Reconfiguring Okazaki fragment start sites on a genome by using a data-driven approach

Adam Soffer[1-3], Morya Ifrach[1], Stefan Ilic[1], Ariel Afek[1], Dan Vilenchik[2,3], Barak Akabayov[1,2,4]

[4]Corresponding author: akabayov@bgu.ac.il

[1]Department of Chemistry, [2]Data Science Research Center, and [3]School of Computer and Electrical Engineering, Ben-Gurion University of the Negev, Beer-Sheva, Israel.

DNA primase, DNA-protein interactions, primase-DNA recognition sequence, machine learning.

ABSTRACT: DNA primase is an essential enzyme that synthesizes short RNA primers on specific DNA sequences. These RNA primers are elongated by DNA polymerase to form Okazaki fragments on the lagging DNA strand. It is therefore reasonable to assume that the binding of DNA primase on a genome marks the start sites of the Okazaki fragments. It has long been known that the frequency of the occurrence of primase trinucleotide recognition on a genome sequence has no influence on the size of the Okazaki fragments. The unresolved enigma that we address in this study is therefore why some, but not all, primase-DNA recognition sequences (PDRSs) become Okazaki fragment start sites. To this end, we applied machine-learning algorithms to analyze a massive amount of data obtained from protein-DNA binding microarrays (PBM) with the aim of identifying the essential elements on DNA that are needed for the binding of bacteriophage T7 primase. A PBM data learning algorithm enabled the prediction of binding values of T7 primase for any given DNA sequence with unprecedented accuracy and flexibility. On the basis of the principles learned about DNA-primase binding, we generated novel DNA sequences with improved binding of T7 primase and improved RNA primer synthesis, as validated experimentally.

INTRODUCTION

During DNA replication the double-stranded DNA is unwound, to expose the two individual DNA strands, one is copied continuously (leading strand) and the other is copied discontinuously (lagging strand). On the lagging DNA strand, an enzyme called DNA primase, recognizes the DNA sequence used as the template for the synthesis of RNA primers, and DNA polymerase is then responsible for elongating these RNA primers into the DNA segments known as Okazaki fragments. This process of RNA-primed DNA synthesis by a DNA polymerase is triggered exclusively by the recognition of a specific DNA sequence by the primase. This recognition is thus fundamental to the establishment of Okazaki fragments and consequently to the whole process of proper DNA replication. In prokaryotes, RNA primer formation occurs on pre-defined sequences on the genome that are specifically recognized by a DnaG-type primase[2]. For example, the activity of bacteriophage T7 DNA primase, comprising the N-domain [residues 1–271] of gene 4 protein, is initiated by sequence-specific binding of DNA primase to 5'-GTC-3' [3,4], which is then followed by the synthesis of a functional primer[5]. Importantly, it is now known that even though DNA primase recognizes a specific trinucleotide sequence, flexibility in the selection of initiation sites for Okazaki fragments is allowed[6], i.e., not every primase-DNA recognition sequence (PDRS) will become an Okazaki fragment start site. The reason for this flexibility is, however, not understood and is the enigma that we address in the current study: Although extensive research has been carried out on the interactions of DNA primase with DNA, it is still not clear why DNA primase ignores the majority of trinucleotide recognition sites. In *Escherichia coli*, for example, DnaG primase ignores ~97% of the trinucleotide recognition sites and initiates Okazaki fragments only every 1.5–2 kb (and not more frequently)[7]. The literature offers two possible explanations for the effect of selective DNA sequence recognition by a primase: The first, well-explored possibility, is that other DNA replication proteins, such as DnaB[8-11], single-stranded DNA-binding protein (SSB)[12], or clamp loader[13], affect the binding of the primase (DnaG) to DNA or even change the preferences for PDRSs on the genome. The second possibility is that a sequence larger than a trinucleotide is required for the specific binding of a primase. On the basis of the mismatch between the frequency of GTC sequences on the bacteriophage T7 genome and the actual size of the Okazaki fragments, it is reasonable to assume that only a sequence larger than a trinucleotide will lead to "effective binding" of a DNA primase, i.e., binding yielding an RNA primer that marks the start site of an Okazaki fragment.

In our first steps to investigate the above-described possibility of a larger DNA binding determinant, we applied high-throughput primase profiling[14], in which binding scores for tens of thousands of DNA primase–DNA binding events on a protein-DNA binding microarray (PBM) were combined with biochemical analysis[14]. This technology facilitated the

analysis of the composition of sequences flanking the specific recognition site and their corresponding binding scores and confirmed that a GTC sequence is not sufficient for the recognition of DNA templates by T7 primase[14,15]. Specifically, we showed that T7 primase has a high affinity for PDRSs containing T/G, C/T, or G/C/T, leading to the formation of longer RNA primers[14]. Although the development of high-throughput primase profiling facilitated the acquisition of massive amounts of data on DNA-primase-binding events, the means to systematically analyze that data in a way that will facilitate a comprehensive understanding of the recognition process are yet to be put in place. Specifically, analysis of this data will throw light on the principles for selection of PDRSs on a genome during DNA replication by enabling us to answer the following questions: 1) Is there information stored in the DNA sequence that is important for T7 primase binding? 2) If we understand the principles of specific primase DNA recognition, can we predict binding scores of T7 primase for a given DNA sequence? And 3) Can we generate new DNA sequences with desired binding scores based on the sequence features embedded in the DNA? Answering the third question will enable us to ascertain which features embedded in specific DNA binding sequences govern the binding of DNA primase.

