Deep metric learning improves lab of origin prediction of genetically engineered plasmids

Igor Soares (✉ igor.muniz.ims@gmail.com)

Amalgam <https://orcid.org/0000-0003-1629-806X>

Fernando Camargo

Amalgam

Adriano Marques

Amalgam <https://orcid.org/0000-0001-8795-9035>

Oliver Crook

Oxford Protein Informatics Group, Department of Statistics, Oxford University

Article

Keywords: lab-of-origin, genetic engineering attribution, metric learning, triplet network, deep learning, dna, rna, plasmid

Posted Date: December 1st, 2021

DOI: <https://doi.org/10.21203/rs.3.rs-1072312/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Deep metric learning improves lab of origin prediction of genetically engineered plasmids

FERNANDO H. F. CAMARGO* and IGOR MUNIZ SOARES*, Amalgam, USA

ADRIANO M. MARQUES, Amalgam, USA

OLIVER M. CROOK, Oxford Protein Informatics Group, UK

Genome engineering is undergoing unprecedented development and is now becoming widely available. To ensure responsible biotechnology innovation and to reduce misuse of engineered DNA sequences, it is vital to develop tools to identify the lab-of-origin of engineered plasmids. Genetic engineering attribution (GEA), the ability to make sequence-lab associations, would support forensic experts in this process. Here, we propose a method, based on metric learning, that ranks the most likely labs-of-origin whilst simultaneously generating embeddings for plasmid sequences and labs. These embeddings can be used to perform various downstream tasks, such as clustering DNA sequences and labs, as well as using them as features in machine learning models. Our approach employs a circular shift augmentation approach and is able to correctly rank the lab-of-origin 90% of the time within its top 10 predictions - outperforming all current state-of-the-art approaches. We also demonstrate that we can perform few-shot-learning and obtain 76% top-10 accuracy using only 10% of the sequences. This means, we outperform the previous CNN approach using only one-tenth of the data. We also demonstrate that we are able to extract key signatures in plasmid sequences for particular labs, allowing for an interpretable examination of the model's outputs.

CCS Concepts: • **Information systems** → **Similarity measures**; **Learning to rank**.

Additional Key Words and Phrases: lab-of-origin, genetic engineering attribution, metric learning, triplet network, deep learning, dna, rna, plasmid

1 INTRODUCTION

Genetic engineering and synthetic biology are fast growing areas of biotechnology. We are now able to transform organisms in highly efficient and sophisticated manners. As this biotechnology becomes more widespread, it is vital that we may attribute genetically engineered organisms to their makers or lab-of-origin. This will prevent plagiarism, encourage responsible development, allow designers to gain due credit and as an approach to holding genetic engineers/designers accountable for their work. Tools for attributing this biotechnology to their owners, often referred to *genetic engineering attribution* (GEA), have only recently become sufficiently well-performing [1, 31, 43]. When making design choices for nucleic-acid sequences, an engineer will impart a design signature which could be detectable by GEA methods. Powerful methods could, in principle, identify the true designer of a biological sequence and hence be excellent tools for accountability.

Several approaches to GEA have now been proposed based on predicting the lab-of-origin of plasmid sequences from the Addgene [18] data repository. The performance of these approaches has quickly improved, rising from 70% top-10 accuracy [31] to 85% top-10 accuracy [2, 43] in recent years. Nielsen and Voigt [31] used Convolutional Neural Networks (CNNs), Alley et al. [2] used Recurrent Neural Networks (RNNs), and Wang et al. [43] used a pan-genome approach, suggesting a variety of methods could perform well on this task. However, further improvements are still

*Both authors contributed equally to this research.

Authors' addresses: Fernando H. F. Camargo, fernando.camargo.ai@gmail.com; Igor Muniz Soares, igor@amalgam.ai/igor.muniz.ims@gmail.com, Amalgam, 11 N Skokie Hwy, Ste 201, Lake Bluff, Illinois, USA, 60044; Adriano M. Marques, adriano@amalgam.ai, Amalgam, 11 N Skokie Hwy, Ste 201, Lake Bluff, Illinois, USA, 60044; Oliver M. Crook, oliver.crook@stats.ox.ac.uk, Oxford Protein Informatics Group, University of Oxford, Oxford, , UK, .

possible, these approach make other downstream tasks challenging and require many training instances to perform well. Here, we present an approach that couples CNNs and metric learning [25] to extract embeddings from DNA sequences, whilst simultaneously learning an embedding of known labs. We also employ circular shift augmentation rather than the typically used reverse-complement augmentation. Together, our method improves over the state-of-the-art by 5 percentage points and allows clustering of sequence and labs for other downstream tasks. Our approach allows us to perform one-shot-learning [24] and so we can predict lab associations with only 1 training instance. Furthermore, using integrated-gradients we can extract design signatures from our model, allowing us to interpret the model outputs.

We evaluate our approach on a dataset from the Addgene repository containing 81,834 DNA sequences along with minimal phenotypic information: antibiotic resistance, copy number, growth temperature, growth strain, selectable markers and in which species (see methods). The plasmids are designed by 3,751 different labs and grouped into 1,313 categories, along with a single additional category to represent "unknown engineered" (see methods). We evaluate the solutions using accuracy and top-10 accuracy metric. Top-10 accuracy means that the model needs to rank the correct lab-of-origin within the ten most likely labs. Ranking often requires a slightly different approach to classification, and so we developed a Metric Learning approach [25] (more specifically, Triplet Networks [15]). We begin the manuscript by outlining our method, demonstrating that it improves over the state-of-the-art. We then demonstrate that the embedding approach allows us to perform other tasks of interest and finally demonstrate how we can perform one/few-shot-learning [9, 10, 24, 45] with our approach.

2 RESULTS

2.1 Metric learning model and model evaluation

Our proposed method uses (deep) metric learning [25], which is where one learns a distance function between objects. Here, this can be thought of learning a similarity between plasmid sequences and labs. The result is an embedding where distances between sequences have their similarity preserved. We use deep learning, in particular a CNN-based approach, to extract embeddings of DNA sequences, while learning the embeddings of the known labs. To demonstrate that using deep metric learning indeed provides an advantage, we also developed a regular classifier with a similar architecture to compare with our deep metric learning approach.

Each model is composed of a CNN with multiple kernels of differing sizes [21]. The CNN is used to extract features from the sequence, which are then concatenated with the phenotypic metadata of that plasmid sequence. The key difference between our two proposed models can be understood by considering the final layers. The classification model's final layer has a softmax activation function, resulting in each lab being given a probability vector associating it with each lab. Instead, our metric learning approach passes these features through a dense layer that generates our sequence embedding. In parallel, we have an embedding layer that learns the lab embeddings. The principle advantage of our metric learning approach is that once it has been trained it can extract embeddings of any DNA sequence, allowing us to group or cluster new sequences to existing labs based on similar characteristics.

