

Predicting higher-order mutational effects in an RNA enzyme by machine learning of high-throughput experimental data

1 **James D. Beck**^{1†}, **Jessica M. Roberts**^{2†}, **Joey Kitzhaber**³, **Ashlyn Trapp**⁴, **Edoardo Serra**¹,
2 **Francesca Spezzano**¹, **Eric J. Hayden**^{2,3,*}

3 ¹Computing PhD Program, Boise State University, Boise, ID, USA

4 ²Biomolecular Sciences Graduate Programs, Boise State University, Boise, ID, USA

5 ³Department of Computer Science, Boise State University, Boise, ID, USA

6 ⁴Department of Biological Sciences, Boise State University, Boise, ID, USA

7 [†]Authors contributed equally to this work and share first authorship

8 * **Correspondence:** Eric J. Hayden erichayden@boisestate.edu

9 **Keywords: ribozyme, fitness landscape, RNA, epistasis, machine learning, long short-term**
10 **memory, random forest. (Min.5-Max. 8)**

11 **Abstract**

12 Ribozymes are RNA molecules that catalyze biochemical reactions. Self-cleaving ribozymes are
13 a common naturally occurring class of ribozymes that catalyze site-specific cleavage of their
14 own phosphodiester backbone. In addition to their natural functions, self-cleaving ribozymes
15 have been used to engineer control of gene expression because they can be designed to alter
16 RNA processing and stability. However, the rational design of ribozyme activity remains
17 challenging, and many ribozyme-based systems are engineered or improved by random
18 mutagenesis and selection (in vitro evolution). Improving a ribozyme-based system often
19 requires several mutations to achieve the desired function, but extensive pairwise and higher-
20 order epistasis prevent a simple prediction of the effect of multiple mutations that is needed for
21 rational design. Recently, high-throughput sequencing-based approaches have produced data sets
22 on the effects of numerous mutations in different ribozymes (RNA fitness landscapes). Here we
23 used such high-throughput experimental data from variants of the CPEB3 self-cleaving ribozyme
24 to train a predictive model through machine learning approaches. We trained models using either
25 a random forest or long short-term memory (LSTM) recurrent neural network approach. We
26 found that models trained on a comprehensive set of pairwise mutant data could predict active
27 sequences at higher mutational distances, but the correlation between predicted and

1 experimentally observed self-cleavage activity decreased with increasing mutational distance.
2 Adding sequences with increasingly higher numbers of mutations to the training data improved
3 the correlation at increasing mutational distances. Systematically reducing the size of the training
4 data set suggests that a wide distribution of ribozyme activity may be the key to accurate
5 predictions. Because the model predictions are based only on sequence and activity data, the
6 results demonstrate that this machine learning approach allows readily obtainable experimental
7 data to be used for RNA design efforts even for RNA molecules with unknown structures. The
8 accurate prediction of RNA functions will enable a more comprehensive understanding of RNA
9 fitness landscapes for studying evolution and for guiding RNA-based engineering efforts.

10

11 **Introduction**

12 RNA enzymes, or ribozymes, are structured RNA molecules that catalyze biochemical reactions.
13 One well-studied class of ribozymes are the small self-cleaving ribozymes that catalyze site
14 specific cleavage of phosphate bonds in their own RNA backbone (Ferré-D'Amaré and Scott,
15 2010). These self-cleaving ribozymes are found in all domains of life, and their biological roles
16 are still being investigated (Jimenez et al., 2015). In addition to their natural functions, these
17 ribozymes have been used as the basis for engineering biological systems. For example, several
18 small ribozymes (hammerhead, twister, pistol and HDV) have been used as genetically encoded
19 gene regulatory elements by combining them with RNA aptamer and embedding them into
20 untranslated regions of genes (Groher and Suess, 2014; Dykstra et al., 2022). This approach
21 continues to gain attention because of the central importance of controlling gene expression and
22 the simple design and build cycles of these small RNA elements. Nevertheless, ribozymes often
23 need optimization for sequence dependent and cell specific effects. This can be achieved by
24 modifying the sequence of the ribozymes, but this often requires multiple mutational changes
25 and the vast sequence space requires extensive trial and error. Given this large sequence space,
26 even the most high-throughput approaches can only find the optimal solutions present in the
27 sequences that can be explored experimentally, which is a fraction of the total possible
28 sequences. The engineering of ribozyme-based systems could benefit from accurate prediction of

1 the effects of multiple mutations in order to narrow the search space towards optimal collections
2 of sequences.