In concert with cutting-edge developments in biochemical technologies, current progress in computational science provides us with the opportunity to construct knowledge-based models that will help us to answer the above questions. Here, we describe an intelligent learning workflow that provides a comprehensive view of the principles that govern the design and activity of PDRSs with unprecedented flexibility and accuracy. We applied this workflow to elucidate the link between the larger context, i.e., the flanking nucleotides, of the primase recognition sequence and the synthesis of RNA primers that initiate Okazaki fragments.

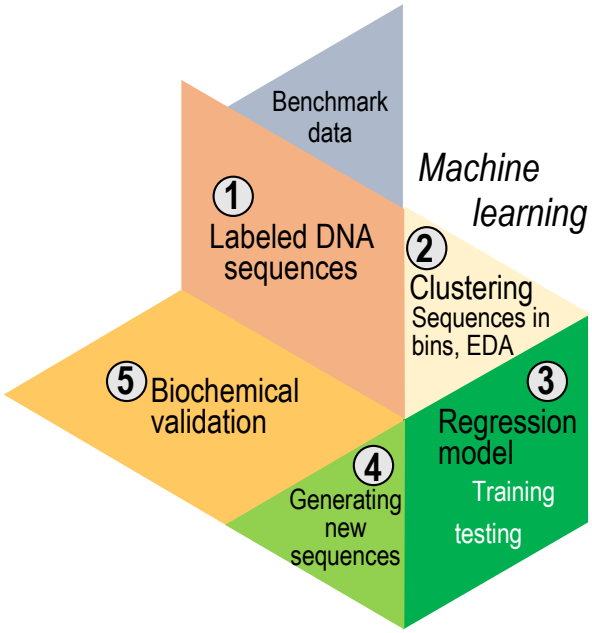

Scheme 1. Analysis workflow after preprocessing the data from the primase–DNA binding microarray. The benchmark dataset containing DNA sequences for the training set was preprocessed (step 1). The DNA sequences were clustered into five bins using exploratory data analysis (EDA), i.e., unsupervised algorithms (step 2). A different regressor was trained for every cluster. Several regression algorithms were used; linear regression with L1 regularization provided the best results. To predict the binding scores of a new DNA sequence, the sequence was assigned to a specific bin, and its score for primase binding was predicted using that bin's regressor (step 3). Novel DNA sequences (PDRSs) with high binding score for primase were generated (step 4). It was then possible to examine the ability of those PDRSs to bind primase and induce the synthesis of RNA primers (step 5).

RESULTS AND DISCUSSION

The overall structure of our study is comprised of the following stages of analysis of the data obtained using PBMs for quantitative measurements of T7 primase-DNA binding, after preprocessing of the data (Scheme 1): step 1) preparation of a PBM-driven benchmark data; step 2) clustering the PDRSs containing DNA sequences; step 3) training a regression model; and step 4) predicting the score of new DNA sequences, and generating novel DNA sequences with desired binding scores for T7 primase. These steps are elaborated below, as are the data preprocessing and step 5 (Scheme 1), which is biochemical validation.

Data preprocessing and vectorization of DNA sequences. Before the data analysis, considerable attention was paid to data preprocessing, as the success of the subsequent application of machine-learning algorithms depended on the explicit presentation of the data in a way that facilitated the extraction of meaningful features and the removal of distracting outliers. The preprocessing of PBM-derived DNA-primase binding data comprised four steps: data cleansing, data filtration, embedding of the sequences into vectors, and data normalization, as follows. Using the PDRSs as meaningful "words" on the basis of their sequence features, where each sequence was assigned to its PBM-driven binding score, we focused on the sequences that could potentially serve as Okazaki fragment start sites. We started with the preparation of a "lexicon" of DNA "words," each comprising a larger context of GTC-containing sequences that allow effective binding of T7 primase, i.e., the binding of T7 primase that yields RNA primers. The starting point for the preprocessing was that while an average size of ~64 nucleotides (Figure 1a) is the expected distance between two GTC sequences, Okazaki fragment size of 1000-6000 nucleotides is obtained experimentally (Figure 1b, marked in the red range box).