Our classification model is trained using regular supervised learning. However, our metric learning approach is trained differently. Specifically, our metric learning method employs the idea of Triplet Networks [15]. Here, we create a triplet (anchors, positives, negatives) as part of the model training process. Our model is anchored around the DNA sequences and hence they are the anchors in this approach. The positive object in this scenario is the true lab-of-origin, whilst the negative object is some other lab. To be concrete suppose s_1 is a plasmid sequence made by the Church lab, then a possible triplet would be $(s_1, \text{Church lab}, \text{Voigt lab})$. Hence, the goal of our approach is to generate embeddings

Fig. 1: Applying Convolutional Neural Networks with Metric Learning on plasmids leads to a new state-of-the-art.

Fig. 1: **a** Genetic sequences and phenotype information are input into our model. All sequences are processed in order to better demonstrate the characteristics of each plasmid and improve the model’s ability to identify patterns. **b** We use convolutional neural networks as the base model for our two approaches. In this model, convolutional operations extract information present in the sequence based on a fixed kernel size. Our method uses convolutional structures with different kernel sizes in parallel, simulating the observation of sequences by pieces of different sizes. The information extracted by each structure is aggregated and added to the phenotype information and then assigned to one of the laboratories. **c** The difference between the two approaches is the training and output of the model. The standard approach treats the genetic attribution problem as a classification problem where a softmax layer is applied to determine the probability that the sequence belongs to each of the laboratories seen by the model in the training phase. On the other hand, a metric learning approach (triplet network) determines how far the features representation of a new sequence is from the sequences cluster of a laboratory in the base. Smaller distances indicate greater similarities between the features of a sequence and a lab-of-origin. **d** Before training, all sequences from a given laboratory are grouped according to their Levenshtein distances. We do not use sequences from the same group in training and validation at the same time, ensuring that sequences that are too close do not cause leakage in training and overfit the model. **e** DNA sequences are compressed by Byte Pair Encoding (BPE) algorithm [11]. It works by looking for common patterns in the sequence and unifying them into tokens, increasing the vocabulary while reducing the sequences’ size. **f** Since plasmids are circular sequences, we randomly shift the starting point of the sequence, increasing our number of training data. This method is performed "online" during sequence loading and preparing for network entry. After all these processing, we limited the sequence size to 1000 characters to optimize the training time and convergence capability of the algorithm. **g** Top-10 prediction accuracy on the test set. We compared our both approaches to Nielsen and Voigt , Alley et al., BLAST baseline [2] and PlasmidHawk [43].

in which the DNA sequences are near their labs-of-origin and far from the sequences of other labs. Using the Addgene dataset (see methods), we trained each approach to perform DNA sequence prediction to one of 1,313 laboratories or "unknown engineered" (see methods). The dataset has a total of 81,834 DNA sequences, where 18,817 of these sequences were separated for testing. During training, we split the data into 85% for algorithm training and 15% for validation. Our complete method can be summarised in Figure 1.

2.2 Metric learning model improves predictions over state-of-the-art

Nielsen and Voigt [31] developed a deep learning model by applying convolutional neural networks. The network was trained on the Addgene plasmid dataset and independently verified in [2]. In brief, their approach can be summarised as follows. First, DNA sequences were one-hot encoded, which were then used as input to the network composed of one convolutional layer of 128 filters, max-pooling operation, and two dense layers. Whilst, they showed it is possible to use machine learning for this task, this seminal approach was hopeful and obtained an accuracy of only 48% and a top-10 accuracy of 70% in predicting the origin lab.

More recently, Alley et al. [2] proposed deteRNNT, a recurrent neural network-based model. The main insight of this approach was to treat the DNA sequence as a text problem, using techniques from the natural language processing field to extract features from the sequence. They tokenized the sequence using Byte Pair Encoding [11], generating larger tokens and decreasing the size of the sequence. These tokens then served as input to a word embedding layer [28] followed by recurrent neural networks [26] [36]. The authors showed that their approach achieves 84.7% top-10 accuracy.

We also compared with BLAST [3], with our test set of 18,817 samples. Despite being a relatively simple tool that employs no modern machine learning and simply finds similar local regions between sequences, it was able to predict source labs by achieving 76.9% top-10 accuracy in our tests: outperforming the approach of Nielsen and Voigt [31]

Fig. 2: The triplet is composed of an anchor (DNA), a positive (the lab-of-origin), and a negative example (another lab). **a** In the beginning, the anchor might be closer to the negative than it is to the positive. During training, we pull the anchor and positive towards each other while pushing the negative away. **b** In the end, labs and their DNA sequences will be nearer to each other, forming groups. We can also expect both labs and DNA sequences to be closer to other similar ones.

PlasmidHawk [43], a recently launched tool, uses Plaster [42], a state-of-the-art pan-genome algorithm to construct a synthetic plasmid resulting in a set of sequence fragments. It then aligns the original plasmid to the synthetic one and makes comparisons to match fragments with the plasmid. This method has so far outperformed other machine learning-based methods by obtaining 85% of top-10 accuracy, whilst employing no machine learning.

Here, we find that our metric learning model alongside our training methodology improves the current state-of-the-art for attributing the lab-of-origin of an engineered DNA sequence, achieving 90.39% top-10 accuracy. For the classic approach of a classification model that predicts the input sequence’s probability from any of the possible labs seen during training, our methodology also surpasses all previous methods reaching 89.36% top-10 accuracy. These methods represent a 4 – 5 absolute percentage improvement in performance over the current state-of-the-art. Whilst our metric learning approach improves over a simple softmax-based method using similar CNN architecture by 1 percentage point.

2.3 Using triplet networks to embed the DNA sequences

We train Triplet Networks to learn embeddings (vector representations with preserved distance) for both labs and DNA sequences. These vectors live in the same vector space, which allows us to compare them in a variety of ways. For example, we can compare: the distance between two labs; between a lab and a sequence; between a sequence and another sequence. This allows us to perform tasks other than ranking the possible origin lab of a given DNA sequence.

Clustering is one common application that can provide insights [19, 32, 46]. Many labs will share information about design techniques, will have been mentored or trained in another lab or will be directly collaborating. However, those relationships and similarities are not always known to us. Even though these similarities are not directly apparent through embeddings themselves, we can also examine which DNA sequences of a particular lab are more similar to those of other labs. Figure 3 showcases how labs can be clustered with our model and Figure 5 shows labs and their

designed DNA sequences. We observe heterogeneous sizes of clusters with some labs generating only similar DNA sequences (evidenced by compact homogenous clusters), whilst some are highly dispersed.

