-19 was declared a pandemic by the world health organization (WHO), there were a few thousand sequences available. This grew to tens of thousands in the late summer, when the Alpha variant emerged in the UK [10]. By the end of 2020 it was hundreds of thousands, and in early 2021 it had reached 1 million; today it is over 2.5 million. This will likely continue to increase exponentially (see [11]) as many countries [12], [13] ramp up their sequencing infrastructure for COVID-19 and future pandemics.

The available SARS-CoV-2 virome sequence data, in databases such as GISAID, has a high *variety*, since it comprises sequences from all over the world. Since COVID-19 has spread all over the world for more than a year now, and viruses continue to mutate over time, there are quite a number of variants of the SARS-CoV-2 virome, and they continue to emerge. Because variants continue to emerge and die off, some epidemiologists have even proposed a dynamic nomenclature system similar to that used for the common cold or flu [14]. We use this so-called “Pango Lineage” nomenclature to identify the variants we study here, since only the very common variants of concern (VoCs) are named. Examples of such named VoCs are the Alpha [10] (Pango Lineage B.1.1.7), Gamma [15] (P.1) and Delta [16] (B.1.617.2) variants (see Table I for a more complete list). The genomic variations (which happen disproportionately in the spike region, see Figure 1) that define these different variants have been associated with increased transmissibility [17], and immune evasion [18].

¹Severe Acute Respiratory Syndrome Coronavirus 2 (SARS-CoV-2), the virus which causes the COVID-19 disease

Because databases such as GISAID [9] collect sequences from all over the world, they come from heterogeneous sources of sequencing technologies and centers, leading to multiple levels of *veracity*. However, the largest source of different veracity in the data is the widely varying degree to which different populations are represented. For example, the UK sequences about 5% of its population of ≈ 70 million, the USA sequences about 1% of its population of ≈ 300 million, while India sequences only a fraction of a percent of its population of ≈ 1.3 billion [19], [20] (see Figure 3). Because of this, for example, even though the Delta variant likely originated in India, the majority of the available sequences of this variant are from the UK and the USA, after it arrived in these countries [9].

Since the genomic sequence of a virus encodes all of its functions such as virulence and transmissibility, the *value* of such massive amounts of genomic data is clear. It is variation in this genomic sequence itself which defines the different variants of SARS-CoV-2 such as Alpha, Delta and Gamma. All of these variants differ from each other in effect (due to their unique genomic variations), yet they all descend from the original SARS-CoV-2 sequence [21]. It is only through a process of evolution and transmission to many parts of the globe for over a year, has it diverged to this extent. The amount of sequence data available today puts us in the age of *genomic surveillance*: tracking the spread of pathogens in terms of genomic content [22], [23].

Approaches for rapidly clustering and classifying sequences will be crucial in these genomic surveillance efforts. A clustering method, when applied to the data on a daily basis, for example, would identify a new and rapidly emerging variant in terms of a cluster which grows abnormally quickly, allowing scientists to focus on this cluster. Classification, on the other hand, would allow us to track the spread of known variants in new municipalities, regions, countries and continents. For example, the USA had a wave of the Alpha variant from the UK in early 2021, and later, a wave of the Delta variant from India and via other intermediaries, such as the UK (see Figure 5). Such patterns of spread can reveal information about the underlying transmission networks between different countries (the UK and USA, or the UK and India), or even parts of different countries. This can help overcome some of the different veracity in the data, such as the widely varying degree to which different countries are represented in terms of sequencing data, due to sampling bias. For example, even though India is very under-sampled compared to the UK, the wave of the Delta variant in the UK, along with information about flights from India to the UK near the beginning of this wave could give us insights on how the Delta variant originated and spread in India.

The development of clustering and classification approaches needs several important considerations, however. For one, the number of sequences is so huge that any way of extracting useful features becomes even more critical. Since the spike protein is the entry point of the virus to the host cell, it is an important characterizing feature of a coronavirus [24], [25].

Most of the variants of SARS-CoV-2 are characterized by mutations which happen disproportionately in the spike region of the genome [10], [15], [16]. Even the mRNA vaccines (*e.g.*, Pfizer and Moderna) for COVID-19 are designed to target only the SARS-CoV-2 spike protein [26] (unlike traditional vaccines which comprise an entire virome). Since the spike region is sufficient to characterize most of the important features of a viral sample, yet is much smaller in length, as depicted in Figure 1, we focus on an embedding approach tailored to the spike region of the sequences.




Fig. 1: The SARS-CoV-2 genome is composed of ≈ 30 Kb nucleotide base pairs, which codes for several proteins, including the spike protein. The region of the genome which corresponds to the spike protein is composed of 3821 (25,384 – 21,563) nucleotide base pairs, hence $3821 (+ 1 \text{ stop codon } *) / 3 = 1274$ amino acids.