We thus posited that GTC-containing PDRSs must be larger than a trinucleotide sequence to meet the frequency on the genome that would allow the creation of Okazaki fragments of sizes that were observed previously (Figure 1c, red lines: larger GTC-containing sequences that match experimental sizes of Okazaki fragments[16]; black lines: frequency of GTC every 64 bases).

Since DNA sequences constitute a form of categorical data represented by nucleotides, the preprocessing step was required to convert the plain representation of DNA sequences into a meaningful numeric representation. Such a representation of DNA sequences was obtained by using One Hot Encoding (OHE). In this way, a categorical sequence was converted into an array of integers in which each nucleotide was represented by four unit vectors: (A=[1000], C=[0100],

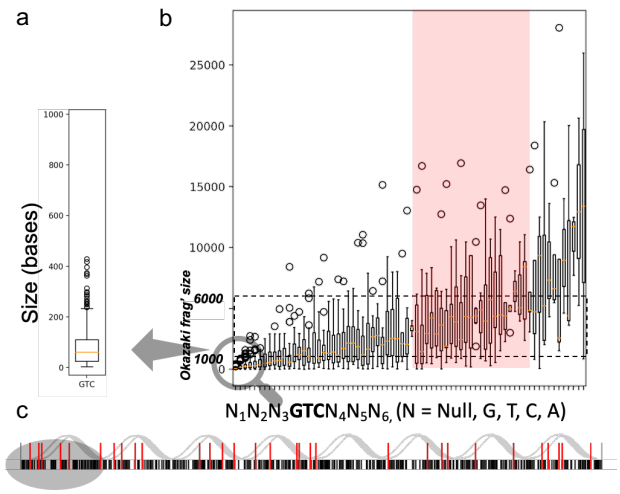

Figure 1. Frequency and distribution of GTC-containing primase-DNA recognition sequences (PDRSs) on the bacteriophage T7 genome. a, Size distribution of GTC-containing DNA sequences on the genome of bacteriophage T7. The frequency of the occurrence of GTC in a random sequence is every $4^3 = 64$ bases (approximately as indicated by the orange line). b, Calculated size distribution of the DNA sequence between GTC-containing PDRSs on the T7 genome that match the actual size of Okazaki fragments. These PDRSs consist of 0-3 nucleotides flanking the GTC sequence and are distributed at an inter-PDRS distance that ranges between 1000-6000 nucleotides, which yield Okazaki fragments of the same size. c, All combinations of possible T7 PDRSs (5′-GTC-3′) on the genome are considered. Black lines represent the frequency of GTCs; red lines represent the frequency of large-context GTC-containing PDRSs that match experimental values for Okazaki fragment sizes.

G=[0010], T=[0001]), e.g., the sequence ACCG was encoded as 1000|01000|01000|0010. Every DNA sequence, represented by a 144-dimensional vector, was fed as an input into both a Kmeans model (using Euclidian distance) and a Ward-method-based[17] hierarchical clustering model.

Defining the mathematical descriptors (features) of the PDRSs. The challenge in the selection of descriptors in the DNA sequences derived from the fact that only a limited number of features that have chemical/physical meaning are useful for model construction and from difficulties in converting DNA sequences into vectors of numbers. Importantly, nucleotides, being categorical variables, cannot be treated in terms of ordinal data. Since "hand-crafted" features extracted from DNA sequences did not improve prediction of primase binding scores (Supplementary Figure S1), we utilized the K-mer method for feature extraction[18]. In brief, the K-mer is a frequency vector that counts all possible combinations of short sequences (of size K) in larger DNA sequences. As the K parameter increases, the number of possible combinations increases, while the frequency of each mer decreases (Supplementary Figure S2), giving sparser, yet more detailed, data. Since the K-mer method can be implemented with different normalization and striding techniques, even with insufficient structural information, we used a 1-step stride to allow overlap between mers and normalized the extracted K-mer counts.

Extracting features from the DNA sequences of the microarray using the K-mer approach allowed us to find association rules for those DNA sequences (unsupervised algorithms). The sequential features obtained were also used to generate a prediction model of primase–DNA binding, based on primase-binding data collected from PBM experiments (supervised algorithms).

Exploratory data analysis. As is customary, we started with exploratory data analysis, which is unsupervised in nature (i.e., the primase binding scores were ignored). The goal here was to produce a meaningful visualization of the data with the aim to obtain new insights. To this end, we reduced the dimensionality of the data using principal component analysis (PCA), and applied various clustering algorithms, which revealed a meaningful cluster structure with respect to the binding score. After applying PCA to the data, i.e., to the 4^K dimensions existing for each 36-mer oligonucleotide, we found that the top three principal components explained 64% of the total variance. Therefore, we concluded that this number of principal components enables the production of a meaningful 3D picture.