Fig. 3: **a** Shows the application of the Elbow Method, a commonly used method to find the best number of clusters. As we can see, we could group the labs with various numbers of clusters, but 17 seems to be the optimal number (with a slight margin). **b** Displays a Hierarchical Clustering Dendrogram, making it possible to see the different points to split clusters further. We can notice that some labs can be easily differentiated, while some are very similar according to euclidian distance. **c** Using the optimal number of 17 clusters, we have the number of labs per cluster. In general, the clusters are very similar in size, but some of them have a significant difference from the average. **d** Shows the labs in a 2D space (after compressing information using t-SNE [14]). The colors represent their clusters.

2.4 Few shot learning

Machine learning algorithms typically require large quantities of training data, and in many applications, this makes it challenging when new classes are added. Sparse training data for these classes can result in poor quality predictions. In the case of GEA, some new labs may only have a single or few appropriate training instances. Training machine learning algorithms to perform well in this task is known as Few-Shot Learning (FSL) [9, 10, 45]. Based on knowledge already acquired by a model trained in a similar task, the few-shot learning method can generalize to a new task using simply a few samples. Our proposed method can straightforwardly be adapted to the FSL situation. Since embeddings

are the representation of features, it is possible to use them as the input for other machine learning algorithms or to calculate approximations between the embedding of a new sequence and all previously observed sequences. Thus, in a scenario of new laboratories with few sequences, we can store these few samples and perform the few-shot learning process by calculating the distance between an unknown sequence and the embeddings of these laboratories. It is sufficient that only one of the stored samples has characteristics similar to the unknown sequence for prediction to be made.

To test the ability of our approach to perform few-shot learning, we undertook the following experiment. We trained our method 100 times, each time we removed all the plasmids from 50 different labs in the training set. For each lab that was left out, we picked a random sample of plasmids to generate the embeddings that represent that lab. All remaining plasmids are used to evaluate our model. In the extreme case (also known as One-Shot Learning), we used a single plasmid to represent each lab and test with all the others. Figure 4a shows the metrics’ mean and standard deviation of our model’s capability with different sample sizes used to represent each lab. As we can see, the larger the sample we have, the greater the top-10 accuracy. However, there are diminishing returns as we increase the sample size. It shows that, in general, only a few representative examples are needed for high accuracy predictions. We see that our approach obtains better top-10 accuracy than the previously published CNN approach of Nielsen and Voigt [31] whilst only using 10% of the training data. Figure 4b and Table 1 show the rank of the lab-of-origin when taking a single plasmid to represent each lab. We observe that if we want to be sure the true lab is within our selected sample with probability 0.9 then we need to include around 685 labs. In other words, we can rule out around 50% of the labs as the origin with confidence of 90% when there was only a single plasmid available for that lab. Furthermore, it demonstrates that if an analyst can pick a single representative plasmid of a lab, using domain knowledge, then our approach has a good chance to attribute an unknown plasmid to that lab without any need to retrain the model.

Quantile	Position
50%	7 (top 0.4%)
60%	17 (top 1.2%)
70%	37 (top 2.7%)
80%	180 (top 13.6%)
90%	685 (top 52%)

Table 1: Given all the experiments, in which we pick a single plasmid to represent each lab, this Table shows the ranked position of the lab for each percentage of the cases. For example, 50% of the time, the lab-of-origin was ranked 7th or lower. It is worth noting that the single plasmid for each lab was picked randomly and we do the same experiment multiple times. If an analyst could select a plasmid that they believe to be representative of the lab, we could expect an even better performance.

2.5 Model interpretability and robustness

Interpreting deep learning models gives us valuable information, such as understanding how the model works and the relative importance of features within the data. It can also reveal why some approaches work better than others, and this can be used to further improve the model. However, interpretation techniques for deep learning are still naive and is an area of active research [6, 8].

In this work, we focus on understanding the differences between a triplet network and a conventional classification model, how robust our model is when performing point mutation, and most importantly, which tokens (parts of the sequence) are critical for identifying a lab.

Fig. 4: **a** Each bar shows the mean and standard deviation of Top 10 Accuracy across all the experiments. It is worth noting that all the labs of the dataset had the chance to be left out of the training more than once with those experiments. The first bar refers to the extreme case where we pick a single plasmid to represent the lab. The other bars refer to picking a percentage of the plasmids to represent the lab and using the rest to evaluate. It is worth noting that the number of plasmids per lab varies a lot. There are labs in which 10% will be one or two, while others that it will be hundreds of examples. For reference, the mean and standard deviation of each percentage are as follow: 10% (5 +- 18), 20% (9 +- 35), 30% (13 +- 53), 40% (17 +- 71), 50% (22 +- 88), 60% (26 +- 106), 70% (31 +- 124), 80% (35 +- 142), 90% (39 +- 159). For last, the dashed line refers to our Top 10 Accuracy when retraining the model. **b** The histogram presents the ranked position of the lab-of-origin when using a single plasmid to represent it. As we can see, in most cases (79%) a single plasmid is enough to rank it at least in the top 100 (given the labs in the dataset). It is worth noting that the median (50% of the cases) ranked position was fifth. The variance can be high because there chosen plasmid may not be representative of the lab

We start by visualizing the differences between the space of features mapped between the two models. For a triplet network, this space is the sequence embeddings. Whilst, for the softmax model, we take the output of the 3072-dimensional last hidden layer. This layer is the concatenation between all convolutional layers and contains all the features used by the model. The two multi-dimensional vectors are reduced to 2D space using t-distributed stochastic neighbor embedding (tSNE). As mentioned in Alley et al. [2], the model is more accurate when the plasmid features are more separable in the latent (unobserved) space. We observe in figure 5 that the triplet network model has better-defined clusters.

It is also important to understand the robustness of the model, as small changes in plasmids can frequently occur. We perform random point mutations in the sequence from a lab and report the ranking of the correct lab generated by the model. Figure 5c shows the mean and median of correct positions after 1000 runs for each perturbation. We found that virtually all runs with up to 100 mutations predicted the correct lab within the top 10 guesses. By increasing the number of mutations, the average of the runs becomes unstable, with the average of the positions being higher for all cases above 400 mutations. However, if we examine the median of the predicted positions, even with 1000 mutations, the median rank for the correct lab remains within the top-10. This indicates that our model is robust to most sequence perturbations, excluding cases where these mutations affect essential features for the model’s prediction.

We proceed to analyze all sequences to discover the importance of each plasmid feature to the model output. Unlike perturbation analysis in each sequence, here we use more recent methods that generate better insights in the interpretation of the model. Our method is based on integrated gradients [39]. The idea is to compute the gradient

Fig. 5: **a** Each circle represents a DNA sequence, with its color highlighting its lab-of-origin. Each plus sign represents a lab. We project all of them from 200D to 2D using t-SNE for presentation purposes. We can see that the DNA sequences group together very well with their lab-of-origin. Some labs display very similar DNA sequences, while others are a bit dispersed. It showcases the differences between large and small labs. **b** tSNE visualization of the 3072D last hidden layer of the softmax model. Although we can see clusters in this feature map, they are not as well defined as in the triplet model. **c** Effect of point mutations on the Triplet model. The mean and median of 100 runs of the position of the correct lab in the model prediction ranking is shown, with the number of mutations ranging from 1 to 1000.

of the model's output relative to the embeddings token layer. This makes it possible to visualize the importance of

each token for the model’s prediction. After calculating the integrated gradient for all sequences, we obtain the token importance of each lab by averaging all sequences in that lab. The same process can be performed with all sequences to get the most seen tokens in the dataset.

