Analysis of B-cell receptor repertoires in COVID-19 patients using deep embedded representations of protein sequences

Inyoung Kim[1,*], Sang Yoon Byun[2], Sangyeup Kim[3], Sangyoon Choi[4], Jinsung Noh[1], Junho Chung[6,7,8], and Byung Gee Kim[1,5]

[1]Bio-MAX Institute, Seoul National University, Seoul, Republic of Korea
[2]Department of Computer Science, Grinnell College, Grinnell, IA, United States of America
[3]Department of Applied Biology and Chemistry, Seoul National University, Seoul, Republic of Korea
[4]College of Veterinary Medicine, Seoul National University, Seoul, Republic of Korea
[5]School of Chemical and Biological Engineering, Seoul National University, Seoul, Republic of Korea
[6]Department of Biomedical Science, Seoul National University College of Medicine, Seoul, Republic of Korea
[7]Department of Biochemistry and Molecular Biology, Seoul National University College of Medicine, Seoul, Republic of Korea
[8]Cancer Research Institute, Seoul National University College of Medicine, Seoul, Republic of Korea

Analyzing B-cell receptor (BCR) repertoires is immensely useful in evaluating one's immunological status. Conventionally, repertoire analysis methods have focused on comprehensive assessment of clonal compositions, including V(D)J segment usage, nucleotide insertion/deletion, and amino acid distribution. Here, we introduce a novel computational approach that applies deep-learning based protein embedding techniques to analyze BCR repertoires. By selecting the most frequently occurring BCR sequences in a given repertoire and computing the sum of the vector representations of these sequences, we represent an entire repertoire as a 100-dimensional vector and eventually as a single data point in vector space. We demonstrate that our new approach enables us to not only accurately cluster repertoires of COVID-19 patients and healthy subjects, but also efficiently track minute changes in immunity conditions as patients undergo a course of treatment over time. Furthermore, using the distributed representations, we successfully trained an XGBoost classification model that achieved over 87% mean accuracy rate given a repertoire of CDR3 sequences.

Correspondence: ***iykim@snu.ac.kr***

Introduction

Analyzing and deciphering biological sequences plays a critical role in gaining a deeper understanding of biological systems. Recent major advances in artificial intelligence have sparked ample interests in adopting natural language processing (NLP) models to extract hidden insights from biological sequences [1]. By reinterpreting protein sequences as sentences and k-mers in these sequences as words, researchers have succeeded in establishing computational methods to represent the language of life. One of the most widely used techniques in the field is Words2Vec, an efficient method to learn and compute word embeddings, which are essentially vectorized representations of words [2].

ProtVec has been the first bioinformatics application to use such word-vector model to embed biological sequences—more specifically, amino acid sequences [3]. The representation, which treats amino acid 3-mers as discrete words and employs the skip-gram model architecture, has been rigorously trained from over 546,790 unique protein sequences from the Swiss-Prot database. Ultimately, ProtVec is highly efficient in transforming an amino acid sequence into a 100-dimensional vector.

B-cell receptors (BCR) are transmembrane proteins on the surface of B-cells and are crucial in producing antibodies that recognize and neutralize myriad foreign objects like antigens [4]. To bind to a wide range of antigens, BCRs most often undergo genetic recombination and somatic hypermutation (SHM), ultimately generating a more diversified pool of paratopes, also known as antigen-binding sites. Each paratope comprises six complementarity-determining regions (CDRs), three from each of the antibody's heavy and light chains. Moreover, CDRs can be categorized into three distinct sections (CDR1, CDR2, and CDR3), where CDR3 is commonly deemed the most variable, and hence the most critical region for recognizing antigens. Therefore, this study will primarily focus on heavy-chain CDR3s when discussing CDRs.

A BCR repertoire refers to a diverse population of BCRs in a given individual or entity. Analyzing BCR repertoires is highly useful in determining and evaluating the overall condition of one's immune system [5]. For example, a healthy individual will generally possess a vastly diverse repertoire of unique sequences to recognize a variety of antigens. A virus-infected individual will similarly have a wide variety of sequences, but it will also most likely have an aberrant increase in the number of certain BCR sequences due to an activation of specific B-cells to bind to the invading antigens [6]. In such ways, BCR repertoire analysis provides a holistic assessment of one's adaptive immune system as well as immune diversity.

Previous studies on BCR repertoires have primarily focused on classifying and clustering BCRs based on its sequence motifs in an attempt to identify neutralizing antibodies against specific viral infections [7–10]. One of the most common and conventional methods for studying BCRs in a wet lab is flow cytometry, which aims to identify antigen-specific B-cells by labelling antigens with fluorescent tags. Microscopy and immunohistochemical staining are other examples of common techniques that are used to detect partic-

ular antigens.

Recently, several computational methods have also been established. Schulteis, for example, employed the GLIPH2 algorithm to cluster T-cell receptor (TCR) sequences and the OLGA algorithm to compute CDR3 sequence similarities from TCR/BCR repertoires of COVID-19 patients [11]. Bashford-Rogers developed a method that utilizes network analysis to compare a wide range of BCR repertoires [12]. Sidhom very recently proposed DeepTCR, a deep learning model that enables enhanced classification of antigen-specific TCRs from human and murine datasets [13].

In this study, we present a novel computational approach for representing BCR repertoires using deep-learning based protein embedding methods. By utilizing ProtVec to embed all BCR amino acid sequences and summing the vectorized values of the most frequently occurring CDR3 sequences in a given repertoire, we represent individual BCR repertoires as single 100-dimensional vectors. Our representation shows promising results in clustering, classifying, and even tracking BCR repertoires of COVID-19 patients and healthy individuals in vector space. Compared to conventional techniques, this computational method provides an entirely new approach for researchers to represent and analyze BCR repertoires.