To interpret the clusters generated using PCA, the binding values obtained for the T7 primase of all the DNA sequences were normalized and used to color code the data points in the clusters (Figure 2a). The most striking result to emerge from the color-coded data was its arrangement into five clusters—one homogeneous cluster of DNA sequences that are strongly bound to T7 primase (colored red in Figure 2a), two homogeneous clusters of DNA sequences with weaker binding to T7 primase (colored blue), and two inhomogeneous clusters with uniformly distributed binding scores. This organization of the data points into meaningful clusters indicates that: 1) there are hidden descriptors within the DNA sequence that are essential for primase binding, and 2) the sequence descriptors obtained by the K-mer approach (Figure 2b) are more than adequate for describing primase–DNA binding.

Using the Kmeans algorithm, we were able to shed light on the distributions of the binding scores within the five clusters and to show that preprocessing using OHE results in clusters containing similar score distributions to the 5-cluster structure obtained by PCA. The Kmeans iterative algorithm partitions the data space into sub-spaces, thereby assigning a matching label (the cluster number) to each instance according to its location.

Clustering of the unlabeled DNA sequences in Kmeans relied on the sequence distances from the corresponding cluster centroids. As each sequence was represented using OHE, the distance between two sequences could be described as the number of changes needed in one sequence to convert it into the other sequence. In the Kmeans analysis, aligning the PBM-driven binding score for each DNA sequence revealed that each cluster exhibited a clear trend, as the group of binding scores to the primase was distributed unevenly with long tails (Figure 2c), which means that each cluster held exceptions.

While Kmeans computes distances of instances from the centroids, clustering of unlabeled DNA sequences using Ward's minimum variance method[17] allowed tracking of the evolution of clusters (Figure 2d). The maximal Ward distance (WD) gap was obtained using 5 clusters, as observed from PCA and Kmeans for the same dataset of DNA sequences. Furthermore, we can see in Figure 2d that each

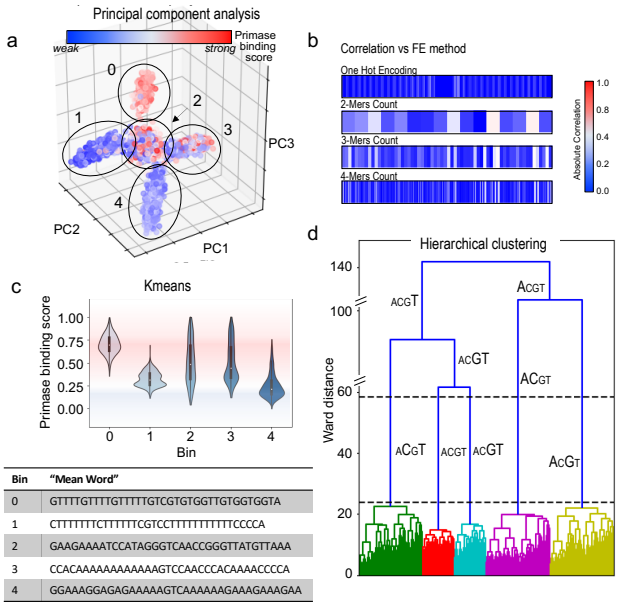

Bin	"Mean Word"
0	GTTTTGTTTTGTTTTTGTCGTGTGGTTGTGGTGGTA
1	CTTTTTTTCTTTTTTCGTCCTTTTTTTTTTTCCCCA
2	GAAGAAAATCCATAGGGTCAACCGGGTTATGTTAAA
3	CCACAAAAAAAAAAAAGTCCAACCCACAAAACCCCA
4	GGAAAGGAGAGAAAAAGTCAAAAAAGAAAGAAAGAA

Figure 2. Inference from DNA sequences without labeled responses for T7 primase binding (unsupervised learning). a, Dimensionality reduction algorithm PCA tri-plot used to visualize the data by selecting the top three principal components. Assignment of the binding scores (labels) to each data point shows an uneven distribution across two clusters (2, 3) and a homogeneous distribution in three clusters (0, for strong binding to T7 primase, and 1 and 4 for weak binding). **b**, Correlation between the primase binding score and feature extraction (FE) using different methods: OHE, 2-mer, 3-mer and 4-mer counts. K-mers were used as descriptors for the PCA analysis. **c,** Kmeans clustering on one-hot encoded DNA sequences. Clustering was performed by measuring pairwise distances of DNA sequences from the centroid of each cluster. Violin plots representing the distribution of the binding scores assigned to each data point in the clusters are shown. Three clusters show evenly distributed scores (0, 1, 4) and two show a less homogeneous score distribution (2, 3). The "mean word" (centroid) of each cluster is presented in the table. **d**, Hierarchical clustering. The x-axis shows DNA sequences emerging into clusters, and the y-axis shows the induced Ward distance of each stage. Letter sizes indicate the letter's frequency in each cluster. The maximal Ward distance gap is indicated between the dashed black lines. The figures were created using the Python package Seaborn and Mat-

colored branch holds a sub-group with two highly repeating letters (CT, GT, AC, AG) or a uniform distribution of the letters (ACGT).