As we can see from figure 6 there are some tokens that appear to be shared by all labs. When generating the token importance of a lab, we can subtract it from the most seen tokens in the dataset to obtain a relative importance. This allows us to examine those tokens and those not to be expected for a particular lab. Furthermore, we can compare the token importance of one lab with the most distant lab from it in the embedded space. The graphs of token importance from the two labs are essentially mirrored, indicating that tokens are quite different in each case. Figure 6 show these analyzes, where in the left-hand column, we plot Normalized Token Importance (NTI) as a function of the token. In the right hand column we highlight the sequences with the largest token importance.

We further explored this model by looking at the token importance for David Root’s lab (Figure 6B). From this analysis, we can see a cluster of sequences that are typical from this lab, allowing us to identify the potential design signatures. Furthermore, this lab is the furthest lab from the "unknown engineered" category. This class is a mixture of possible labs and so has poorly defined features. The fact that David Root’s lab is the most different from this class suggests it has well-defined and perhaps has highly unique design choices. We note that for the "unknown engineered" class that the scale for the normalised token importance is shallow, and the colour gradient mostly red (see Figure 6D). This demonstrates that there is not a clear design choice or discriminating feature of this category, which is to be expected as it is a mixture of many possible labs (see methods). This analysis could be repeated for any of the labs in the dataset to identify key signatures or potential collaborations based on token proximity.

We next examined the use of integrated gradients for a single sequence. One of the major goals of GEA approaches, is to examine plasmids with unknown origin and be able to extract valuable sequence information, leading to correct assignment or further importation avenues to explore. Figure 7 shows that with our approach we can obtain the importance of each token within an unknown sequence. When comparing the sequence token importance with the lab predicted by the model, we see a concordant behaviour in the plots demonstrating similarity in the highlighted features. This allows us to carefully examine the sequence features the model is using for prediction and hence allows secondary expert evaluation on the veracity of the prediction.

3 DISCUSSION

This manuscript presents a new state-of-the-art for genetic engineering attribution using convolutional neural networks and metric learning. We achieve 89.4% top-10 lab prediction with a conventional classification model, by simply improving training details for better convergence and making use of ensemble learning by using parallel convolutions with different kernel sizes. Furthermore, we show that it is possible to treat the genetic engineering attribution problem as a metric learning problem, creating a vector space where genetic sequences with similar characteristics lie next to one another. Metric Learning is quite common in other areas such as recommendation systems [16, 47]. However, this is the first application to genetic engineering attribution. This methodology further improves the accuracy of our model, reaching 90.4% top-10 accuracy, a 5.4% improvement compared to the former state of the art and has several new advantages such as creating vector representation of labs, comparison and clustering of DNA sequences and labs-of-origin and the ability to examine design style and robustness to unseen labs. For example, a plasmid sequence might be too distant from known labs, resulting in low similarity values. Furthermore, we also have a particular embedding for unknown labs. If a new plasmid sequence is nearer to this embedding than to the known labs, it is possible that this sequence is from a currently unobserved lab. Meanwhile, a classifier model does not usually know

Fig. 6: **a** Normalized Token Importance (NTI) for all labs in the dataset obtained by averaging the token importance of all sequences and normalizing them between 0 and 1. On the right, the top 30 tokens for all data. **b** The normalized token importance for David Root's lab shows that it has similar tokens to all labs in the TOKEN ID range between 600 and 700, but in some other regions it differs a lot. **c** Difference between the token importance of David Root's lab and all labs. This graph highlights which tokens make a difference for this particular lab, whether considering the presence or absence of a token when compared to other labs. On the right, the tokens that should be observed when analyzing this lab. **d** The token importance of the furthest lab to David Root's lab in the embedding space. The NTI values of the two laboratories are practically mirrored, indicating that these two laboratories have opposite characteristics.

Fig. 7: Plasmid features importance of unknown sequences. **a** Normalized Token Importance for an unknown sequence. This may help to investigate specific patterns. **b** The NTI for Bernard Moss's lab which was assigned the sequence author by our model. **c** Comparison between tokens highlighted for the sequence and important tokens from the predicted lab. The red line represents the sequence. **d** Plotting the difference between the NTI of the sequence and the NTI of the predicted laboratory, we can see that few tokens stand out in this sequence beyond the usual presented by the laboratory.

how to handle such uncertainty. Typically, it spreads probabilities for each lab, summing up to 1.0. Hence, any plasmid sequence is assigned to known labs, even if it is from a completely unknown lab.

Additionally, we demonstrated that following our methodology is also possible to perform few-shot learning. We achieve 58.1% top-10 accuracy using only one sample, and with only 10% of the sequences, we outperform the previous CNN approach without training a new model, simply by comparing embeddings vectors. This training methodology allows the possibility of identifying new laboratories with few samples and even just a single genetic sequence. Clearly, there is a tradeoff between sample quantity and model accuracy, but we believe that such methodology could be useful in extreme cases. Finally, these embeddings are also feature-rich, which means we can use them as input for other machine learning models, tackling other problems. For example, we are able to extract the defining signatures for labs and compare them to others using our approach.

Although we present a new state-of-the-art classifier and a new training methodology with interpretability, we believe that there are possible improvements to our model. This could include the use of more advanced machine learning architectures, other pre-processing methods and new data augmentation techniques. This would lead to better convergence and algorithm training. We also believe that new techniques like Transformers [41] and Graph Convolutional Networks [35] would be good candidates for this task, since patterns inside the sequence can be considered contextualized, for which Transformers generally show good performance [41]. We hope our methodology and results encourage new architectures to tackle this problem.