Methods

Embedding amino acid sequences using ProtVec.

ProtVec regards amino acid 3-mers as 'biological words' and effectively embeds each of such words (over 9,048 words in total) into a 100-dimensional vector. As shown in Fig. 1, an original sequence is first split into three separate lists of non-overlapping 3-mers. The 3-mers in these lists are then converted into vectors based on the pretrained data available at Harvard Dataverse [3]. Finally, by summing all the vector representations of the non-overlapping 3-mers, we obtain a single 100-dimensional vector that represents a single protein sequence. The pretrained data file available at Harvard Dataverse is a comma-separated values (CSV) file that includes 100 float values assigned to each of the 9,048 3-mers. The model had been trained from 546,790 sequences of the Swiss-Prot database using the Skip-gram Word2Vec architecture. The CSV file was used as a look-up table to search for 100-dimensional vectors based on a given 3-mer.

Pre-processing the data acquired from OAS.

The raw data acquired from the Observed Antibody Space (OAS) BCR dataset included heavy and light chains as well as several different isotypes. We decided to use sequences from immunoglobulin heavy G chain (IGHG) due to its strong association with SARS-CoV-2 [7]. Additionally, the raw data did not include read counts of unique CDR3 sequences, but rather read counts of unique BCR sequences. As a result, we added a pre-processing step that counted the number of unique CDR3 sequences in a given repertoire.

Visualizing sequences and repertoires using t-SNE.

Visualizing the embedded representations (sequences and/or repertoires) on a 2-dimensional vector space requires reducing the 100-dimensional ProtVec data to 2-dimensional x, y coordinates so that one can ultimately generate a scatter plot. We utilized the t-Distributed Stochastic Neighbor Embedding technique (TSNE implementation from scikit learn) in reducing the dimensions. For all the t-SNE charts shown here, we set the specific hyper-parameters to 'perplexity=30.0' and 'learning_rate=100'.

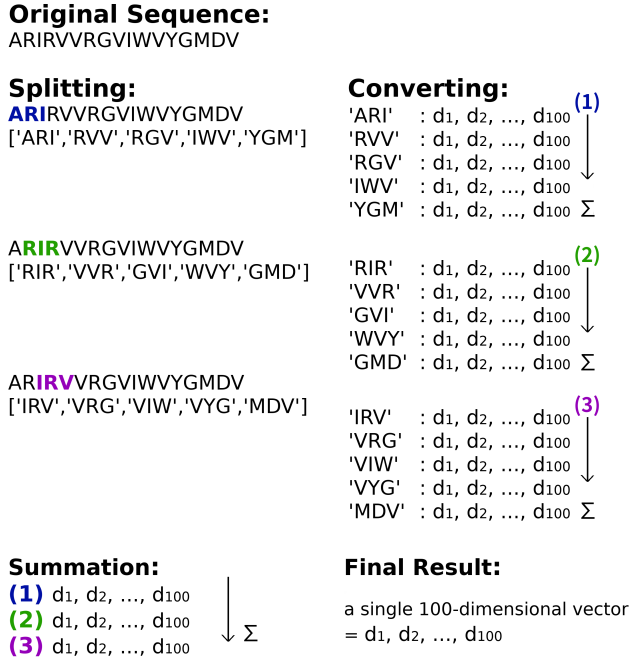

Fig. 1. Embedding protein sequences using ProtVec. An illustration of embedding an amino acid sequence using ProtVec. This technique enables one to express a protein sequence as a single 100-dimensional vector while retaining its unique physical/chemical properties.

Training binary classification models.

In training the models, we utilized the SVM implementation from scikit learn library (version 0.22.2.post1) and XGBClassifier from the XGBoost Python package (version 0.90). For the SVM model that classified the two selected protein families (Fig. 2), we used 1,528 embedded representations of amino acid sequences from each family (3,056 in total). For the XGBoost model that classified COVID-19 patients and healthy subjects (Fig 4b), we used 106 embedded representations of BCR repertoires from each of the two classes (212 in total). We evaluated both models using 5-fold cross validation (StratifiedKFold implementation from scikit-learn) and ROC/AUC computation (scikit-learn). All sequences and/or repertoire data that were involved in training, testing, and validating the models were embedded using ProtVec.

Results

Assessment of the efficacy of ProtVec representation.

To ensure that ProtVec can indeed accurately and precisely represent protein sequences, we initially assessed the efficacy of the embedding technique on a public dataset provided by Google AI [14]. Details on the specific protein fam-