Previously, some efforts have been done to perform classification and clustering of SARS-CoV-2 spike sequences [27]–[29]. However, those methods are not scalable to the amount of data we use in this study. Although they were successful in getting high predictive accuracy, it is not clear if the proposed methods are robust and will give the same predictive performance on larger datasets. In this paper, we propose Spike2Vec, an efficient and scalable feature vector generation approach for SARS-CoV-2 spike sequences, to which we can apply different machine learning tasks downstream, such as classification and clustering. Our contributions in this paper are as follows:

- 1) We propose an embedding approach, called Spike2Vec that outperforms the baseline classification method in terms of predictive accuracy.
- 2) We show that our method is scalable on larger datasets by using ≈ 2.5 million spike sequences.
- 3) We prove from the results that the machine learning models used in [27]–[29] are not scalable on these larger datasets. This robust checking helps us to analyze the machine learning models in detail in terms of their appropriateness for SARS-CoV-2 spike sequences.
- 4) We also show that in terms of clustering, our embedding approach is better than the baseline model.

The rest of the paper is organized as follows: Section II contains a discussion on the previous studies related to our research problem. Section III contains a detailed description of our Spike2Vec approach. Section IV contains the implementation details of the experimental evaluation of Spike2Vec, along with the dataset statistics and discussion of the baseline models. We present and discuss the results of this experimental evaluation in Section V. Finally, we conclude our paper in Section VI.

II. LITERATURE REVIEW

Because of the rapid spread of COVID-19 since December 2019, volumes of sequence data are available for this virus. This new source of information has attracted researchers from all fields to perform analysis on this data to better understand the diversity and dynamics of this virus. With the number of sequences in the millions, this is far out of reach for approaches based (at least purely) on phylogenetic reconstruction, which can handle at most thousands [30], or tens of thousands [31] of sequences. Researchers have hence recently turned to clustering and other machine learning (ML) approaches as an alternative to studying patterns in this data. Authors in [29] propose a one-hot encoding based approach to classify different coronavirus hosts using the spike portion of the virus rather than the entire sequence, obtaining near-optimal prediction accuracy. Ali *et al.* in [27] perform classification of different variants of the human SARS-CoV-2. Although they were successful in achieving higher accuracy than in [29], the kernel method used in their approach, however, is not scalable to the size of the data we use in this study. This drawback makes it difficult to use this approach in real-world scenarios such as the current scenario.

Supervised and unsupervised feature selection methods such as ridge regression [32], lasso regression, and principal component analysis (PCA) [33] are very popular for not only reducing the runtime, but also for improving the predictive performance of the underlying machine learning algorithms. Authors in [28] perform clustering on SARS-CoV-2 spike sequences and show that clustering performance could be improved by simply using lasso and ridge regression. Although they were also able to get significant improvement in terms of clustering quality as compared to the baseline, using feature selection methods like ridge regression and lasso regression scales very poorly on the larger datasets, such as the one we use in this study. Melnyk *et al.*, in [34] perform clustering of the entire SARS-CoV-2 genome (rather than just the spike sequence) using CliqueSNV [35], a method originally designed for identifying haplotypes in an intra-host viral population. Although they obtained good overall F_1 scores, our (clustering) approach tends to obtain better overall F_1 scores. It would be interesting to know whether that is because of our feature vector representation, or because we leverage more (and more up-to-date) data, or both.

Existing work on fixed length numerical representation of the data successfully perform different data analytics tasks. It has applications in different domains such as graphs [36], [37], nodes in graphs [38], [39], and electricity consumption [33], [40]. This vector-based representation also achieves significant success in sequence analysis, such as texts [41]–[43], electroencephalography and electromyography sequences [44], [45], networks [46], and biological sequences [32], [47]. However, most of the existing sequence classification methods require the input sequences to be aligned. Although sequence alignment help to analyze the data better, it is a very costly process. Several efforts have been made previously to identify

the transmission patterns of different variants that can help the appropriate authorities to devise appropriate public health interventions so that the rapid spread of viruses can be avoided [48]–[51].

Farhan *et al.*, in [52] propose an efficient approach to compute a similarity matrix (kernel matrix). The computed kernel matrix is proven to be efficient for sequence classification. However, since their approach requires to save an entire $n \times n$ dimensional kernel matrix (where n is the total number of sequences), this makes their method expensive in terms of space. Authors in [53] use the random feature method to map the original input into a low dimensional feature space so that the inner product of the low dimensional data is approximately equal to the inner product of the original data points. Peng *et al.*, in [54] use the random feature attention model for text classification. Their approach is linear in terms of runtime and space, and uses random feature methods to approximate the softmax function.

While dealing with *Big Data*, it is important to analyze the trade-off between the prediction accuracy and the runtime [55]. Although Ali *et al.*, in [27] use the kernel method for spike sequence classification, since the kernel computation is, however, expensive in terms of time and space, their approach is only a proof of concept, and not feasible in a real-world scenario.