4 METHOD

4.1 Addgene dataset description and data splitting

The addgene data was the same used by [2], and comprised all plasmids deposited in the Addgene repository up to July 27th 2018 – a total of 81,834 entries. For each plasmid, the dataset included a DNA sequence, along with metadata on growth strain, growth temperature, copy number, host species, bacterial resistance markers, and other selectable markers. Each of these categorical metadata fields was re-encoded as a series of one-hot feature groups:

- Growth strain: growth_strain_ccdb_survival, growth_strain_dh10b, growth_strain_dh5alpha, growth_strain_neb_stable, growth_strain_other, growth_strain_stbl3, growth_strain_top10, growth_strain_xl1_blue
- Growth temperature: growth_temp_30, growth_temp_37, growth_temp_other
- Copy number: copy_number_high_copy, copy_number_low_copy, copy_number_unknown
- Host species: species_budding_yeast, species_fly, species_human, species_mouse, species_mustard_weed, species_nematode, species_other, species_rat, species_synthetic, species_zebrafish
- Bacterial resistance: bacterial_resistance_ampicillin, bacterial_resistance_chloramphenicol, bacterial_resistance_kanamycin, bacterial_resistance_other, bacterial_resistance_spectinomycin
- Other selectable markers: selectable_markers_blasticidin, selectable_markers_his3, selectable_markers_hygromycin, selectable_markers_leu2, selectable_markers_neomycin, selectable_markers_other, selectable_markers_puromycin, selectable_markers_trp1, selectable_markers_ura3, selectable_markers_zeocin

In addition to the sequence and the above metadata fields, the raw dataset also contained unique sequence IDs, as well as separate IDs designating the origin lab. Both sequence and lab IDs were obfuscated through 1:1 replacement with random alphanumeric strings. The number of plasmids deposited in the dataset by each lab was unbalanced, with many labs depositing one or a few sequences. To deal with this problem, Alley et al. [1] grouped labs with fewer than 10 data points into a single auxiliary category labelled “Unknown Engineered”. This reduced the number of categories

from 3751 (the number of labs) to 1314 (1313 unique labs + Unknown Engineered). In addition to issues with small labs, the dataset also contains “lineages” of plasmids. That is, sequences that were derived by modifying other sequences in the dataset. If unmitigated this introduces unintended correlations between the test and validation set. To overcome this, Alley et al. [1] inferred lineage networks among plasmids in the dataset, based on information in the complete Addgene database acknowledging sequence contributions from other entries. Lineages were identified by searching for connected components within the network of entry-to-entry acknowledgements in the Addgene database and we refer to Alley et al. [1] for more details. The data were partitioned into train, validation, and test sets, with the constraints that (i) every category have at least three data points in the test set, and (ii) all plasmids in a given lineage be assigned to a single dataset. Following the split, the training set contained 63,017 entries (77.0%); the validation set contained 7,466 entries (9.1%); and the test set contained 11,351 entries (13.9%).

4.2 Grouping sequences by their Levenshtein distance

Genetic sequences from the same lab display large degrees of similarity. These sequences can make it easier to identify similar sequences in the training and validation sets. However, when training a machine learning algorithm, this may be perceived as data leakage between these sets [20], as the model does not need to learn to extract different features to identify such sequences. To ameliorate this issue, we developed a more robust model by group sequence from each lab based on their Levenshtein distance [5]. The Levenshtein formula used can be seen in equation 1.

$$\text{lev}_{a,b}(i, j) = \begin{cases} \max(i, j) & \text{if } \min(i, j) = 0, \\ \min \begin{cases} \text{lev}_{a,b}(i-1, j) + 1 \\ \text{lev}_{a,b}(i, j-1) + 1 \\ \text{lev}_{a,b}(i-1, j-1) + 1_{(a_i \neq b_j)} \end{cases} & \text{otherwise.} \end{cases} \quad (1)$$

After grouping, each laboratory has N groups of sequences. We then split the dataset, ensuring that there will be sequences from the same group only in training or validation, which means they will never be present in both sets at the same time. This entire process was performed using python and the python-Levenshtein library (<https://github.com/ztane/python-Levenshtein>). As this is a costly algorithm and there are thousands of sequences to be grouped together, the entire process was performed on a machine with 128GB RAM and AMD EPYC 7401P 24-Core processor. This approach complements the lineage-based strategy, which also avoids data leakage.

4.3 Training cross-validation

We performed a 5-fold cross-validation strategy [12] for each experiment within the training set. To be precise, for each hyperparameter setting, we split the data into k parts (we used $k = 5$), using one of them to validate and the remaining to train, repeating this process k times. After that, we evaluate each model, by taking the mean of the metrics for that experiment. This approach helps to avoid overfitting and improve generalisation. It also enables us to ensemble the k models, to further improve generalisation.

4.4 Byte pair encoding

Similar to Alley et al., we use the Byte Pair Encoding algorithm [11] to process all sequences in the dataset, grouping subsequences into new tokens, increasing the vocabulary while reducing the sequences’ size. The BPE algorithm first examines all the sequences in the dataset to learn how to perform the grouping. The trained algorithm is saved and used

to transform the sequences during the convolutional neural network training. This last process is performed "online". This means that, while loading our batch of samples, we transform the batch into new small sequences. We converted our vocabulary size from 4 DNA bases into 1001 different tokens: 1000 tokens from the new vocabulary generated by the BPE algorithm plus 1 for an unknown token. The training and inference of the BPE algorithm were performed using the sentencepiece package (<https://github.com/google/sentencepiece>), and, as in the section 4.2, we use python and the same machine to perform this operation.

4.5 Circular data augmentation

Machine learning models and especially deep learning models are highly dependent on large amounts of data. One of the fundamental methods for adding variance to those models, increasing generalizability, and reducing overfitting [7] is data augmentation [27]. Generally, data augmentation performs transformations on the sample, considerably changing some characteristics. In this work, performing such transformations can be dangerous as it may end up modifying some essential parts to assign the sequence to a lab-of-origin. However, it is possible to take advantage of the fact that plasmids are circular and create a circular shift data augmentation process. This contrasts with a reverse complement augmentation one might use. During training, we show different versions of the same DNA sequence by shifting it circularly as shown in 1. This approach helps the model understand that the same pattern can happen at different positions within the sequence, increasing the generalizability of the training. Further, to decrease the model's complexity, we also limited the sequence to 1000 tokens.

We also perform Test Time Augmentation [29] which helps to improve the model's prediction capability. To perform this analysis, during inference we run the model multiple times. Each time the model sees a shifted version of the sequence and makes a prediction of the same sample seen from different angles. We then take the average of the outputs (class probabilities for the classifier and embeddings for our proposed method).

4.6 CNN base architecture and training details

Both types of models are composed of a Convolutional Neural Network with multiple kernels of different sizes, as proposed by Kim. We use it to extract features from the sequence and concatenate them with the binary features provided in the dataset. The difference between the classification and triplet network models is on the final layers. The final base structure is composed of an embedding layer, several convolutional layers in parallel with different kernel sizes and custom dropout layer for regularization. The embedding layer has the shape of 1001x200, where 1001 is our vocabulary size and 200 the vector embedding dimension found empirically. Its purpose is to map each token into a 200-dimensional vector containing the features representation of that token [48]. For the convolutional layers, we have a total of 12 layers in parallel, where the first layer has kernel size 1, the second has kernel size 2 and so on, until the last layer has kernel size 12. All convolutional layers are followed by a SeLU activation function [23] and a max pooling operation. We concatenate the features extracted by each of them, obtaining the final representation with different windowings of the sequence. We also implemented a custom Dropout Layer. A standard Dropout Layer [38] randomly masks out parts of a tensor to regularize the neural network. But if we did that on the embeddings before applying a similarity function, the output would be too unstable. So, we created a layer that randomly masks out the same parts of all the embeddings involved before applying the similarity function. We found this approach instrumental in regularizing our model.