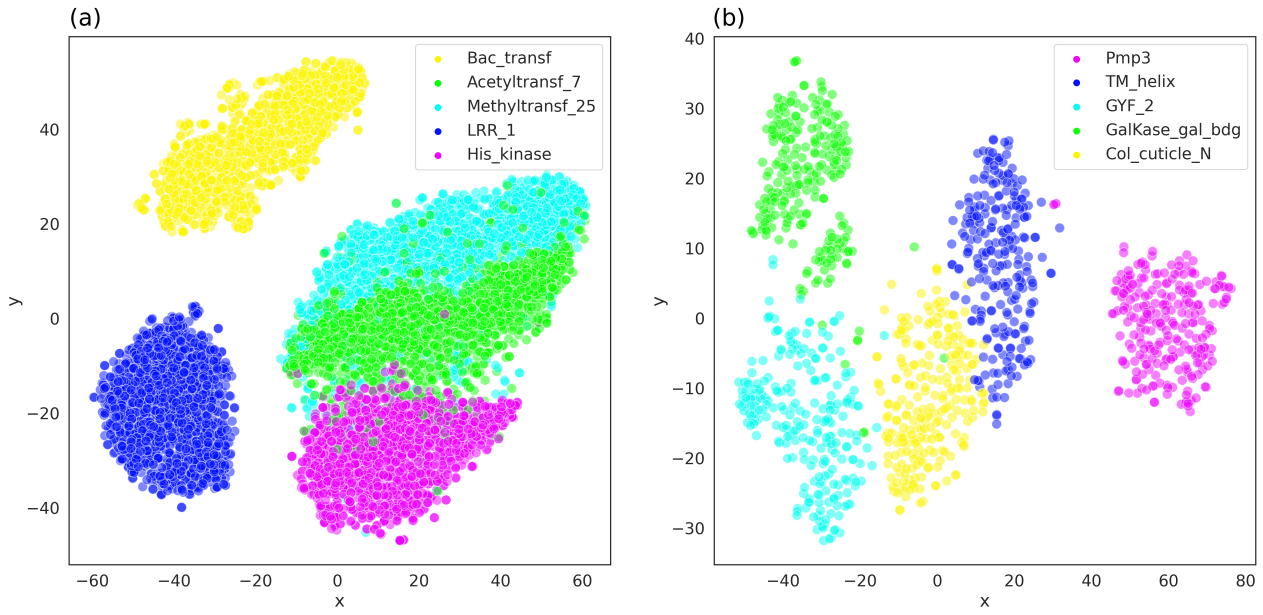

Fig. 2. Visualization of embedded protein sequences. Result of embedding the acquired protein sequences and visualizing them on a 2- dimensional vector space using t-SNE. **(a).** Visualization of amino acid sequences from five protein families that had the largest number of sequences in the dataset. A total of 7,640 sequences (1,528 sequences from 5 families) were used. Detailed list of selected protein families, including their corresponding family ID, the number of sequences, and family accession numbers, is summarized in Table 1a. **(b).** Visualization of sequences from five similar (in terms of mean sequence length and standard deviation) protein families. A total of 1,210 sequences (242 sequences for 5 families) were used. Detailed information about the protein families is summarized in Table 1b.

ilies and the number of sequences used for the assessment are shown in Table 1. After embedding each of the protein sequences as 100-dimensional vectors, we used t-Distributed Stochastic Neighbor Embedding (t-SNE) to reduce their dimensions. We observed that the protein families of interest were successfully clustered into five distinct classes when plotted in 2-dimensional vector spaces (Fig. 2a and 2b).

(a) Five protein families with the largest number of sequences

Family ID	No. of CDR3s	Family Accession
Methyltransf_25	3,637	PF13649.6
LRR_1	1,927	PF00560.33
Acetyltransf_7	1,927	PF13508.7
His_kinase	1,537	PF06580.13
Bac_transf	1,528	PF02397.16

(b) Five protein families with similar mean sequence length

Family ID	No. of CDR3s	Accession	Mean Length
GalKase_gal_bdg	448	PF10509	49.82
Col_cuticle_N	242	PF01484	49.41
TM_helix	247	PF05552	49.38
GYF_2	307	PF14237	49.16
Pmp3	286	P87284	49.14

Table 1. Protein families. Details regarding family ID, the number of CDR3 sequences, family accession numbers, and mean sequence lengths of the selected sequences from the Google AI public protein dataset.

For Fig. 2a, we noticed that two classes, Methyltransf_25 and Acetyletransf_7, overlapped each other in the t-SNE chart, possibly hinting that t-SNE was unable to distinguish the two families. When we used Support Vector Machines (SVM) to further analyze the two families, however, we obtained a classification model that successfully classified the two families with a 97% mean accuracy rate and an AUC value of 0.99. Thus, we eventually came to a conclusion that the overlap was a result of some negligible limitation that may have occurred during the process of dimensionality reduction.

More specifically, for the five protein families with similar lengths, we first examined the mean lengths of protein sequences from each of the protein families and ultimately selected five families that had similar mean lengths (49 amino acids). We also searched for families that had low standard deviations to ensure that sequences from the selected families did not vary widely in lengths. For Fig. 2b, which plotted sequences from five protein families similar in mean sequence lengths, we observed five distinct clusters with minimal overlap.

The distribution of frequencies of CDR3 read counts.

Before closely analyzing the repertoires, we first briefly explored the arrangements of few different repertoires to gain a better understanding of the distribution of CDR3 sequences. Fig. 3a illustrates the histogram of CDR3 amino acid sequences from a BCR repertoire of a COVID-19 patient (Subject-A_TP2; Kim et al.) and that of a healthy subject (700010756; Roskin et al.). In comparing the two histograms, we observed that the patient repertoire (purple) possessed more unique CDR3 sequences in general and espe-