3 One way to think of the ribozyme optimization problem is in terms of fitness landscapes.
4 Molecular fitness landscapes of protein and RNA molecules are studied by measuring the effects
5 of numerous mutations on the function of a given reference molecule (Athavale et al., 2014;
6 Blanco et al., 2019). Recently, the fitness landscapes of RNA molecules have been studied
7 experimentally by synthesizing large numbers of sequences and using high-throughput
8 sequencing to evaluate the relative activity of the RNA in vitro, or the growth effect of the RNA
9 in a cellular system, both of which are termed “RNA fitness” (Kobori and Yokobayashi, 2016; Li
10 et al., 2016; Pressman et al., 2019). The goal of in vitro evolution is often to find the highest peak
11 in the landscape, or one of many high peaks, by introducing random mutations and selecting for
12 improved activity. However, the RNA fitness landscapes that have been experimentally studied
13 so far have revealed rugged topographies with peaks of high relative activity and adjacent valleys
14 of low activity. Landscape ruggedness is an impediment to finding desired sequences through in
15 vitro evolution approaches (Ferretti et al., 2018). Epistasis, defined as the non-additive effects of
16 mutations, is the cause of ruggedness in fitness landscapes, and epistasis has been used to
17 quantify the ruggedness of fitness landscapes (Szendro et al., 2013). More frequent and more
18 extreme epistasis indicates that a landscape is more rugged. Importantly, more epistasis also
19 means that the effect of combining multiple mutations is challenging to predict even if the effects
20 of each individual mutation are known. In addition, experimental fitness landscapes can only
21 study a limited number of sequences, except for very small RNA molecules (Pressman et al.,
22 2019). It is often not possible to know if the process of in vitro evolution discovered a sequence
23 that is globally optimal, or just a local optimum. For these reasons, it has become a goal to
24 accurately predict the activity of sequences in order to streamline RNA evolution experiments
25 and to study fitness landscapes in a more comprehensive manner (Groher et al., 2019; Schmidt
26 and Smolke, 2021).

27 Here, we use high-throughput experimental data of mutational variants of a self-cleaving
28 ribozyme to train a model for predicting the effect of higher-order combinations of three or more
29 mutations. The ribozyme used in this study is the CPEB3 ribozyme (Figure 1A). This ribozyme
30 is highly conserved in the genomes of mammals, where it is found in an intron of the CPEB3

1 gene (Salehi-Ashtiani et al., 2006). For training purposes, we generated a new data set that
2 includes all possible individual and pairs of mutations to the reference CPEB3 ribozyme
3 sequence (Figure 1B). These mutations were made by randomization of the CPEB3 ribozyme
4 sequence with a 3% per nucleotide mutation rate during chemical synthesis of the DNA
5 template. We reasoned that given the extensive amount of pairwise epistasis in RNA (Bendixsen
6 et al., 2017), this data set might be sufficient for predicting higher-order mutants. In addition, we
7 used a second, previously published data set that included 27,647 sequences comprised of
8 random permutations of mutations found in mammals that include up to 13 mutational
9 differences from the same reference ribozyme (Bendixsen et al., 2021). This second data set not
10 only contains higher-order mutational combinations, but also a broad range of self-cleaving
11 activity (Figure 1D). In both data sets, the relative activity of each sequence was determined by
12 the deep sequencing of co-transcriptional self-cleavage data, as previously described. Briefly, the
13 mutated DNA template was transcribed in vitro with T7 RNA polymerase. The transcripts were
14 prepared for Illumina sequencing by reverse transcription and PCR. Relative activity was
15 determined as the *fraction cleaved*, defined by the fraction of sequencing reads that mapped to a
16 specific sequence variant in the shorter, cleaved form relative to the total number of reads for
17 that sequence variant.

18 We set the goal of being able to predict the activity of the higher-order mutants in the
19 phylogenetically derived fitness landscape (Figure 1D). In addition, we wanted to guide future
20 experiments aimed at producing additional data for training models of ribozyme-based systems.
21 The number of possible sequences increases exponentially with the number of variable
22 nucleotide positions. In addition, the probability of finding active ribozymes at higher mutational
23 distances becomes increasingly unlikely. Experiments aimed at training predictive models will
24 need to choose realistic numbers of sequences that can have the highest impact on model
25 performance. We therefore evaluated the effect of adding to the training data sequences with
26 increasing mutational distances from the wild-type sequence as well as the effect of reducing the
27 number of sequences in the training data. The results of these experiments were expected to be
28 useful in guiding the choice of which sequence variants, and how many, to analyze
29 experimentally in order to produce effective training data sets.