The entire architecture was developed and trained using Python and Pytorch [33]. Although the training methodology is different between the two approaches, all the training details, such as optimizer, learning rate scheduler and

regularization techniques remain the same. We use the Adam [22] optimizer together with the One Cycle learning rate scheduler [37]. This scheduler was essential to achieve a better convergence in training and its settings were maximum learning rate of $1e^{-3}$ and cycle execution in 200 epochs. To regularize our model and prevent overfitting, we used a weight decay of $1e^{-5}$ during training, and dropped 20% of the embedded sequences using our custom dropout layer.

4.7 Triplet network learning

To generate the triplets, we use the labeled dataset to provide us with the anchor and positive. We then use a technique known as Hard Negative Mining [13] to select the negative (an incorrect lab). This means that rather than choosing a random lab as a negative example, we choose the most challenging one given the current state of the embeddings. Thus, in our case it would be the nearest lab to our sequence in the latent space.

One of the most challenging parts of this work was the implementation of the algorithm to mine the negative examples during training efficiently. We could have used a library called PyTorch Metric Learning [30]. They have Hard Negative Mining implemented per batch (it does not take the whole dataset into account while finding the negative) and implemented Cross-Batch Memory for Embedding Learning [44]. However, this library only supports a single entity type. Furthermore, we also have easy access to the whole lab embeddings, since we use an embedding layer. So, the approach we re-implemented for our specific needs as described in the Algorithm 1. It is worth noting that we implemented it using tensors to make it as efficient as possible. Our source code provides a PyTorch implementation, and it should be straightforward to implement it in Tensorflow and other frameworks.

<p>Input: <i>lab_indices</i> (B,), <i>anchor_embeddings</i> (B, L), <i>positive_embeddings</i> (B, E)</p> <p>Output: Batch of <i>negatives</i> of shape (B, E)</p> <pre> 1 <i>all_lab_indices</i> $\leftarrow$ indices of all labs repeated by B // (B, L) 2 <i>negative_labs_mask</i> $\leftarrow$ matrix of boolean values with <i>True</i> for every lab // (B, L) 3 <i>negative_labs_mask</i>[:, <i>lab_indices</i>] $\leftarrow$ <i>False</i> // (B, L) 4 // At this point, we have a mask with True for the negative labs and False otherwise 5 <i>all_negative_lab_indices</i> $\leftarrow$ <i>all_labs</i>[<i>negative_labs_mask</i>] reshaped // (B, L-1) 6 <i>all_negative_labs_embeddings</i> $\leftarrow$ <i>lab_embeddings</i>[<i>all_negative_labs</i>] L2 normalized // (B, L-1, E) 7 <i>anchor_similarities</i> $\leftarrow$ dot product between <i>anchor_embeddings</i> and <i>all_negative_labs_embeddings</i> // (B, L-1) 8 <i>hardest_negative_lab_indices</i> $\leftarrow$ argmax of <i>anchor_similarities</i> // (B,) 9 <i>negative_embeddings</i> $\leftarrow$ <i>all_negative_labs_embeddings</i>[:, <i>hardest_negative_lab_indices</i>] // (B, E) 10 return <i>negative_embeddings</i> </pre>
--

Algorithm 1: Algorithm for Hard Negative Mining using tensors. The shape of each tensor is at the end of each line as a comment, being B batch size, E the embedding dimension, and L is the number of labs. It is worth noting that we L2 normalize all the embeddings.

4.8 Cosine similarity

We used the cosine similarity as the metric to measure how similar the vectors were in the embeddings. Mathematically, it measures the cosine of the angle between two vectors projected in a multi-dimensional space, resulting in a value inside the range of -1 and 1, where -1 indicates opposite vectors and 1 indicates equal vectors. Given two non-zeros vectors of embeddings, A and B, the cosine similarity is:

$$\text{similarity} = \cos(\theta) = \frac{A \cdot B}{\|A\| \|B\|} = \frac{\sum_{i=1}^n A_i B_i}{\sqrt{\sum_{i=1}^n A_i^2} \sqrt{\sum_{i=1}^n B_i^2}} \quad (2)$$

4.9 t-SNE and K-means

Throughout this work, we have used 2 techniques to analyze and visualize the results, the *t*-SNE algorithm, and the K-means cluster algorithm. T-Distributed Stochastic Neighbor Embedding [40] is a dimensionality reduction technique that can reduce dimensions with non-linear relationships. It is particularly well suited for the visualization of high-dimensional complex real-world datasets. Making use of this, we are able to reduce the embeddings generated by the triplet network and visualize them in 2D space. K-means [4] is a clustering algorithm that attempts to organize the data into K clusters. The objective of the algorithm is to group similar data together by their Euclidean distance to K centroids, where K is a value chosen by the user, while keeping all centroids distant from each other. Each sample will be linked (or allocated/assigned) to the cluster with the nearest centroid.

4.10 Interpreting the model

To visualize the mapped features of both models, we took different approaches since they are different models. For the triplet network model, we first inferred all sequences from the validation set and got their embeddings. These embeddings were 200-dimensional and they were reduced to 2 dimensions using tSNE in scikit-learn project [34] with default parameters. For the classification model, we extracted the activations of the last hidden layer that maps features from all convolutional layers, before concatenating those features with the extra inputs (sequence metadata) and passing through the last layer which outputs logits. These hidden features were 3072-dimensional, reduced to 2 dimensions in the same way as embeddings. The visualization was done using matplotlib [17] and we colored each point by the corresponding lab.

To analyze the influence of perturbations on model prediction, we took a specific plasmid from the validation set and randomly generated perturbation in its sequence. As the process is random, we perform 100 experiments for each number of mutations, ranging from 1 to 1000. To make these mutations, we use a random integer function to select the specific position to be changed and a random choice operation to choose one of the 5 possible bases in the sequence (N, G, C, T, A). For each of the 100,000 runs, we took the position of the correct laboratory in the model's prediction ranking. The figure 5c presents the mean and median of those positions.