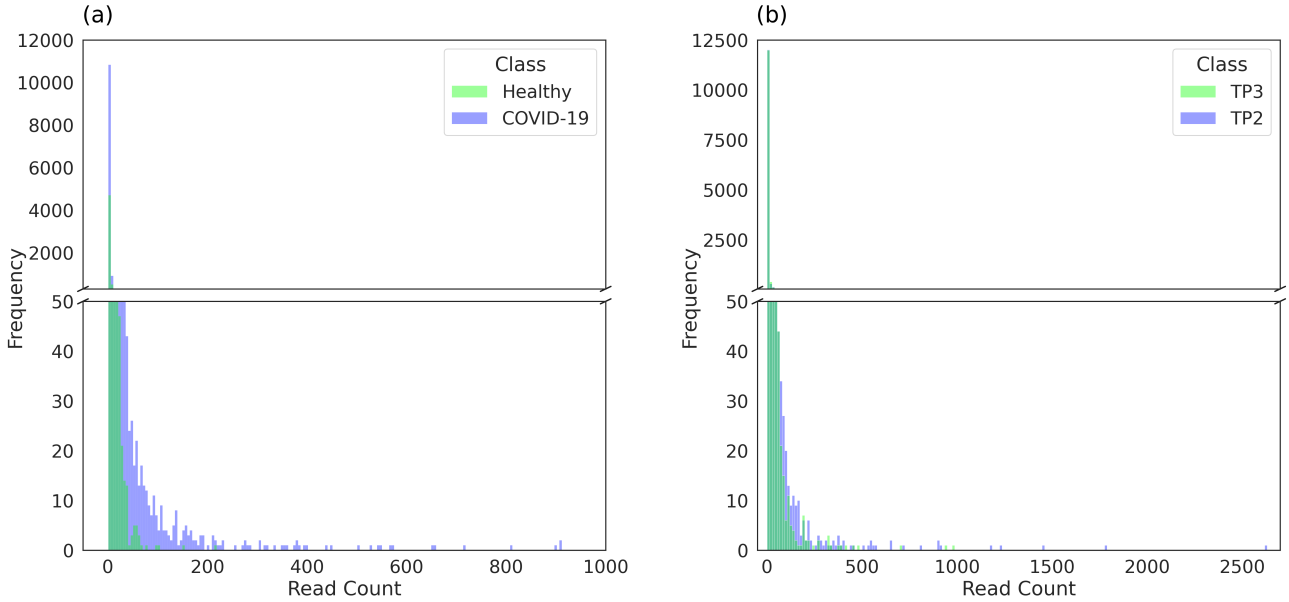

Fig. 3. The distribution of CDR3 read count. The distribution of frequencies of CDR3 sequence read counts comparing a patient BCR repertoire with a healthy repertoire. Read count (x-axis) represents the number of times a specific CDR3 sequence appears within a repertoire, and frequency (y-axis) represents the number of times the corresponding read count occurred. **(a)**. A histogram (bins = 2,000) of repertoires from a COVID-19 patient and a healthy subject. For the sake of clear visualization, we only displayed sequences with read counts up to 1,000. More details can be found in Table 2. **(b)**. A histogram (bins = 200) of repertoires from Subject A's two time points (from a study by Kim et al.). TP2 is when the subject experienced severe symptoms and where TP3 is when the subject was fully recovered.

cially many more CDR3 sequences with higher read counts compared to the healthy repertoire (green). While the healthy repertoire did have a sequence with a read count of 9,901, which was abnormally high, all other sequences in the repertoire had significantly lower read counts than the read counts recorded in the patient repertoire (Table 2). Conversely, the patient repertoire had a sequence with the highest read count of 2,631, and the succeeding four sequences also had read counts higher than 1,000, suggesting that sequences in the patient repertoire generally has a much higher read count.

Furthermore, Fig. 3b compares two repertoires from a single COVID-19 patient, Subject-A. Subject-A's TP2 (severe symptoms) repertoire represents the patient repertoire, and TP3 (fully recovered) repertoire represents the healthy repertoire. Once again, we observed that the TP2 repertoire had more CDR3 sequences with significantly higher read counts. The highest read count of the TP2 repertoire was 2,631, whereas that of the TP3 repertoire was 979 (Table 2). In short, both Fig. 3a and 3b highlight that the patient repertoires contain more CDR3 sequences with higher read counts.

The visualization of individual sequences using t-SNE.

Now using our selected repertoires from OAS, we visualized each individual CDR3 sequence on a t-SNE chart [15]. We first extracted and embedded the ten most frequently occurring sequences from each repertoire (106 patient repertoires and 106 healthy repertoires), and simply projected each vectorized sequence on to a 2-dimensional vector space using t-SNE (Fig. 4). As shown in Fig. 4a, the sequences did not seem to have any clear patterns or form distinct clusters within the vector space.

Index	Read counts		
	Subject-A_TP2 (Patient)	Subject-A_TP3 (Healthy)	700010756 (Healthy)
0	2,631	979	9,901
1	1,777	375	215
2	1,459	700	153
3	1,236	475	102
4	1,177	457	97
5	909	416	80
6	908	386	70
7	899	377	64
8	811	340	63
9	716	319	62

Table 2. CDR3 Read counts from patient and healthy subject. Details regarding the ten highest CDR3 read counts recorded in each repertoire. For example, the most frequently occurring CDR3 sequence in Subject-A_TP2's repertoire appeared 2,631 times in total. Subject-A_TP2 represents a patient repertoire while Subject-A_TP3 and 700010756 represent healthy repertoires. The read counts were sorted in descending order, and the table displays the first ten of them.

The visualization of individual repertoires using t-SNE.