III. PROPOSED APPROACH

In this section, we give a step by step description of the Spike2Vec approach. From the SARS-CoV-2 spike sequences, we first generate k -mers so that we can preserve some ordering information of the sequences. It is interesting to note that *de novo* genome assembly (when no reference is present) involves inferring sequence order by assembling k -mers obtained from short reads [56]. After the k -mers are generated, to convert the alphabetical information of k -mers into a fixed length numerical representation (so that ML algorithms can be applied), we generate frequency vectors, which count the number of occurrences of each k -mer in the spike sequence. We then map the high dimensional frequency vectors to a low dimensional embedding using an approximate kernel approach. Each step of Spike2Vec is explained in detail below.

A. k -mers Generation

The first step of Spike2Vec is to compute all k -mers for the spike sequences. The main idea behind using k -mers is to allow some ordering information of the sequence to be preserved. The total number of k -mers, which we can be generated from a spike sequence of length N is $N - k + 1$, where $N = 1274$ and k is a user-defined parameter for the size of each mer. See Figure 2 for an example. In our experiments we use $k = 3$; this was decided using a standard validation set approach [57].

B. Frequency Vectors Generation

Since k -mers is an alphabets-based representation of a spike sequence, we need to convert the k -mers into a numerical




Fig. 2: Example of 4-mers of the amino acid sequence “MD-PEGRKMLSV”.

representation. Therefore, we design a feature vector that contains the count of each k -mer in its respective spike sequence. Each sequence A is over an alphabet Σ . Note that alphabets in our dataset represent amino acids of the spike sequence. These fixed length frequency vectors have length $|\Sigma|^k$ — the number of possible k -mers of a spike sequence. Since the total number of alphabets in our data are 21 (the number of amino acids), the length of each frequency vector becomes $21^3 = 9261$.

C. Low Dimensional Representation

In large scale machine learning (ML) tasks such as classification and clustering, typical supervised and unsupervised dimensionality reduction methods such as principal component analysis, ridge regression, and lasso regression, etc., are not suitable because they take a lot more time to execute. Therefore, in a real world scenario where we can have a huge amount of data, the scalability of any underlying algorithm could be one of the major issues. One option is to use kernel based algorithms that compute a similarity matrix which can later be used for the underlying ML tasks. To compute the kernel matrix (gram matrix), the kernel trick is used.

Definition 3.1 (Kernel Trick): The Kernel Trick is used to generate features for an algorithm which depends on the inner product between only the pairs of input data points. The main idea is to avoid the need to map the input data (explicitly) to a high-dimensional feature space.

The Kernel Trick relies on the following observation: *Any positive definite function $f(a,b)$, where $a, b \in \mathcal{R}^d$, defines an inner product and a lifting ϕ so that we can quickly compute the inner product between the lifted data points* [53]. More formally:

$$\langle \phi(a), \phi(b) \rangle = f(a, b) \quad (1)$$

The major drawback of kernel methods is that in case of large training data, the kernel method suffers from large initial computational and storage costs.

To overcome these computational and storage problems, we use an approximate kernel method called random Fourier features (RFF) [53], which maps the input data to a randomized low dimensional feature space (euclidean inner product space). More formally:

$$z : \mathcal{R}^d \rightarrow \mathcal{R}^D \quad (2)$$

In this way, we approximate the inner product between a pair of transformed points. More formally:

$$f(a, b) = \langle \phi(a), \phi(b) \rangle \approx z(a)'z(b) \quad (3)$$

In Equation (3), z is low dimensional (unlike the lifting ϕ). In this way, we can transform the original input data with z , which acts as the approximate low dimensional embedding for the original data. This low dimensional representation is then used as an input for different ML tasks like classification and regression.

IV. EXPERIMENTAL EVALUATION

We now detail the experiments we performed to evaluate Spike2Vec in terms of both the downstream classification and clustering results obtained.

A. Experimental Setup

All experiments are conducted using an Intel(R) Xeon(R) CPU E7-4850 v4 @ 2.10GHz having Ubuntu 64 bit OS (16.04.7 LTS Xenial Xerus) with 3023 GB memory. Implementation of Spike2Vec is done in Python and the code is available online for reproducibility². Our pre-processed data is also available online³, which can be used after agreeing to terms and conditions of GISAID [9]. For the classification algorithms, we use 10% of the data for training and 90% for testing. The purpose of using smaller training dataset is to show how much performance gain we can achieve while using minimal training data. Note that our data split and pre-processing follow those of [27].

B. Dataset Statistics

We used the (aligned) amino acid sequences corresponding to the spike protein from the largest known database of SARS-CoV-2 sequences, GISAID [9]. In our dataset, we have 2,519,386 spike sequences along with the COVID-19 variant information (in our data, we have 1327 variants in total) for each spike sequence. The information about some of the more well-represented variants is given in Table I. Since most of the variants are new, we do not have all the information available for all them. Therefore, we put “-” in any field of Table I for which we do not have any information available.