30

1 **Results**

2 We first evaluated our new training data set that contained all single and double-mutants of the
3 CPEB3 ribozyme. We found that the data did in fact contain full coverage of the possible 207
4 single mutants and the 21,114 double mutants. While the number of reads that mapped to each of
5 these sequences varied, we found that, on average, 170 reads mapped to each double mutant, and
6 ~18,000 reads mapped to each single mutant (Supplemental Figure 1). This read depth was
7 sufficient for the determination of the fraction cleaved for all single and double mutants (Figure
8 1B). Mapping the fraction cleaved to base paired structural elements showed expected patterns of
9 activity caused by compensatory base pairs. Mutations that break a base pair typically showed
10 low activity, but a second mutation that restored the base pair showed high activity. To further
11 evaluate this data, we calculated the non-additive pairwise epistasis in this data set (Figure 1C).
12 Together, this analysis indicated that this data set contained a wide range of ribozyme activity
13 and the effects of all pairwise intramolecular epistatic interactions.

14 In order to determine the training potential of the comprehensive double-mutant data, we first
15 trained models using only the fraction cleaved data for sequences with two or fewer mutations
16 including the wild-type reference sequence. We then tested the models' performance in
17 predicting the fraction cleaved for sequences with increasing numbers of mutations. We trained
18 two models with two approaches (see Materials and Methods). The first approach used a
19 Random Forest regressor. In the second approach, we added a Long Short-Term Memory
20 (LSTM) recurrent neural network to extract hidden features from the data. We then fed the
21 hidden features with associated fraction cleaved to a Random Forest regressor. We will refer to
22 this approach as "LSTM". We found that models trained on 2 or fewer mutations with Random
23 Forest outperformed LSTM at predicting the activity of sequences with five or fewer mutations
24 (Figure 2 A-C), but LSTM performed better when predicting the activity of sequences with six or
25 more mutations relative to the wild-type (Figure 2 D-I). However, both approaches showed a
26 decrease in the correlation between predicted and observed when challenged to predict the
27 activity of sequences with higher numbers of mutations, and both resulted in relatively low
28 correlation (Pearson $r < 0.7$) for sequences with seven or more mutations when trained only on
29 this double mutant data (Figure 2 and Supplementary Table 2). We concluded that models
30 trained on simple random mutagenesis containing all double mutants can be useful for predicting

1 lower mutational distances, but we anticipated that additional data might improve the ability to
2 predict the effect of higher numbers of mutations.

3 To determine the effect of adding higher-order mutants to the training data, we divided the
4 phylogenetic derived sequence data by mutational distance and re-trained models with increasing
5 orders of mutations in the training set. As expected, adding higher-order mutants improved the
6 predicted to observed correlation at higher mutational distances (Figure 3 and Supplemental
7 Figures 2-14). Interestingly, we found that the Random Forest approach outperformed the LSTM
8 approach when sequences with more mutations were included in the training data. This is
9 especially apparent for predicting the activity of sequences with 8-10 mutations. The Random
10 Forrest approach resulted in models with high correlation between predicted and observed for all
11 mutational distances when trained with data from sequences with four or more mutations (Figure
12 3 A-C). For both approaches, the largest improvements in the correlations occurred when
13 sequences with 3 mutations (relative to wild-type) were added to the data. Subsequently
14 appending additional sequences with greater numbers of mutations had diminishing
15 improvements on the correlation. We note that all the testing data was set aside prior to training
16 and identical testing data was used for all models. The results demonstrate that adding higher
17 order mutants to the training data improves the Pearson correlation of sequences at higher
18 distances in this data set. It is important to note that the phylogenetically derived data has
19 different numbers of sequences for each class of mutations (Table 1), and sequences with higher
20 numbers of mutations in our data show mostly low activity (Supplementary Figure 15). This
21 helps interpret the effect of sequentially adding higher-order mutant sequences to the training
22 data. It is also important to note that the phylogenetic derived sequences only contain mutations
23 at thirteen different positions. The higher order sequences in this data are therefore combinations
24 of the lower order sequences. For example, a sequence with six mutations can be constructed by
25 combining two sequences with three mutations, both of which would be in the “3 mutations”
26 training data. Our model is therefore predicting the effects of combining sets of mutations, and
27 adding precise sets of lower order mutations that re-occur in higher order mutations clearly
28 improves the correlations between prediction and experimental observation in our data.