Rather than plotting individual sequences (Fig. 4), we decided to plot sets of sequences altogether, which, as we have previously discussed, represent BCR repertoires. We

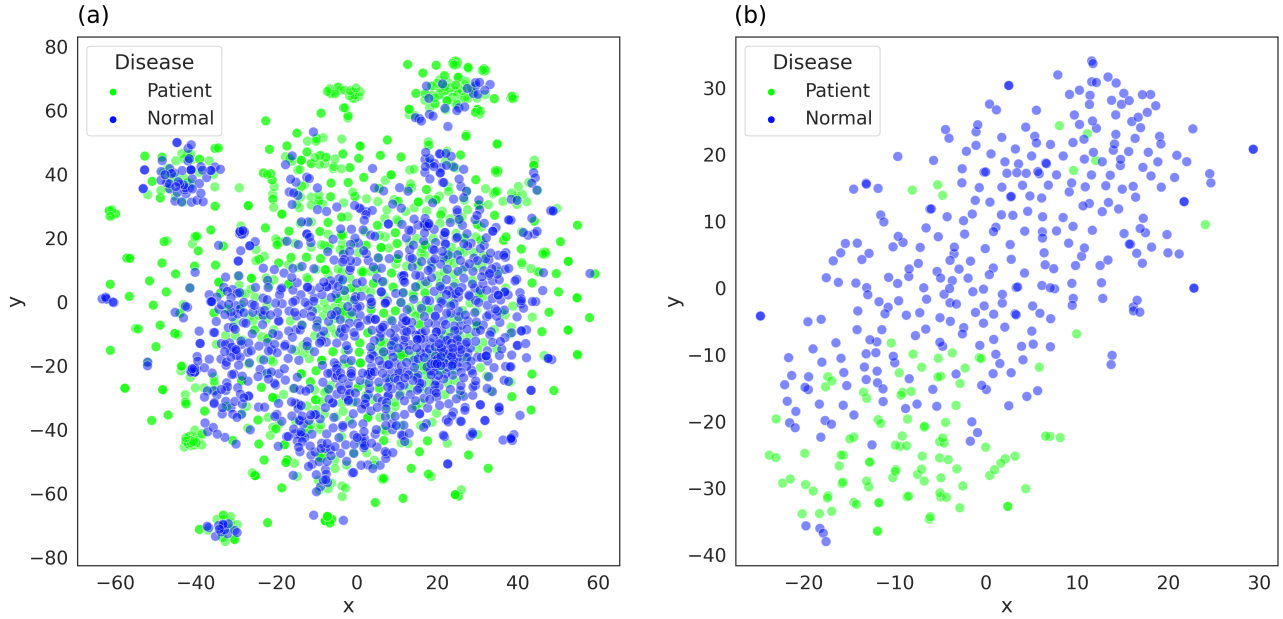

Fig. 4. Visualization of embedded CDR3 sequences and BCR repertoire. (a). Result of embedding and visualizing individual CDR3 sequences of COVID-19 patients and healthy subjects using t-SNE. A total of 2,120 sequences (top 10 frequent CDR3 sequences from 1,060 patients and 1,060 healthy subjects) were used. Every data point in the figure represents an individual CDR3 sequence—blue and green representing healthy subjects and COVID-19 patients, respectively. **(b).** Result of embedding and visualizing individual BCR repertoires of COVID-19 patients and healthy subjects using t-SNE. In both figures, a total of 428 repertoires (106 from COVID-19 patients and 322 from healthy subjects) were plotted. Details regarding specific datasets and studies used are summarized in Table S1.

assumed that CDR3 sequences that occur more frequently within a repertoire carry more significance to COVID-19 viral infections. We thus sliced the 100 most frequently occurring CDR3 amino acid sequences from each repertoire and computed the sum of the vector representations of these sequences. This method allowed us to represent an entire repertoire as a single 100-dimensional vector. When we visualized the newly attained vector representations using t-SNE, we were able to successfully distinguish between 106 patient repertoires and 322 healthy repertoires (Fig. 4b).

In Fig. 4b, we observed that the repertoires of COVID-19 patients gathered largely on the bottom left corner, while those of healthy subjects clustered more towards the upper right region of the figure. There were, however, a few patient repertoires among the healthy repertoires. The positions of these repertoires and possible explanations will be discussed later on. Additionally, using the embedded representations of repertoires, we trained a binary classification model with XGBoost. The model achieved 87% mean accuracy rate in classifying COVID-19 patients and healthy subjects given a 100-dimensional vector representation of a repertoire.

Tracking the spatial and temporal movements of COVID-19 patient repertoires.

Using the same public BCR dataset from OAS, we once again sliced the 100 most frequently occurring CDR3 amino acid sequences from each repertoire to represent repertoires as 100-dimensional vectors. The t-SNE visualizations of these repertoire representations are shown in Fig. 5. This time, we decided to focus on tracking the spatial and temporal movements of specific repertoires from Kim's study [7].

We discovered that the repertoires can be largely categorized into two main groups—disease and healthy clusters. The disease cluster, as shown in Fig. 5a, is primarily composed of repertoires of COVID-19 patients 9 to 15 days after symptoms onset. On the other hand, the healthy cluster, as captured in Fig. 5b, mainly consists of repertoires of healthy subjects as well as patients 0 to 9 days and/or 21 to 45 days after symptoms onset.

Focusing on Subject-A and Subject-E, both of whom were COVID-19 patients with three time points, we observed that their repertoires had remarkably interesting movements based on their time points, suggesting a potential association between our t-SNE charts and Kim's chronological analysis of IGH repertoires of the two patients [7]. Kim, in his study, highlighted that antibody clonotypes that were reactive against SARS-CoV-2 receptor-binding domains (RBD) underwent swift class switching with minimal somatic mutations and divergence from the germ line. Isotypes and subtypes with low divergence thus suggest that the antibodies are actively reacting to the viral infection.