Figure 3 shows the total number of spike sequences for the top 10 countries worldwide. In our GISAID dataset, a total of 219 countries are represented. Since the USA has the highest number of spike sequences, we use it in a case study to analyze the spread patterns of different variants in Section IV-C1.

C. Data Visualization

To see if there is any (hidden) clustering in the data, we mapped the data to 2D real vectors using the t-distributed stochastic neighbor embedding (t-SNE) approach [61]. Since it was not possible to run the t-SNE algorithm on all ≈ 2.5 million spike sequences, we obtained a representative subset

²<https://github.com/sarwanpasha/Spike2Vec>

³<https://drive.google.com/drive/folders/1-YmIM8ipFpj-gl9hSF3t6VuofrpgWUa?usp=sharing>

Pango Lin.	Region	Labels	No. Mut. S/Gen.	No. sequences
B.1.1.7	UK [10]	Alpha	8/17	976077
B.1.1.351	South Africa [10]	Beta	9/21	20829
B.1.617.2	India [16]	Delta	8/17	242820
P.1	Brazil [15]	Gamma	10/21	56948
B.1.427	California [58]	Epsilon	3/5	17799
AY.4	India [59]	Delta	-	156038
B.1.2	-	-	-	96253
B.1	-	-	-	78741
B.1.177	-	-	-	72298
B.1.1	-	-	-	44851
B.1.429	-	-	-	38117
AY.12	India [59]	Delta	-	28845
B.1.160	-	-	-	25579
B.1.526	New York [60]	Iota	6/16	25142
B.1.1.519	-	-	-	22509
B.1.1.214	-	-	-	17880
B.1.221	-	-	-	13121
B.1.258	-	-	-	13027
B.1.177.21	-	-	-	13019
D.2	-	-	-	12758
B.1.243	-	-	-	12510
R.1	-	-	-	10034

TABLE I: The SARS-CoV-2 variants which were represented in more than 10,000 sequences (of the ≈ 2.5 million sequences). The S/Gen. column represents the number of mutations in the Spike (S) region / entire genome. The total number of amino acid sequences in our dataset is 2,519,386. The variants listed in this table comprise 1,995,195 sequences.




Fig. 3: Country-wise distribution (for the top 10 countries) of spike sequences.

of sequences containing 7000 randomly selected sequences such that the proportion of each variant in this subset is equal to its proportion in the original data. The t-SNE plot for Delta, Beta, Iota, Epsilon, and Gamma variants is shown in Figure 4.




Fig. 4: t-SNE embeddings of spike sequences

1) *USA Case Study*: Figure 5 shows the COVID-19 spread pattern for three variants in the USA from March 2020 to

July 2021. We can see in Figure 5 that after the coronavirus spread hit its peak in April 2021, the number of cases of the coronavirus started decreasing. That was the point where a significant proportion of the population of the USA was vaccinated (hence peak spread reduced).




Fig. 5: Spread pattern of Alpha (blue line), Delta (orange line), and Gamma (black line) variants in the USA from March 2020 to July 2021. The y-axis shows the total number of COVID-19 infected patients.

To evaluate Spike2Vec, we perform classification and clustering on the low dimensional feature vectors that it produces. For the classification task, we use naive Bayes (NB), logistic regression (LR), and ridge classifier (RC) [62]. For the clustering analysis, we use the k -means algorithm. We report accuracy, precision, recall, weighted F_1 , macro F_1 , and ROC-AUC for classification and weighted F_1 score (with the major variant in a cluster as the label) for clustering evaluation. For k -means, we use 22 as the number of clusters. We selected this number of clusters using the elbow method [28]. For this, we perform clustering with different numbers of clusters from 2 to 100 and then see the trade-off between the *sum of squared error* (distortion score) and the runtime. After analyzing this trade-off, we use the knee point detection algorithm (KPDA) [63] to find the optimal value of k for the k -means.

D. Baseline Algorithm

As a baseline, we use the one-hot encoding (OHE) method [29]. In spike sequences, we have 21 unique amino acids (unique alphabets forming Σ) namely “ACDEFGHIKLMNPQRSTVWXY”. Also, the length of each spike sequence is 1273 plus a terminating character * at the 1274th location. After getting the OHE for each spike sequence, we get a feature vector of length 26,733 corresponding to each spike sequence ($21 \times 1273 = 26,733$). We then use RFF on the OHEs to get the low dimensional feature vector representation (to avoid the curse of dimensionality).

V. RESULTS AND DISCUSSION

In this section, we present the classification and clustering results for Spike2Vec and its comparison with OHE and report the importance of each amino acid using information gain.