29 In order to inform future experiments for collecting training data, we next set out to determine
30 the effect of decreasing the amount of data in the training sets. Starting from the 80% of data

1 used as prior training data, we randomly sampled sequences from this data to create new training
2 data sets with 60%, 40%, 20%, 10% and 1% of the total data. These subsampled data sets were
3 used to train models using the random forest regressor. The same testing data was set aside for
4 all models and used to compare the Pearson correlation coefficient of each model trained with
5 decreasing amounts of data. As an illustrative example, we focused on a model trained with
6 sequences with 5 or fewer mutations relative to wild-type used to predict the activity of
7 sequences with 7 mutations (Figure 4 and Supplementary Table 1). We chose this example
8 because it achieved very high correlation (Pearson $r = 0.99$) when trained with 80% (25,733
9 unique sequences) of the data and therefore provided an opportunity to observe how rapidly the
10 correlation decreased with less data. We found that the models trained on 5 or fewer mutations
11 predicted with high correlation when as little as 40% (12,866) of the data was used for training
12 (Pearson $r = 0.97$). With only 20% (6,433) and 10% (3,217) of the data, the model still showed
13 good prediction accuracy with a Pearson correlation $r \cong 0.9$. Surprisingly, we still observed
14 reasonably high correlation when including only 1% (322) of the training data, and this was
15 reproducible over five different models trained with different random samples of the data
16 (Pearson $r = 0.81$, $stdev = 0.046$, $n = 5$). Similar results were observed with other training and
17 testing scenarios. To illustrate general trends, we have plotted the Pearson correlation for the
18 same model trained on 5 or fewer mutations when predicting the activity of sequences with 6, 7,
19 8 or 9 mutations, and for a model trained on 9 or fewer mutations used to predict sequences with
20 5, 6, 7, or 8 mutations (Figure 4). This analysis suggests that the total amount of training data is
21 not critical for predicting the activity of sequences in our data set. When combined with the
22 diminishing returns of adding more higher order mutations (Figure 3), this analysis emphasizes
23 the importance of collecting appropriate experimental data sets for training that include
24 ribozymes with more mutations that still maintain relatively high activity. However, given the
25 low probability of finding higher-order sequences with higher activity, an iterative approach with
26 several cycles of predicting and testing might be necessary to acquire such data.

27 While the primary goal was to predict the relative activity of RNA sequences, we wondered if
28 the models might also be useful for predicting structurally important nucleotides. To address this
29 question, we analyzed the “feature importance” in several of our Random Forest models. Feature
30 importance is a method to assign importance to specific input data. Because our data only uses

1 sequence as input, the features in our data are specific nucleotides (A, G, C or U) at specific
2 positions. We found that for the Random Forest models, the most important feature all clustered
3 around the active site of the ribozyme (Supplemental Figures 16 and 17). Further, the CPEB3
4 ribozyme uses metal ion catalysis and several of the most important features were nucleotides
5 that have been observed coordinated to the active site magnesium ion in the CPEB3 ribozyme, or
6 the analogous nucleotides in the structurally similar HDV ribozyme (Kapral et al., 2014;
7 Skilandat et al., 2016). For example, for all the models trained with some higher order mutants,
8 the most important feature was G1, which positions the cleaved phosphate bond in contact with
9 the catalytic magnesium ion. The second most important feature was G25, which forms a wobble
10 base pair with U20 (Lévesque et al., 2012), another important feature (top 4-6), and this
11 nucleotide pair coordinates the active site magnesium ion through outer sphere contacts. The
12 catalytic nucleotide C57 binds the same catalytic magnesium as the G25:U20 wobble pair, and
13 had a high feature importance similar to U20. Most of the other important features are involved
14 in base pairs that stack or interact with the metal ion coordinating bases. Interestingly, we found
15 that nine of the ten most important features were identical for models trained with only single
16 and double mutants or with increasing amounts of higher-order mutants. However, the G1 and
17 G25 features became increasingly more important as sequences with higher mutational distance
18 were added to the training data. This indicates that the higher-order mutants in the training data
19 helped emphasize structurally critical nucleotides. We conclude that the machine learning
20 models presented identified nucleotides involved in forming the active sites of the CPEB3
21 ribozyme. Because we did not use structural data to train our models, the results suggest that
22 similar data could identify active sites in RNA molecules with unknown structures.