From Kim's study, Subject-A's IgG1 and IgG3 both had significantly low divergence on day 11 and even lower values on day 17, but ultimately had a large increase by day 45. These observations indicate that during a certain period of time, most likely around day 11 and 17, subject's IgG1/3 antibodies have potentially bound to the RBD of SARS-CoV-2 and neutralized the virus. An increased divergence on day 45

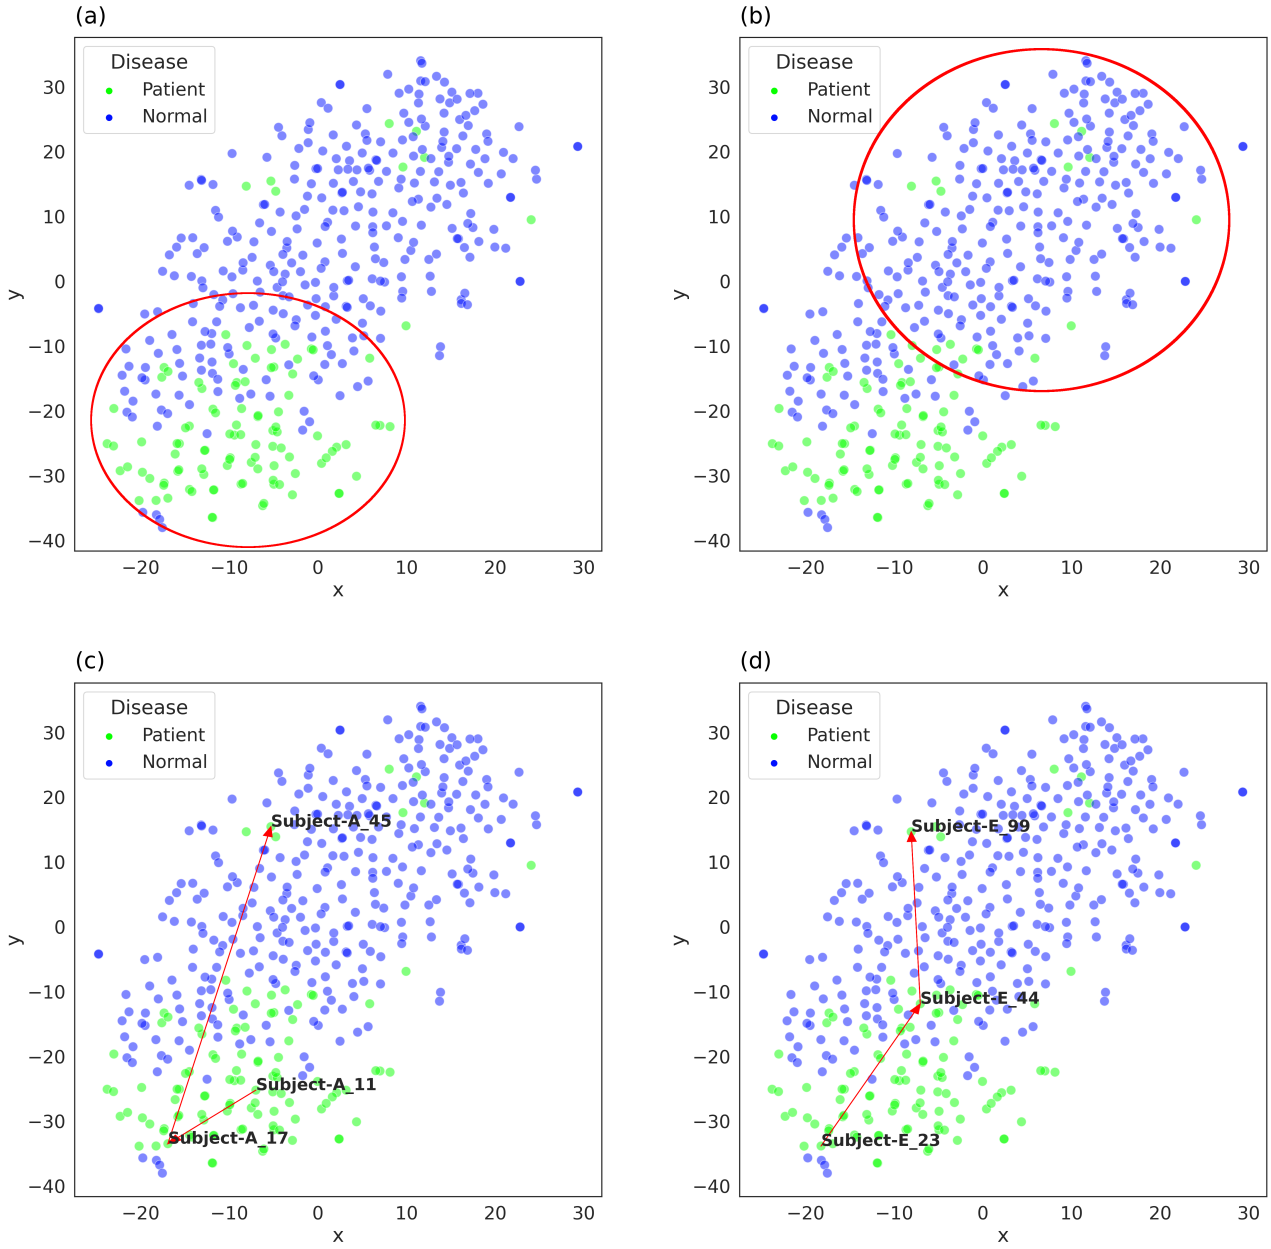

Fig. 5. Visualization of embedded repertoire. Visualization of embedded repertoire representations of COVID-19 patients and healthy subjects from a study by Kim et al. **(a).** Repertoires of COVID-19 patients 9 to 15 days after symptoms onset formed their own cluster on the upper-left corner of the chart. The red circle indicates the disease cluster. **(b).** Repertoires of COVID-19 patients 0 to 9 days and/or 21 to 45 days after symptoms onset. They were gathered at the bottom-right corner. The red circle indicates the healthy cluster. **(c).** Movement of Subject-A's repertoires within the vector space over three time points (day 11, 17, and 45 after symptoms onset). **(d).** Movement of Subject-E's repertoires within the vector space over three time points (day 23, 44, and 99 after symptoms onset).