A. Classification Results

Results for different classification algorithms are shown in Table III. Note that overall logistic regression is a clear

Variant	k -means (Cluster IDs)																					
	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21
Epsilon	109	0	3186	432	113	219	67	0	160	9	134	0	18	14	0	78	9	3792	113	0	48	41
Alpha	6061	1	175923	23353	5846	11754	3376	0	9466	1041	6889	0	734	1281	0	4160	329	205848	5730	0	3193	2136
Gamma	344	0	10403	1312	327	686	205	0	534	63	390	0	52	77	0	254	14	11977	324	0	182	137
Beta	144	0	3853	436	115	237	64	0	191	8	148	0	19	25	0	81	7	4435	119	0	71	44
Delta	1432	1	43691	5732	1391	2832	831	0	2342	241	1777	0	172	315	0	1016	77	51596	1400	0	836	541

TABLE II: Contingency tables of variants vs. clusters.

winner in case of Spike2Vec. All of the classifiers in case of Spike2Vec clearly outperform the corresponding classifiers with OHE. This performance for different evaluation metrics shows the effectiveness of using k -mers instead of OHE for representing the spike sequences. Also, we can observe that although the performance of RC in the case of Spike2Vec is not better than LR, it is significantly better than LR and NB, however, in terms of training runtime. Therefore, we can conclude that overall LR is better in terms of prediction performance with Spike2Vec while RC is better in terms of runtime along with comparable performance to LR.

Approach	ML Algo.	Acc.	Prec.	Recall	F_1 (Weig.)	F_1 (Macro)	ROC-AUC	Training time (sec.)
OHE	NB	0.30	0.58	0.30	0.38	0.17	0.59	566.09
	LR	0.56	0.49	0.56	0.49	0.19	0.57	1309.06
	RC	0.56	0.47	0.56	0.48	0.17	0.56	110.76
Spike2Vec	NB	0.42	0.79	0.42	0.52	0.39	0.68	457.54
	LR	0.68	0.68	0.68	0.64	0.49	0.69	830.63
	RC	0.67	0.68	0.67	0.62	0.44	0.67	95.73

TABLE III: Variants Classification Results (10% training set and 90% testing set) for the top 22 variants (1995195 spike sequences) listed in Table I. Best values are shown in bold.

B. Clustering Results

We also test the performance of Spike2Vec using the k -means clustering method. The contingency table for k -means (for some of the more well represented variants) computed using Spike2Vec is given in Table II. Since we have a total of 1327 variants, it is not possible to show results for all variants. Therefore, we only present results for some of the common variants. The weighted F_1 score for OHE and Spike2Vec are shown in Table IV. We can observe that Spike2Vec clearly outperforms OHE in case of all but one variant. The reason for the poor performance in the case of Beta and Epsilon variants is due to the fact that they are in comparatively smaller proportion in the dataset (see Table I). Because of this, Spike2Vec is not able to design a rich feature vector representation of these variants.

Methods	F_1 Score (Weighted) for Different Variants				
	Alpha	Beta	Delta	Gamma	Epsilon
OHE	0.0410	0.0479	0.5942	0.6432	0.0571
Spike2Vec	0.9997	0.0300	0.8531	0.9680	0.2246

TABLE IV: F_1 scores for five variants from the k -means clustering algorithm on all 1327 variants (2519386 spike sequences) in the GISAID dataset. Best values are in bold.

C. Importance of each Amino Acid

To evaluate the importance of amino acid a in a set V of variants, we compute the information gain $IG(V, a) = H(V) - H(V|a)$, where $H = \sum_{v \in V} -p_v \log p_v$. Note that H is the entropy, and p_v is the probability of the class $v \in V$. We extracted a sample dataset of 20000 spike sequences and computed IG (see Figure 6). The USA based Centers for Disease Control and Prevention (CDC) catalogs mutations that take place in different variants [59]. We compare the mutation information from the CDC with our (high) IG values. According to the CDC, R452L is present in Epsilon and Delta lineages and sub-lineages while K417N, E484K, and N501Y substitutions are present in the Beta variant. Note that R452L means that amino acid at position 452 mutated from 'R' to 'L'. Similarly, K417T, E484K, and N501Y substitutions are present in the Gamma variant [59]. We can see in Figure 6 that we obtained the maximum IG values for many of the same amino acids positions mentioned by CDC.




Fig. 6: Information gain for each amino acid position (x-axis) with respect to the variants.

VI. CONCLUSION

We propose an embedding approach that can be used to perform different machine learning tasks on the SARS-CoV-2 spike sequences. We show that our model can scale to several million sequences, and it also outperforms the baseline models significantly. Since the COVID-19 disease is relatively new, we do not have enough information available for different coronavirus variants so far. We will explore the new (and existing) variants in more detail in the future. We will also use deep learning models to enhance the prediction performance of Spike2Vec.

A. Acknowledgments

This research was supported by a Georgia State University startup grant.