Predicting the binding score. After clustering the DNA sequences in the microarray into groups with common features, the next step was to predict the outcome of T7 primase–DNA binding for a given DNA sequence. To increase accuracy, each cluster was fitted separately. We used GTC-containing DNA sequences and their corresponding PBM-driven binding scores as input and output pairs for the training set, respectively. The PBM-driven data comprises the continuous numerical binding values for DNA sequences, i.e., the type of data that regression models are aimed to solve.

We extracted sequence-based features (SBF) inspired by pseudo K-tuple nucleotide composition[19]. We modified the ordinary method for SBF extraction by neglecting locality-based features, since distance was interpreted as the number of nucleotides between a mer and its closest 5′-GTC start site. Our modified method for feature extraction provided us with sequence-wise normalized K-mers. We then tested different regression algorithms using criteria that can differentiate between uninformative and highly informative SBFs. Thereafter, we applied the L1 regularized linear regression [least absolute shrinkage and selection operator (Lasso)[20]] model on each bin separately with the aim to emphasize meaningful mers and to prevent overfitting of the model. (Lasso's output is a closed form expression that is generated under fewer, yet more meaningful, coefficients by applying a penalty for each variable.) In addition, we extracted an expected performance measure separately for each bin obtained by Kmeans, as the mean absolute errors (MAE, an error estimate for the regression) for our results were uneven across bins (Table 1). Lasso's performance differed across bins, where bins 0 and 1 each generated an error of about 15%, and bins 2, 3 and 4 generated a relatively small error of 6% (Table 1). We found that pre-treating the data using Kmeans decreased the prediction error by approximately 10% (Table 2).

Table 1. Results summary of 5-Fold-MCCV in each clusters

Bin	MAE	STD	Mean MAE [%]	Mean STD [%]	Bin Weight [% of instances from data]
0	0.168	0.005	16.8	0.5	15%
1	0.138	0.007	13.8	0.7	19.4%
2	0.067	0.004	6.7	0.4	13.27%
3	0.060	0.003	6	0.3	25.3%
4	0.077	0.003	7	0.3	26.8%

Table 2. Effect of Kmers on model prediction

MAE	K=1	K=2	K=3	K=4
BNS = 1	0.171	0.112	0.103	0.084
BNS = 5	0.110	0.102	0.093	0.079
Ratio $\frac{error\ no\ bins}{error\ with\ bins}$	1.55	1.09	1.10	1.06
Decrease in error in % $100 * \left(1 - \frac{error\ with\ bins}{error\ no\ bins}\right)$	35%	9%	10%	6%

To force all the coefficients of the Lasso model to be positive for every K-mer feature, we extracted the largest 10 coefficients of each bin and investigated the effect of the K-mer features on the model. The trained models were cross validated on 5-folds of the training dataset and tested on a small test-set taken from the PBM results; the set consisted of 16 sequences, divided into two groups with significantly different primase binding signals (Figure 3a), namely, 8 sequences that showed weak binding to primase, and 8 sequences with strong binding. Given the training data distribution, our models predicted binding of primase with any GTC-containing sequence, with an MAE of < 12% (Figure 3a).

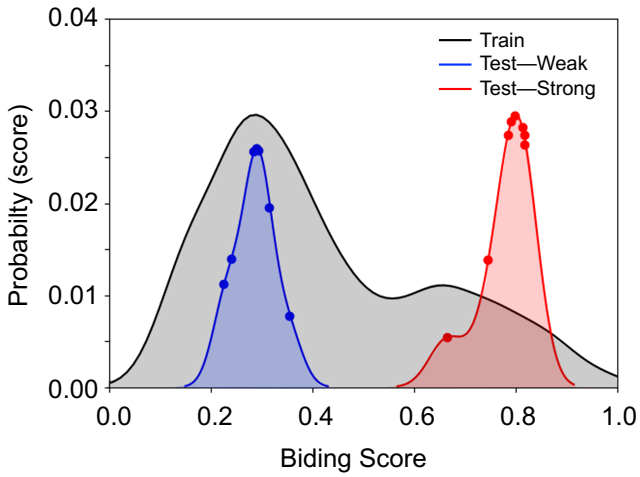

Figure 3 | Results for linear regression with L1 regularization (Lasso). After cleaning, the training set contained 3150 instances (DNA sequences), whereas the test set contained 16 instances. Prediction of scores by using the regression model was performed on 16 DNA sequences with known scores, eight of which showed weak binding to T7 primase (blue graph) and eight showed strong binding to T7 primase (red graph). In accordance with the training-set double distribution (black graph), the predicted binding of the two test groups are distributed at weak and strong binding scores areas, respectively. Although the probability of finding DNA sequences with strong binding to primase is low, the model accurately predicted all DNA sequences that belong to the strong binding group. DNA sequences, their empirical and predicted scores are presented in Supplementary Table S1.