further suggests that the antibodies were no longer reacting as strongly as they did before. We observed a surprisingly similar pattern in our t-SNE chart (Fig. 5c). Subject-A was initially located near the center of the disease cluster on day 11 (A_TP1). By day 17 (A_TP2), the subject has moved farther towards the disease cluster and away from the healthy cluster. However, by day 45 (A_TP3), the repertoire has moved to the other side, near the center of the healthy cluster.

Similarly, Kim's study shows a gradual increase in divergence for Subject-E's IgG1 and IgG3 across the three time points (day 23, 44, and 99). Over this time period, the reactivity of subject's IgG1/3 antibodies has potentially subsided after initially reacting strongly to the SARS-CoV-2 virus. Moreover, compared to Subject-A, Subject-E experienced a steadier and gradual increase in divergence. These observations align well with the t-SNE visualization (Fig. 5d). Subject-E was inside the disease cluster on day 23 (E_TP1), but as time passed, the subject moved closer towards the

healthy cluster (day 44; E_TP2) and eventually rested safely within the healthy cluster (day 99; E_TP3).

Discussion

In this study, we introduced a novel computational approach of utilizing ProtVec, a deep-learning based embedding technique for protein sequences, to represent a BCR repertoire as a 100-dimensional vector. More specifically, we discovered that by slicing the most frequently occurring CDR3 sequences and using the summation of vector representations of these sequences, we were not only able to effectively represent a subject's entire BCR repertoire as a single data point inside a vector space, but also track the repertoire as its immunological status changed over time.

Firstly, we assessed the efficacy of ProtVec to ensure that we can apply the embedding technique to represent BCR sequences in our study. As demonstrated in Fig. 2, the two t-SNE charts confirmed how effective ProtVec can be in representing different protein families. These selected protein families were grouped into their respective clusters remarkably well, suggesting that ProtVec is more than capable of capturing unique features and characteristics from each protein family. It is especially important to note that ProtVec was also able to successfully distinguish protein families similar in mean amino acid sequence lengths (Fig. 2b). Moreover, we found that the embedded representations can have other useful applications like training a supervised learning model. A simple binary classification model we built using ProtVec representations was highly capable of accurately classifying protein families. Based on these results, we were confident that ProtVec would be an apposite embedding method for our BCR repertoire analysis.

The distribution of frequencies of CDR3 amino acid sequences provided an insightful overview of the data we were handling (Fig. 3). A single BCR repertoire is typically composed of thousands, if not tens of thousands, of unique BCR amino acid sequences, each with its own frequency or read count [6]. The histogram of the CDR3 read counts revealed that while the vast majority of the sequences have low read counts, it was not uncommon to discover several sequences with abnormally high read counts. Fig. 3a compared repertoires from two different subjects, both in their 50s and with similar number of unique BCR sequences. Fig. 3b, on the other hand, compared two repertoires from a single subject. Both figures ultimately revealed that patient repertoires, when compared to healthy repertoires, possessed a much greater number of unique CDR3 sequences with significantly higher read counts.

We deduced that such differences can be associated with affinity maturation of B cells. Patients suffering from a viral infection most often experience a rapid increase in the read counts of particular sequences due to an activation of B-cells when their BCRs bind to specific antigens [5]. Therefore, the CDR3 sequences on the right-hand side of Fig. 3a and 3b were most likely the sequences that burgeoned as a result of having higher affinity to these antigens. We thus decided to focus on small selections of (100 most frequently appearing) BCR sequences because we hypothesized that they are the sequences that have been rapidly expanded by the activation and hence possess much more significance in comparing COVID-19 patients and healthy subjects.

Our first attempt at using ProtVec to analyze BCR repertoires involved visualizing individual BCR amino acid sequences in a 2-dimensional vector space (Fig. 4). However, as Fig. 4 shows, the embedded sequences did not form any meaningful patterns or clusters that distinguished one class from another. While there were small, indistinct clusters around the largest clump, because they included sequences from both patients and healthy subjects, we deduced that their clustering was a result of other trivial external factors. We eventually arrived at a conclusion that we cannot possibly cluster or classify COVID-19 patients and healthy subjects when we solely focus on individual CDR3 amino acid sequences.

Afterwards, we represented entire repertoires by using the summation of vector representations of the 100 most frequently occurring CDR3 sequences (Fig. 5). The results highlight that this method can accurately distinguish between COVID-19 patients and healthy subjects. This discovery is especially meaningful because although we have previously demonstrated that ProtVec can capture unique characteristics of a BCR sequence (Fig. 2), there was no practical way for ProtVec to represent BCR repertoires themselves. Our new method of summing up these vector representations, however, has shown to capture unique features and characteristics of an entire repertoire, allowing us to visualize and analyze a repertoire in vector space. The ability to represent repertoires in low dimensions further enables us to track a subject's repertoire based on their time points.