To find out which tokens are most important within a sequence, we decided to use a similar methodology to integrated gradients [39]. Integrated gradient is an interpretability technique for deep neural networks which finds the input features that contribute the most to the model prediction. We started by computing gradients between model predictions with respect to the sequence embedding layer, getting a matrix of gradients in the shape of 1001x200 (number of tokens x embedding dimension). Each gradient measures the relationship between the embedding weight and the output. We then calculate the absolute value element-wise and sum them up in the second axis, generating a final vector of 1001 positions containing the summed importance of each token.

$$\text{TokenImportance}_i(x) = \sum_{k=1}^{200} \left\| \frac{\partial F(x)}{\partial x_{ik}} \right\| \times \frac{1}{200} \quad (3)$$

where:

$F()$ = model's prediction function

$\frac{\partial F}{\partial x_{ik}}$ = gradient of model F 's prediction function relative to each embedding feature x_{ik}

i = number of token

k = embedding dimension position

After generating the token importance of each sequence in the validation set, we took the token importance of each laboratory by averaging the token importance of all sequences in that laboratory. The visualization was done using matplotlib, and to better present the figure, the token importance values were normalized between 0 and 1 (NTI). To compare the NTI of a specific lab with the further lab from it, we compute the cosine similarity between the analyzed lab embedding and all other lab embeddings by performing a dot product. The lowest value indicates the least similar laboratory.

ACKNOWLEDGMENTS

To Amalgam and XNV for providing us the necessary infrastructure and financial support. OMC acknowledges funding from a Todd-Bird Junior Research Fellowship from New College, Oxford, as well as Open Philanthropy.

REFERENCES

- [1] Ethan C. Alley, Miles Turpin, Andrew Bo Liu, Taylor Kulp-McDowall, Jacob Swett, Rey Edison, Stephen E. Von Stetina, George M. Church, and Kevin M. Esvelt. 2020. Attribution of genetic engineering: A practical and accurate machine-learning toolkit for biosecurity. *bioRxiv* (2020). <https://doi.org/10.1101/2020.08.22.262576> arXiv:<https://www.biorxiv.org/content/early/2020/08/22/2020.08.22.262576.full.pdf>
- [2] Ethan C. Alley, Miles Turpin, Andrew Bo Liu, Taylor Kulp-McDowall, Jacob Swett, Rey Edison, Stephen E. Von Stetina, George M. Church, and Kevin M. Esvelt. 2020. A machine learning toolkit for genetic engineering attribution to facilitate biosecurity. *Nature Communications* 11, 1 (08 Dec 2020), 6293. <https://doi.org/10.1038/s41467-020-19612-0>
- [3] Stephen F. Altschul, Warren Gish, Webb Miller, Eugene W. Myers, and David J. Lipman. 1990. Basic local alignment search tool. *Journal of Molecular Biology* 215, 3 (Oct. 1990), 403–410. [https://doi.org/10.1016/s0022-2836\(05\)80360-2](https://doi.org/10.1016/s0022-2836(05)80360-2)
- [4] David Arthur and Sergei Vassilvitskii. 2007. K-Means++: The Advantages of Careful Seeding. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms* (New Orleans, Louisiana) (SODA '07). Society for Industrial and Applied Mathematics, USA, 1027–1035.
- [5] Bonnie Berger, Michael Waterman, and Yun Yu. 2020. Levenshtein Distance, Sequence Comparison and Biological Database Search. *IEEE Transactions on Information Theory* PP (05 2020), 1–1. <https://doi.org/10.1109/TIT.2020.2996543>
- [6] Supriyo Chakraborty, Richard J. Tomsett, Ramya Raghavendra, Daniel Harborne, Moustafa Alzantot, F. Cerutti, Mani B. Srivastava, Alun David Preece, Simon J. Julier, Raghuvir M. Rao, Troy D. Kelley, Dave Braines, M. Sensoy, Chris J. Willis, and Prudhvi K. Gurram. 2017. Interpretability of deep learning models: A survey of results. *2017 IEEE SmartWorld, Ubiquitous Intelligence & Computing, Advanced & Trusted Computed, Scalable Computing & Communications, Cloud & Big Data Computing, Internet of People and Smart City Innovation (SmartWorld/SCALCOM/UIC/ATC/CBDCom/IOP/SCI)* (2017), 1–6.
- [7] Davide Chicco. 2017. Ten quick tips for machine learning in computational biology. *BioData Mining* 10, 1 (Dec. 2017). <https://doi.org/10.1186/s13040-017-0155-3>
- [8] Finale Doshi-Velez and Been Kim. 2017. Towards A Rigorous Science of Interpretable Machine Learning. arXiv:1702.08608 [stat.ML]
- [9] Li Fei-Fei, R. Fergus, and P. Perona. 2006. One-shot learning of object categories. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 28, 4 (April 2006), 594–611. <https://doi.org/10.1109/tpami.2006.79>
- [10] Michael Fink. 2005. Object Classification from a Single Example Utilizing Class Relevance Metrics. In *Advances in Neural Information Processing Systems*, L. Saul, Y. Weiss, and L. Bottou (Eds.), Vol. 17. MIT Press. <https://proceedings.neurips.cc/paper/2004/file/ef1e491a766ce3127556063d49bc2f98-Paper.pdf>
- [11] Philip Gage. 1994. A New Algorithm for Data Compression. *C Users J.* 12, 2 (Feb. 1994), 23–38.
- [12] Trevor Hastie, Robert Tibshirani, and Jerome Friedman. 2009. *The elements of statistical learning: data mining, inference and prediction* (2 ed.). Springer. <http://www-stat.stanford.edu/~tibs/ElemStatLearn/>
- [13] Alexander Hermans, Lucas Beyer, and Bastian Leibe. 2017. In Defense of the Triplet Loss for Person Re-Identification. arXiv:1703.07737 [cs.CV]
- [14] Geoffrey Hinton and Sam Roweis. 2002. Stochastic Neighbor Embedding. In *Proceedings of the 15th International Conference on Neural Information Processing Systems (NIPS'02)*. MIT Press, Cambridge, MA, USA, 857–864.
- [15] Elad Hoffer and Nir Ailon. 2018. Deep metric learning using Triplet network. arXiv:1412.6622 [cs.LG]