Biochemical validation. On the genome, the initial step of sequence specific (PDRS) binding is followed by synthesis of a dinucleotide (5′-AC-3′), which is then extended into a functional primer by DNA primase. It has previously been shown that A/G-containing sequences that flank the specific recognition site increase primase-DNA binding affinity in comparison to T/G-containing sequences[14]. Since binding to DNA is a pre-requisite for primase activity, the strength of binding affects the magnitude of the catalytic activity of the primase and the yield of the RNA product. We used qualitative biochemical assays to experimentally validate the prediction model described above. The validation provided insights into the features embedded in the DNA sequences that are important for binding and catalytic activity of DNA primase.

I. Correlation between prediction of primase binding to PDRS and catalytic activity. The eight sequences with strong binding and the eight with weak binding to primase, which were used as the test set in the supervised learning part of this study, yielded RNA primers, as was expected from their PBM-driven binding values. This finding shows, for the first time, that the sequence descriptors embedded in the DNA sequence are sufficient to predict binding scores and that prediction of a binding sequence correlates well (96.9 % Pearson correlation coefficient) with the formation of RNA primers (Supplementary Figure S3). The understanding of how sequential features embedded in the DNA is related to the binding of primase allows us not only to predict binding scores of a given PDRS, but also to design novel PDRSs that yield high primase binding scores.

II. Exhaustive search for flanking sequences that yield novel PDRSs. Features important for DNA-primase binding were used in formulating design principles to generate novel GTC-containing DNA sequences with desired binding scores. Assuming that the DNA sequences originate from 5 different clusters (Figure 2) that require 5 different models, we generated two types of DNA sequences as follows: 1) We selected DNA sequences from two homogenous clusters of primase binding scores (Figure 2) and exhaustively altered the non-"GTC" nucleotides to generate primase recognition DNA sequences (new PDRSs). Three altered sequences that did not exist in the training set and yielded the 10th, 50th, and 90th percentile binding strengths were selected from the two clusters (clusters 0 and 4, Figure 2) for further biochemical evaluation. 2) DNA sequences were generated in the same way, and two novel DNA sequences from each Kmean cluster, one that represented the strongest binding prediction and another that represented weakest, were selected (overall 10 sequences. DNA sequence and their predicted scores are presented in Supplementary Table S2) for biochemical analysis (Supplementary Figure S4).

To characterize the effect of the novel DNA sequences (PDRSs) generated as described above, we quantified and compared RNA primer formation by T7 primase, where the generated PDRSs were used as templates for the synthesis (Figure 4). Specifically, we used [γ–^{32}P]ATP to 5′ end-label the RNA primers, which ensured that each primer was labeled only once, and thus the absolute amounts of RNA primers could be quantified (Figure 4). For the 10 DNA templates that represented weak/strong binding to primase from each one of the five clusters (Figure 2), we found that the newly designed DNA sequence flanking the 5′-GTC-3′ sequence with higher binding scores for T7 primase showed improved RNA primer synthesis activity, as was to be expected (Supplementary Figure S4). These results confirm our machine learning prediction model and indicate that higher binding affinity for PDRS recognition sequences are dictated by features embedded in the DNA sequence. Using the exhaustive search for flanking sequences that yield novel primase-DNA binding sequences developed here, we are now in a position to design DNA templates that yield: 1) larger amounts of RNA primers, and 2) longer RNA primers that can serve as functional primers for T7 DNA polymerase [21]. Both the length and the quantity of the RNA primers are likely to be essential for the decision to start Okazaki fragments by DNA polymerase on the lagging DNA strand.

CONCLUSIONS

On the basis of PBM data for primase binding previously obtained for > 150,000 DNA sequences, this study set out to: 1) develop the means to predict the binding score of T7 primase for a given DNA sequence, 2) describe the DNA sequence features essential for binding of the enzyme, and 3) generate novel sequences with a high propensity for T7 primase binding. The K-mers approach for feature selection in the DNA sequential data that was used here appears to cover all possible combinatorial pieces of information hidden in the DNA sequences and serves as an efficient strategy for feature extraction. The Kmers method, which simply counts explicit combinations of nucleotides in a DNA sequence, was superior to other accepted methods of “hand-

Sample			SEQUENCE
P			GGGTCAAAAAAAAAA
Min.			GACGAAGACGACGAAGAGTCCGAGGAAGCAGACGAA
Max			TGTGGTGGGTGTGTGTGGTCTTTTGTTTGTGGTGGT
	BINDING (%)	PREDICTION	SEQUENCE
0-10	10	0.588396629	TGGGTTGTGTGGATTTTGTCGGTGTGGGTGTTTGTG
0-50	50	0.642564063	TGGGTTTGTGGGGTTGTGTCGGGTTGTTGTTTGGTG
0-90	90	0.825284774	TTGTGGTGTGGTGTTTGGTCTGGTGTGTCTTTTGGG
4-10	10	0.107192565	AGAGAAGAGAGAAGAGAGTCGTAGAGAGAGAGAGAG
4-50	50	0.228822138	AGGAAAGAAGAGGAGAGGTCGGAGTAAAGGAAGAAG
4-90	90	0.36810782	GGGGAAAGAGGGAGGAGGTCAGGAAAAGAAAGGGAA