REFERENCES

- [1] R. Sharma, A. Mateush, and J. Übi, "Tale of three states: analysis of large person-to-person online financial transactions in three baltic countries," in *IEEE BigData*, 2019, pp. 1497–1505.
- [2] Y. Dong, D. Yan, A. I. Almudaifer, S. Yan, Z. Jiang, and Y. Zhou, "BELT: A pipeline for stock price prediction using news," in *IEEE BigData*, 2020, pp. 379–410.
- [3] J. Kremer, K. Stensbo-Smidt, F. Gieseke, K. S. Pedersen, and C. Igel, "Big universe, big data: Machine learning and image analysis for astronomy," *IEEE Intelligent Systems*, vol. 32, no. 2, pp. 16–22, 2017.
- [4] M. Mace, "Comet NEOWISE sizzles as it slides by the sun, providing a treat for observers," *Infrared Processing and Analysis Center*, 2020.
- [5] C. K. Leung, Y. Chen, C. S. H. Hoi, S. Shang, and A. Cuzzocrea, "Machine learning and olap on big covid-19 data," in *IEEE BigData*, 2020.
- [6] C. K. Leung, Y. Chen, S. Shang, and D. Deng, "Big data science on covid-19 data," in *IEEE BigData*, 2020.
- [7] S. Ali, Y. Zhou, and M. Patterson, "Efficient analysis of covid-19 clinical data using machine learning models," *arXiv preprint arXiv:2110.09606*, 2021.
- [8] Z. Tayebi, S. Ali, and M. Patterson, "Robust representation and efficient feature selection allows for effective clustering of sars-cov-2 variants," *arXiv preprint arXiv:2110.09622*, 2021.
- [9] "GISAID," <https://www.gisaid.org/>, accessed: 2021-09-06.
- [10] S. E. Galloway, P. Paul *et al.*, "Emergence of SARS-CoV-2 b.1.1.7 lineage," *Morbidity and Mortality Weekly Report*, vol. 70, no. 3, p. 95, 2021.
- [11] S. ZD, L. SY, F. F. C. RH, Z. C. E. MJ *et al.*, "Big data: Astronomical or genomic?," *PLoS Biology*, vol. 13, no. 7, p. e1002195, 2015.
- [12] Genome Web, <https://www.genomeweb.com/infectious-disease/cdc-commits-90m-create-public-health-pathogen-genomics-research-centers>, [Online; accessed 5-September-2021].
- [13] Reuters, <https://www.reuters.com/article/us-china-genomics-state/chinese-state-fund-invests-in-gene-firm-bgi-idUSKBN2AM0AT>, [Online; accessed 5-September-2021].
- [14] A. Rambaut, E. C. Holmes, A. O'Toole, V. Hill, J. T. McCrone, C. Ruis, L. du Plessis, and O. G. Pybus, "A dynamic nomenclature proposal for SARS-CoV-2 lineages to assist genomic epidemiology," *Nature Microbiology*, vol. 5, no. 11, pp. 1403–1407, 2020.
- [15] F. Naveca *et al.*, "Phylogenetic relationship of SARS-Cov-2 sequences from amazons with emerging brazilian variants harboring mutations e484k and n501y in the spike protein," *Virological.org*, vol. 1, 2021.
- [16] P. Yadav *et al.*, "Neutralization potential of covishield vaccinated individuals sera against b.1.617.1," *bioRxiv*, vol. 1, 2021.
- [17] E. Volz, S. Mishra, M. Chand, J. C. Barrett, R. Johnson, L. Geidelberg, W. R. Hinsley, D. J. Laydon, G. Dabrera, A. O'Toole *et al.*, "Assessing transmissibility of sars-cov-2 lineage b.1.1.7 in england," *Nature*, vol. 593, pp. 266–269, 2021.
- [18] M. McCallum, J. Bassi, A. Marco, A. Chen, A. Walls, J. Iulio, M. Tortorici, M. Navarro, C. Silacci-Fregni, C. Saliba, M. Agostini, D. Pinto, K. Culap, S. Bianchi, S. Jaconi, E. Cameroni, J. Bowen, S. Tilles, M. Pizzuto, S. Guastalla, G. Bona, A. Pellanda, C. Garzoni, W. Van Voorhis, L. Rosen, G. Snell, A. Telenti, H. Virgin, L. Piccoli, D. Corti, and D. Veessler, "Sars-cov-2 immune evasion by variant b.1.427/b.1.429," vol. 1, 2021, doi:10.1101/2021.03.31.437925.
- [19] Y. Furuse, "Genomic sequencing effort for sars-cov-2 by country during the pandemic," *International Journal of Infectious Diseases*, vol. 103, pp. 305–307, 2021. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1201971220325571>
- [20] A. Maxmen, "Why US coronavirus tracking can't keep up with variants," *Nature*, vol. 592, 2021.
- [21] F. Wu, S. Zhao, B. Yu, Y.-M. Chen, W. Wang, Z.-G. Song, Y. Hu, Z.-W. Tao, J.-H. Tian, Y.-Y. Pei *et al.*, "A new coronavirus associated with human respiratory disease in china," *Nature*, vol. 579, pp. 265–269, 2020.
- [22] F. Wu, A. Xiao, J. Zhang, K. Moniz, N. Endo, F. Armas, M. Bushman, P. R. Chai, C. Duvallet, T. B. Erickson, K. Foppe, N. Ghaeli, X. Gu, W. P. Hanage, K. H. Huang, W. L. Lee, K. A. McElroy, S. F. Rhode, M. Matus, S. Wuertz, J. Thompson, and E. J. Alm, "Wastewater surveillance of sars-cov-2 across 40 u.s. states from february to june 2020," *Water Research*, vol. 202, p. 117400, 2021. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0043135421005984>