23

24 **Discussion**

25 We have shown that a model trained on ribozyme activity data can accurately predict the self-
26 cleavage activity of sequences with numerous mutations. This approach can be used to guide
27 experiments based on a relatively small set of initial data. Importantly, the approach did not use
28 structural information such as X-ray crystallography or cryo-EM, and used only sequence and
29 activity data, which can be obtained with common molecular biology approaches (in vitro

1 transcription, RT-PCR, and sequencing). In addition, the training data starts with small amounts
2 of synthetic DNA. The comprehensive double mutant data and the phylogenetic derived data
3 each started from a single DNA oligo synthesis that used doped phosphoramidites at the variable
4 positions. Each data set was collected on a single lane of an Illumina sequencer. The approach
5 presented in this paper is therefore accessible, rapid and inexpensive as compared to approaches
6 that use structural data to train their models.

7 Sequence conservation of naturally occurring RNA molecules has been another useful data type
8 for training models to predict RNA structure from sequence (De Leonardis et al., 2015; Weinreb
9 et al., 2016). This approach is based on the observation that nucleotide positions that form a base
10 pair often show co-evolutionary patterns of sequence conservation. In some cases, this co-
11 evolutionary data has been combined with thermodynamic predictions or structural data from
12 chemical probing, such as SHAPE experiments (Calonaci et al., 2020). Numerous ribozymes,
13 aptamers and aptazymes have been discovered through in vitro evolution experiments and
14 conservation data is not available unless sequencing experiments were applied during the
15 selection process. Our approach could be used to expand functional information of non-natural
16 RNA molecules which could then be used to guide structure prediction of these molecules in a
17 way similar to how naturally occurring sequence conservation has been used. In addition,
18 sequence conservation does not necessarily predict relative activity. For example, while the
19 CPEB3 ribozyme is highly conserved in nature, not all of the sequence are equally proficient at
20 catalyzing self-cleavage (Chadalavada et al., 2010; Bendixsen et al., 2021). Our approach using
21 machine learning from experimentally derived data may prove useful for guiding experiments
22 with non-natural RNA molecules discovered through in vitro selection or SELEX-like
23 approaches. However, adopting this machine learning approach will require that each
24 experimenter acquire specific data for their system necessary to train and test sequences with the
25 functions they are investigating.

26 With future work, it may be possible to produce more general models of ribozyme activity. For
27 example, a model trained on data sets from several different self-cleaving ribozymes with
28 different nucleotide lengths might learn to predict the activity of sequences of arbitrary length
29 and sequence composition. In fact, recent advances in RNA structure prediction have used the
30 crystal structures of several different self-cleaving ribozymes as training data to develop

1 predictive modes that achieve near-atomic level resolution of arbitrary sequences (Townshend et
2 al., 2021). Alternatively, models trained on ribozymes with different activities beyond self-
3 cleavage might be able to classify sequences as ribozymes of various functions. There has been
4 some success with generating general models for predicting protein functions. The latent features
5 identified by deep generative models of protein sequences are being used to better understand the
6 complex, higher-order amino acid interactions necessary to achieve a functional protein structure
7 (Riesselman et al., 2018; Detlefsen et al., 2022). We hypothesize that latent features could aid in
8 the identification of generalized parameters that govern the epistatic interactions of higher-order
9 mutants of RNA sequences as well. We hope that the accuracy and accessibility of the approach
10 presented here will inspire others to carry out similar experiments and initiate the data sharing
11 that will be needed to develop more general models, similar to what is being accomplished for
12 protein functional predictions (Biswas et al., 2021).