- [16] Cheng-Kang Hsieh, Longqi Yang, Yin Cui, Tsung-Yi Lin, Serge Belongie, and Deborah Estrin. 2017. Collaborative Metric Learning. International World Wide Web Conferences Steering Committee. <https://doi.org/10.1145/3038912.3052639>
- [17] J. D. Hunter. 2007. Matplotlib: A 2D graphics environment. *Computing in Science & Engineering* 9, 3 (2007), 90–95. <https://doi.org/10.1109/MCSE.2007.55>
- [18] Joanne Kamens. 2014. The Addgene repository: an international nonprofit plasmid and data resource. 43, D1 (Nov. 2014), D1152–D1157. <https://doi.org/10.1093/nar/gku893>
- [19] Md Rezaul Karim, Oya Beyan, Achille Zappa, Ivan G Costa, Dietrich Rebholz-Schuhmann, Michael Cochez, and Stefan Decker. 2020. Deep learning-based clustering approaches for bioinformatics. 22, 1 (Feb. 2020), 393–415. <https://doi.org/10.1093/bib/bbz170>
- [20] Shachar Kaufman, Saharon Rosset, and Claudia Perlich. 2011. Leakage in Data Mining: Formulation, Detection, and Avoidance. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (San Diego, California, USA) (KDD '11). Association for Computing Machinery, New York, NY, USA, 556–563. <https://doi.org/10.1145/2020408.2020496>
- [21] Yoon Kim. 2014. Convolutional Neural Networks for Sentence Classification. arXiv:1408.5882 [cs.CL]
- [22] Diederik P. Kingma and Jimmy Ba. 2014. Adam: A Method for Stochastic Optimization. <http://arxiv.org/abs/1412.6980> cite arxiv:1412.6980Comment: Published as a conference paper at the 3rd International Conference for Learning Representations, San Diego, 2015.
- [23] Günter Klambauer, Thomas Unterthiner, Andreas Mayr, and Sepp Hochreiter. 2017. Self-Normalizing Neural Networks. *CoRR* abs/1706.02515 (2017). arXiv:1706.02515 <http://arxiv.org/abs/1706.02515>
- [24] Gregory Koch, Richard Zemel, and Ruslan Salakhutdinov. 2015. Siamese Neural Networks for One-shot Image Recognition.
- [25] Brian Kulis. 2013. Metric Learning: A Survey. *Foundations and Trends® in Machine Learning* 5, 4 (2013), 287–364. <https://doi.org/10.1561/22000000019>
- [26] Zachary Chase Lipton. 2015. A Critical Review of Recurrent Neural Networks for Sequence Learning. *CoRR* abs/1506.00019 (2015). arXiv:1506.00019 <http://arxiv.org/abs/1506.00019>
- [27] Agnieszka Mikolajczyk and Michal Grochowski. 2018. Data augmentation for improving deep learning in image classification problem. In *2018 International Interdisciplinary PhD Workshop (IIPHDW)*. IEEE. <https://doi.org/10.1109/iiphdw.2018.8388338>
- [28] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed Representations of Words and Phrases and their Compositionality. In *Advances in Neural Information Processing Systems*, C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger (Eds.), Vol. 26. Curran Associates, Inc. <https://proceedings.neurips.cc/paper/2013/file/9aa42b31882ec039965f3c4923ce901b-Paper.pdf>
- [29] Nikita Moshkov, Botond Mathe, Attila Kertesz-Farkas, Reka Hollandi, and Peter Horvath. 2020. Test-time augmentation for deep learning-based cell segmentation on microscopy images. *Scientific Reports* 10, 1 (March 2020). <https://doi.org/10.1038/s41598-020-61808-3>
- [30] Kevin Musgrave, Serge Belongie, and Ser-Nam Lim. 2020. PyTorch Metric Learning. arXiv:2008.09164 [cs.CV]
- [31] Alec A. K. Nielsen and Christopher A. Voigt. 2018. Deep learning to predict the lab-of-origin of engineered DNA. *Nature Communications* 9, 1 (07 Aug 2018), 3135. <https://doi.org/10.1038/s41467-018-05378-z>
- [32] Mahamed Omran, Andries Engelbrecht, and Ayed Salman. 2007. An overview of clustering methods. *Intell. Data Anal.* 11 (11 2007), 583–605. <https://doi.org/10.3233/IDA-2007-11602>
- [33] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimeshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems* 32, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Eds.). Curran Associates, Inc., 8024–8035. <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>
- [34] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12 (2011), 2825–2830.
- [35] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. 2008. The graph neural network model. *IEEE transactions on neural networks* 20, 1 (2008), 61–80.
- [36] Alex Sherstinsky. 2018. Fundamentals of Recurrent Neural Network (RNN) and Long Short-Term Memory (LSTM) Network. *CoRR* abs/1808.03314 (2018). arXiv:1808.03314 <http://arxiv.org/abs/1808.03314>
- [37] Leslie N. Smith and Nicholay Topin. 2017. Super-Convergence: Very Fast Training of Residual Networks Using Large Learning Rates. *CoRR* abs/1708.07120 (2017). arXiv:1708.07120 <http://arxiv.org/abs/1708.07120>
- [38] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *J. Mach. Learn. Res.* 15, 1 (Jan. 2014), 1929–1958.
- [39] Mukund Sundararajan, Ankur Taly, and Qiqi Yan. 2017. Axiomatic Attribution for Deep Networks. *CoRR* abs/1703.01365 (2017). arXiv:1703.01365 <http://arxiv.org/abs/1703.01365>
- [40] Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing Data using t-SNE. *Journal of Machine Learning Research* 9 (2008), 2579–2605. <http://www.jmlr.org/papers/v9/vandermaaten08a.html>
- [41] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*. 5998–6008.
- [42] Qi Wang, RA Elworth, Tian Rui Liu, and Todd J Treangen. 2019. Faster pan-genome construction for efficient differentiation of naturally occurring and engineered plasmids with plaster. In *19th International Workshop on Algorithms in Bioinformatics (WABI 2019)*. Schloss Dagstuhl-Leibniz-Zentrum

- fuer Informatik.
- [43] Qi Wang, Bryce Kille, Tian Rui Liu, R. A. Leo Elworth, and Todd J. Treangen. 2021. PlasmidHawk improves lab of origin prediction of engineered plasmids using sequence alignment. *Nature Communications* 12, 1 (26 Feb 2021), 1167. <https://doi.org/10.1038/s41467-021-21180-w>
 - [44] Xun Wang, Haozhi Zhang, Weilin Huang, and Matthew R. Scott. 2019. Cross-Batch Memory for Embedding Learning. *CoRR* abs/1912.06798 (2019). arXiv:1912.06798 <http://arxiv.org/abs/1912.06798>
 - [45] Yaqing Wang, Quanming Yao, James T. Kwok, and Lionel M. Ni. 2020. Generalizing from a Few Examples. *Comput. Surveys* 53, 3 (July 2020), 1–34. <https://doi.org/10.1145/3386252>
 - [46] Dongkuan Xu and Yingjie Tian. 2015. A Comprehensive Survey of Clustering Algorithms. 2, 2 (June 2015), 165–193. <https://doi.org/10.1007/s40745-015-0040-1>
 - [47] Junliang Yu, Min Gao, Wenge Rong, Yuqi Song, and Qingyu Xiong. 2017. A Social Recommender Based on Factorization and Distance Metric Learning. 5 (2017), 21557–21566. <https://doi.org/10.1109/access.2017.2762459>
 - [48] Quan Zou, Pengwei Xing, Leyi Wei, and Bin Liu. 2018. Gene2vec: gene subsequence embedding for prediction of mammalian N6-methyladenosine sites from mRNA. *RNA* 25, 2 (Nov. 2018), 205–218. <https://doi.org/10.1261/rna.069112.118>

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [SoaresEPCflat.pdf](#)