- [23] J. Gardy and N. Loman, "Towards a genomics-informed, real-time, global pathogen surveillance system," *Nature Reviews Genetics*, vol. 19, pp. 9–20, 2018.
- [24] F. Li, "Structure, function, and evolution of coronavirus spike proteins," *Annual Review Virology*, vol. 3, no. 1, pp. 237–261, 2016.
- [25] A. Walls, Y. Park, and M. Tortorici, "Structure, function and antigenicity of the sars-cov-2 spike glycoprotein," *Cell*, vol. 181, no. 2, pp. 281–292, 2020.
- [26] K. S. Corbett, D. K. Edwards, S. R. Leist *et al.*, *Nature*, vol. 586, pp. 567–571, 2020.
- [27] S. Ali, B. Sahoo, N. Ullah, A. Zelikovskiy, M. Patterson, and I. Khan, "A k-mer based approach for sars-cov-2 variant identification," *To Appear at: International Symposium on Bioinformatics Research and Applications (ISBRA)*, 2021.
- [28] S. Ali, M. A. Khan, I. Khan, M. Patterson *et al.*, "Effective and scalable clustering of SARS-CoV-2 sequences," *To appear in: International Conference on Big Data Research (ICBDR)*, 2021.
- [29] K. Kuzmin *et al.*, "Machine learning methods accurately predict host specificity of coronaviruses based on spike sequences alone," *Biochemical and Biophysical Research Communications*, vol. 553, no. 3, pp. 553–558, 2020.
- [30] J. Hadfield, C. Megill, S. Bell, J. Huddleston, B. Potter, C. Callender, P. Sagulenko, T. Bedford, and R. Neher, "Nextstrain: real-time tracking of pathogen evolution," *Bioinformatics*, vol. 34, pp. 4121–4123, 2018.
- [31] B. Minh, H. Schmidt, O. Chernomor, D. Schrempf, M. Woodhams, A. von Haeseler, and R. Lanfear, "IQ-TREE 2: New models and efficient methods for phylogenetic inference in the genomic era," *Molecular Biology and Evolution*, vol. 37, pp. 1530–1534, 2020.
- [32] S. Ali, S. Ciccolella, L. Lucarella, G. Della Vedova, and M. D. Patterson, "Simpler and faster development of tumor phylogeny pipelines," *To appear at: Journal of Computational Biology (JCB)*, 2021.
- [33] S. Ali, H. Mansoor, N. Arshad, and I. Khan, "Short term load forecasting using smart meter data," in *International Conference on Future Energy Systems*, 2019, pp. 419–421.
- [34] A. Melnyk, F. Mohebbi, S. Knyazev, B. Sahoo, R. Hosseini, P. Skums, A. Zelikovskiy, and M. D. Patterson, "From alpha to zeta: Identifying variants and subtypes of sars-cov-2 via clustering," *To appear at: Journal of Computational Biology (JCB)*, 2021.
- [35] S. Knyazev, V. Tsyvina, A. Shankar, A. Melnyk, A. Artyomenko, T. Malygina, Y. B. Porozov, E. M. Campbell, S. Mangul, W. M. Switzer *et al.*, "Accurate assembly of minority viral haplotypes from next-generation sequencing through efficient noise reduction," *Nucleic Acids Research*, p. 264242, 2020.
- [36] Z. R. Hassan, M. Shabbir, I. Khan, and W. Abbas, "Estimating descriptors for large graphs," in *Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD)*, 2020, pp. 779–791.
- [37] Z. R. Hassan, I. Khan, M. Shabbir, and W. Abbas, "Computing graph descriptors on edge streams," *CoRR*, vol. arXiv:2109.01494, 2021.
- [38] S. Ali, M. H. Shakeel, I. Khan, S. Faizullah, and M. A. Khan, "Predicting attributes of nodes using network structure," *ACM Transactions on Intelligent Systems and Technology (TIST)*, vol. 12, no. 2, pp. 1–23, 2021.
- [39] A. Grover and J. Leskovec, "node2vec: Scalable feature learning for networks," in *International Conference on Knowledge Discovery & Data Mining (KDD)*, 2016, pp. 855–864.
- [40] S. Ali, H. Mansoor, I. Khan, N. Arshad, M. A. Khan, and S. Faizullah, "Short-term load forecasting using AMI data," *arXiv preprint arXiv:1912.12479*, 2019.
- [41] M. H. Shakeel, A. Karim, and I. Khan, "A multi-cascaded model with data augmentation for enhanced paraphrase detection in short texts," *Information Processing & Management*, vol. 57, no. 3, p. 102204, 2020.
- [42] M. H. Shakeel, S. Faizullah, T. Alghamidi, and I. Khan, "Language independent sentiment analysis," in *International Conference on Advances in the Emerging Computing Technologies (AECT)*, 2020, pp. 1–5.
- [43] M. H. Shakeel, A. Karim, and I. Khan, "A multi-cascaded deep model for bilingual sms classification," in *International Conference on Neural Information Processing*, 2019, pp. 287–298.
- [44] M. Atzori, A. Gijsberts, C. Castellini, B. Caputo, A.-G. M. Hager, S. Elsig, G. Giatsidis, F. Bassetto, and H. Müller, "Electromyography data for non-invasive naturally-controlled robotic hand prostheses," *Scientific Data*, vol. 1, no. 1, pp. 1–13, 2014.
- [45] A. Ullah, S. Ali, I. Khan, M. A. Khan, and S. Faizullah, "Effect of analysis window and feature selection on classification of hand