As shown in Fig. 5a and 5b, visualizing repertoires of all subjects from all time points generates two main clusters (disease and healthy clusters). When we focused on two particular patients, Subject-A and Subject-E, we were able to track the movements of their repertoires over the course of three time points. We discovered that the timeline and relative positions of subjects' repertoires—whether they were located in the disease cluster or healthy cluster—aligned very well with the time points in which IGH clonotypes that were reactive against SARS-CoV-2 RBD have been observed [7]. When the subject's IgG1 and IgG3 had low divergence from the germ line, indicating that the antibodies were reacting against the virus, the repertoire was located within the disease cluster. On the other hand, when subject's IgG1 and IgG3 had relatively higher divergence, the repertoire was back within the healthy cluster. This result is immensely enlightening because it suggests that our repertoire representations can go beyond merely expressing a BCR repertoire in vector space and further help us examine a subject's protective immunity.

In conclusion, our novel approach enables researchers to computationally express and analyze subject's BCR repertoires within a vector space. More importantly, it allows us to pinpoint one's immunological status and track their spatial and temporal movement within vector space as patients undergo a course of treatment and recovery. We expect our

unconventional approach of utilizing protein embedding in repertoire analysis to bring new possibilities for further BCR repertoire research in the near future.

Bibliography

1. Hitoshi Iuchi, Taro Matsutani, Keisuke Yamada, Natsuki Iwano, Shunsuke Sumi, Shion Hosoda, Shitao Zhao, Tsukasa Fukunaga, and Michiaki Hamada. Representation learning applications in biological sequence analysis. *Computational and Structural Biotechnology Journal*, 19:3198, 2021.
2. Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. *arXiv*, 2013.
3. Ehsaneddin Asgari and Mohammad RK Mofrad. Continuous distributed representation of biological sequences for deep proteomics and genomics. *PloS one*, 10(11):e0141287, 2015.
4. J. Punt, S. Stranford, P. Jones, and J.A. Owen. *Kuby Immunology*. Macmillan Learning, 2018. ISBN 9781319114756.
5. Ofir Lindenbaum, Nima Nouri, Yuval Kluger, and Steven H Kleinstein. Alignment free identification of clones in b cell receptor repertoires. *Nucleic acids research*, 49(4):e21–e21, 2021.
6. Jacob D Galson, Sebastian Schaetzle, Rachael JM Bashford-Rogers, Matthew IJ Raybould, Aleksandr Kovaltsuk, Gavin J Kilpatrick, Ralph Minter, Donna K Finch, Jorge Dias, Louisa K James, et al. Deep sequencing of b cell receptor repertoires from covid-19 patients reveals strong convergent immune signatures. *Frontiers in immunology*, 11:3283, 2020.
7. Sang Il Kim, Jinsung Noh, Sujeong Kim, Younggeun Choi, Duck Kyun Yoo, Yonghee Lee, Hyunho Lee, Jongtak Jung, Chang Kyung Kang, Kyoung-Ho Song, et al. Stereotypic neutralizing vh antibodies against sars-cov-2 spike protein receptor binding domain in patients with covid-19 and healthy individuals. *Science translational medicine*, 13(578), 2021.
8. Huang Huang, Chunlin Wang, Florian Rubelt, Thomas J Scriba, and Mark M Davis. Analyzing the mycobacterium tuberculosis immune response by t-cell receptor clustering with gliph2 and genome-wide antigen screening. *Nature biotechnology*, pages 1–9, 2020.
9. Aleksandr Kovaltsuk, Matthew IJ Raybould, Wing Ki Wong, Claire Marks, Sebastian Kelm, James Snowden, Johannes Trück, and Charlotte M Deane. Structural diversity of b-cell receptor repertoires along the b-cell differentiation axis in humans and mice. *PLoS computational biology*, 16(2):e1007636, 2020.
10. Zachary Montague, Huibin Lv, Jakub Otwinowski, William S DeWitt, Giulio Isacchini, Garrick K Yip, Wilson W Ng, Owen Tak-Yin Tsang, Meng Yuan, Hejun Liu, et al. Dynamics of b-cell repertoires and emergence of cross-reactive responses in patients with different severities of covid-19. *Cell Reports*, page 109173, 2021.
11. Christoph Schultheiß, Lisa Paschold, Donjete Simnica, Malte Mohme, Edith Willscher, Lisa von Wenserski, Rebekka Scholz, Imke Wieters, Christine Dahlke, Eva Tolosa, et al. Next-generation sequencing of t and b cell receptor repertoires from covid-19 patients showed signatures associated with severity of disease. *Immunity*, 53(2):442–455, 2020.
12. RJM Bashford-Rogers, Laura Bergamaschi, EF McKinney, DC Pombal, Federica Mescia, JC Lee, DC Thomas, SM Flint, P Kellam, DRW Jayne, et al. Analysis of the b cell receptor repertoire in six immune-mediated diseases. *Nature*, 574(7776):122–126, 2019.
13. John-William Sidhom, H Benjamin Larman, Drew M Pardoll, and Alexander S Baras. Deeptcr is a deep learning framework for revealing sequence concepts within t-cell repertoires. *Nature communications*, 12(1):1–12, 2021.
14. Google Research. Kaggle pfam seed random split, 2019.
15. Aleksandr Kovaltsuk, Jinwoo Leem, Sebastian Kelm, James Snowden, Charlotte M Deane, and Konrad Krawczyk. Observed antibody space: a resource for data mining next-generation sequencing of antibody repertoires. *The Journal of Immunology*, 201(8):2502–2509, 2018.