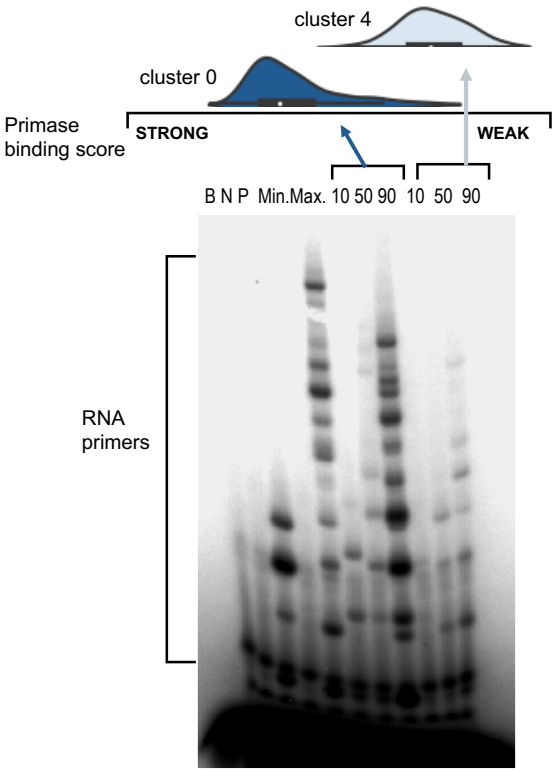

Figure 4. **RNA primer synthesis catalyzed by the T7 primase on computer generated GTC-containing DNA templates.** Three DNA sequences from Kmeans clusters #0 and #4 (distributions are presented in the middle) that show the 10th, 50th, and 90th percentile binding scores were selected in each cluster. Note that cluster #0 shows stronger primase binding values, on average, than cluster #4. The standard reaction mixture contained oligonucleotides with the primase recognition sequence, a control oligonucleotide 5'-GGGTCA10-3', and γ-^{32}P-ATP, CTP, GTP, and UTP. After incubation, the radioactive products were analyzed by electrophoresis on a 25% polyacrylamide gel containing 7 M urea, and visualized using autoradiography. Table present the DNA sequences and their corresponding binding values.

crafted" feature extraction from DNA sequences. Features obtained by Kmers clearly bear DNA properties important for primase binding, as demonstrated by the unsupervised analysis in which clustered groups coincided with experimental binding scores.

Although this study focused on DNA sequence recognition by T7 primase, the findings may well have bearing on rules hidden in DNA sequences that are crucial for other specific DNA-protein interactions. These findings thus contribute to our understanding of how DNA primase selects Okazaki fragments start sites on the genome and why only some of the possible priming sites initiate during DNA replication, while others do not, resulting in Okazaki fragments with a larger-than-expected average length. The implications of this study are that design principles for any DNA sequence with a desired binding affinity to T7 primase can indeed be generated computationally on the basis of our analysis. Furthermore, PDRSs could be designed to yield an RNA primer with a particular content. In conclusion, state-of-the-art carefully selected learning methods, like those used here, have enormous analytical potential for predicting specific protein-DNA interactions, but require large amounts of data, a requirement than can indeed be met by using PBMs.

ASSOCIATED CONTENT

Supporting Information

Experimental and computational methods, supplementary tables and figures.

AUTHOR INFORMATION

Corresponding Author akabayov@bgu.ac.il

Author Contributions

The manuscript was written through contributions of all authors. All authors have given approval to the final version of the manuscript.

Funding Sources

This research was supported by the Israel Science Foundation (ISF, Grant No. 1023/18).

ABBREVIATIONS

PBM, protein-DNA binding microarray; PDRS, primase-DNA recognition sequence; EDA, exploratory data analysis; PCA, principal component analysis; WD, Ward distance.

REFERENCES

(1) Kornberg, A. & Baker, T. A. *DNA replication*. 2nd edn, (University Science Books, 2005).

(2) Frick, D. N. & Richardson, C. C. DNA primases. *Annual review of biochemistry* **70**, 39-80, doi:70/1/39 [pii]10.1146/annurev.biochem.70.1.39 (2001).

(3) Stratling, W. & Knippers, R. Function and purification of gene 4 protein of phage T7. *Nature* **245**, 195-197, doi:10.1038/245195a0 (1973).