- movements using EMG signal,” in *Intelligent Systems Conference*, 2020, pp. 400–415.
- [46] S. Ali, M. K. Alvi, S. Faizullah, M. A. Khan, A. Alshantiti, and I. Khan, “Detecting ddos attack on sdn due to vulnerabilities in openflow,” in *International Conference on Advances in the Emerging Computing Technologies (AECT)*, 2020, pp. 1–6.
- [47] P. P. Kuksa, I. Khan, and V. Pavlovic, “Generalized similarity kernels for efficient sequence classification,” in *SIAM International Conference on Data Mining (SDM)*, 2012, pp. 873–882.
- [48] M. Ahmad, J. Tariq, M. Farhan, M. Shabbir, and I. Khan, “Who should receive the vaccine?” in *Australasian Data Mining Conference (AusDM)*, 2016, pp. 137–145.
- [49] M. Ahmad, J. Tariq, M. Shabbir, and I. Khan, “Spectral methods for immunization of large networks,” *Australasian Journal of Information Systems*, vol. 21, 2017.
- [50] J. Tariq, M. Ahmad, I. Khan, and M. Shabbir, “Scalable approximation algorithm for network immunization,” in *Pacific Asia Conference on Information Systems (PACIS)*, 2017, p. 200.
- [51] M. Ahmad, S. Ali, J. Tariq, I. Khan, M. Shabbir, and A. Zaman, “Combinatorial trace method for network immunization,” *Information Sciences*, vol. 519, pp. 215–228, 2020.
- [52] M. Farhan, J. Tariq, A. Zaman, M. Shabbir, and I. U. Khan, “Efficient approximation algorithms for strings kernel based sequence classification,” in *Advances in neural information processing systems (NeurIPS)*, 2017, pp. 6935–6945.
- [53] A. Rahimi, B. Recht *et al.*, “Random features for large-scale kernel machines,” in *NIPS*, vol. 3, no. 4, 2007, p. 5.
- [54] H. Peng, N. Pappas, D. Yogatama, R. Schwartz, N. A. Smith, and L. Kong, “Random feature attention,” in *International Conference on Learning Representations (ICLR)*, 2021.
- [55] S. Ali, “Cache replacement algorithm,” *arXiv preprint arXiv:2107.14646*, 2021.
- [56] R. Rizzi, S. Beretta, M. Patterson, Y. Pirola, M. Previtali, G. D. Vedova, and P. Bonizzoni, “Overlap graphs and de Bruijn graphs: data structures for de novo genome assembly in the big data era,” *Quantitative Biology*, vol. 7, no. 4, pp. 278–292, 2019.
- [57] P. A. Devijver and J. Kittler, *Pattern recognition: A statistical approach*. Prentice hall, 1982.
- [58] W. Zhang *et al.*, “Emergence of a novel SARS-Cov-2 variant in Southern California,” *Jama*, vol. 325, no. 13, pp. 1324–1326, 2021.
- [59] SARS-Cov-2 Variant Classifications and Definitions, “” <https://www.cdc.gov/coronavirus/2019-ncov/variants/variant-info.html>, Accessed: 08-01-2021.
- [60] A. West Jr *et al.*, “Detection and characterization of the SARS-Cov-2 lineage b. 1.526 in new york,” *bioRxiv*, 2021.
- [61] L. Van der Maaten and G. Hinton, “Visualizing data using t-SNE.” *Journal of Machine Learning Research (JMLR)*, vol. 9, no. 11, 2008.
- [62] A. Singh, B. S. Prakash, and K. Chandrasekaran, “A comparison of linear discriminant analysis and ridge classifier on twitter data,” in *2016 International Conference on Computing, Communication and Automation (ICCCA)*, 2016, pp. 133–138.
- [63] V. Satopaa, J. Albrecht, D. Irwin, and B. Raghavan, “Finding a” kneedle” in a haystack: Detecting knee points in system behavior,” in *International conference on distributed computing systems workshops*. IEEE, 2011, pp. 166–171.