(4) Wolfson, J. & Dressler, D. Regions of single-stranded DNA in the growing points of replicating bacteriophage T7 chromosomes. *Proceedings of the National Academy of Sciences of the United States of America* **69**, 2682-2686, doi:10.1073/pnas.69.9.2682 (1972).

(5) Richardson, C. C. *et al.* Replication of bacteriophage T7 DNA by purified proteins. *Cold Spring Harbor symposia on quantitative biology* **43 Pt 1**, 427-440 (1979).

(6) Lee, S. J., Zhu, B., Hamdan, S. M. & Richardson, C. C. Mechanism of sequence-specific template binding by the DNA primase of bacteriophage T7. *Nucleic acids research* **38**, 4372-4383, doi:10.1093/nar/gkq205 (2010).

(7) Corn, J. E., Pease, P. J., Hura, G. L. & Berger, J. M. Crosstalk between primase subunits can act to regulate primer synthesis in trans. *Molecular cell* **20**, 391-401, doi:S1097-2765(05)01604-7 [pii]10.1016/j.molcel.2005.09.004 (2005).

(8) Corn, J. E., Pelton, J. G. & Berger, J. M. Identification of a DNA primase template tracking site redefines the geometry of primer synthesis. *Nature structural & molecular biology* **15**, 163-169, doi:10.1038/nsmb.1373 (2008).

(9) Andrilenas, K. K., Penvose, A. & Siggers, T. Using protein-binding microarrays to study transcription factor specificity: homologs, isoforms and complexes. *Briefings in functional genomics* **14**, 17-29, doi:10.1093/bfgp/elu046 (2015).
(10) Soultanas, P. The bacterial helicase-primase interaction: a common structural/functional module. *Structure* **13**, 839-844, doi:10.1016/j.str.2005.04.006 (2005).
(11) Thirlway, J. & Soultanas, P. In the Bacillus stearothermophilus DnaB-DnaG complex, the activities of the two proteins are modulated by distinct but overlapping networks of residues. *Journal of bacteriology* **188**, 1534-1539, doi:10.1128/JB.188.4.1534-1539.2006 (2006).
(12) Naue, N., Beerbaum, M., Bogutzki, A., Schmieder, P. & Curth, U. The helicase-binding domain of Escherichia coli DnaG primase interacts with the highly conserved C-terminal region of single-stranded DNA-binding protein. *Nucleic acids research* **41**, 4507-4517, doi:10.1093/nar/gkt107 (2013).
(13) Chintakayala, K. *et al.* Domain swapping reveals that the C- and N-terminal domains of DnaG and DnaB, respectively, are functional homologues. *Molecular microbiology* **63**, 1629-1639, doi:10.1111/j.1365-2958.2007.05617.x (2007).
(14) Afek, A. *et al.* DNA Sequence Context Controls the Binding and Processivity of the T7 DNA Primase. *iScience* **2**, 141-147, doi:10.1016/j.isci.2018.03.019 (2018).
(15) Ilic, S. *et al.* DNA Sequence Recognition by DNA Primase Using High-Throughput Primase Profiling. *J Vis Exp*, doi:10.3791/59737 (2019).
(16) Balakrishnan, L. & Bambara, R. A. Okazaki fragment metabolism. *Cold Spring Harbor perspectives in biology* **5**, doi:10.1101/cshperspect.a010173 (2013).
(17) Ward, J. H. J. Hierarchical Grouping to Optimize an Objective Function. *J Am Stat Assoc* **58**, 236–244 (1963).
(18) Chen, W., Lei, T. Y., Jin, D. C., Lin, H. & Chou, K. C. PseKNC: a flexible web server for generating pseudo K-tuple nucleotide composition. *Anal Biochem* **456**, 53-60, doi:10.1016/j.ab.2014.04.001 (2014).
(19) Dao, F. Y., Lv, H., Wang, F. & Ding, H. Recent Advances on the Machine Learning Methods in Identifying DNA Replication Origins in Eukaryotic Genomics. *Front Genet* **9**, 613, doi:10.3389/fgene.2018.00613 (2018).
(20) Tibshirani, R. Regression Shrinkage and Selection via the lasso. *Journal of the Royal Statistical Society* **58**, 267–288 (1996).
(21) Romano, L. J. & Richardson, C. C. Characterization of the ribonucleic acid primers and the deoxyribonucleic acid product synthesized by the DNA polymerase and gene 4 protein of bacteriophage T7. *The Journal of biological chemistry* **254**, 10483-10489 (1979).
(22) Zhu, B., Lee, S. J. & Richardson, C. C. An in trans interaction at the interface of the helicase and primase domains of the hexameric gene 4 protein of bacteriophage T7 modulates their activities. *The Journal of biological chemistry* **284**, 23842-23851, doi:10.1074/jbc.M109.026104 (2009).
(23) Pedregosa, F. *et al.* Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* **12**, 2825-2830 (2011).
(24) Brown, C. E. *Coefficient of Variation*. (Springer, 1998).