13 One challenge to our predictive models appears to be the low frequency of active sequences at
14 higher mutational distances. In our phylogenetically derived data the vast majority of sequences
15 have very low activity (Figure 1D), and the probability of finding sequence with high fraction
16 cleaved decreases with the number of mutations relative to wild-type. As a consequence, models
17 trained on lower-order mutant variants tend to overestimate the activity of sequences at higher
18 mutational distances. It has been previously observed that experimental RNA fitness landscapes
19 are dominated by *negative epistasis*, which means that mutations in combination tend to have
20 lower fitness than would be expected from the additive effects of individual mutations
21 (Bendixsen et al., 2017). The overestimation of fraction cleaved at higher mutational distances
22 suggests that our models have a difficult time learning to predict negative epistasis. It has been
23 previously observed that mutations with “neutral” or “beneficial” effects on protein function
24 often have destabilizing effects on protein structure (Soskine and Tawfik, 2010). We postulate
25 that the same effect is causing negative epistasis in the RNA data. This suggests that additional
26 information, such as measurements or estimates of thermodynamic stability of helices, might be
27 necessary for increasing accuracy at even higher distances beyond those offered by this data set
28 (Groher et al., 2019; Yamagami et al., 2019). For example, we have recently demonstrated that
29 our sequencing based approach to measuring ribozyme activity can be extended to include
30 magnesium titrations in order to evaluate RNA folding/stability (Peri et al., 2022). In the future,

1 combining structural and functional information might be the best approach to accurately design
2 RNA molecules with desired functional properties.

3

4 **Materials and Methods**

5 ***Ribozyme activity data***

6 Ribozyme activity was determined as previously described (Bendixsen et al., 2021). Briefly,
7 DNA templates were synthesized with the promoter for T7 RNA polymerase to enable in vitro
8 transcription. Templates were synthesized with mixtures of phosphoramidites at variable
9 positions. For the comprehensive double-mutant data set, templates were synthesized with 97%
10 wild-type nucleotides and 1% each of the other three nucleotides. For the phylogenetic derived
11 data set, the template was synthesized with an equal mixture of the naturally occurring
12 nucleotides that were found at 13 positions that varied across 99 mammalian genomes. During in
13 vitro transcription, RNA molecules self-cleaved at different rates. The reaction was stopped at 30
14 minutes, and the RNA was concentrated and reverse transcribed with a 5'-RACE protocol that
15 appends a new primer site to the cDNA of both cleaved and uncleaved RNA (SMARTScribe,
16 Takara). The cDNA was PCR amplified with primers that add the adaptors for Illumina
17 sequencing. This procedure was done in triplicate with unique dual-indexes for each replicate.
18 DNA was combined equimolar and sent for sequencing (GC3F, University of Oregon).
19 Sequencing was performed on a single lane of a HiSeq 4000 using paired-end 150 reads.

20 ***Ribozyme activity from sequence data***

21 FastQ sequencing data were analyzed using custom Julia and Python scripts. Briefly, the scripts
22 identified the reverse transcription primer binding site at the 3'-end to determine nucleotide
23 positions and then determined if the sequence was cleaved or uncleaved by the absence or
24 presence of the 5'-upstream sequence. For the single and double mutants, all possible sequences
25 were generated and stored in a list, and reads that matched the list elements were counted and
26 cleaved or uncleaved was determined by the presence or absence of the 5'-upstream sequence.
27 For the phylogenetically derived data, nucleotide identities were determined at the expected 13
28 variable positions by counting the string character position from the fixed regions. Sequencing

1 reads were discarded if they contained unexpected mutations in the primer binding site, the
2 uncleaved portion, or the ribozyme sequence. For each unique genotype in the library the number
3 of cleaved and uncleaved sequences were counted and ribozyme activity (fraction cleaved) was
4 calculated as $\text{fraction cleaved} = \text{counts}_{\text{cleaved}} / (\text{counts}_{\text{cleaved}} + \text{counts}_{\text{uncleaved}})$.

5 ***Machine Learning***

6 Random Forest regression uses an ensemble of decision trees to improve prediction accuracy.
7 Each tree in the ensemble is created by partitioning the sequences within a sample into groups
8 possessing little variation. Each sample is drawn with replacement and the resulting trees are
9 aggregated into forests that best predict the cleavage rates of the sequences. The Random Forest
10 regression was performed using the python package scikit-learn. Each sequence was transformed
11 into a 69 by 4 one-hot encoding representation of the sequence. Each of the four possible
12 nucleotides within the sequence was represented by a vector of length 4 possessing a uniquely
13 located “1” within the vector to signify the nucleotide’s identity. Each sequence in the training
14 set was fit using scikit-learn’s RandomForestRegressor ensemble module. Feature importance
15 was computed via a forest of randomized trees using the *features_importances* function in the
16 module under default settings. Briefly, the relative importance of a feature was determined by the
17 depth of the feature when it was used as a decision node in a tree. Features used at the top of the
18 tree contribute to the final prediction decision of a larger fraction of the input samples. The
19 expected fraction of the samples they contributed to was used as an estimate of the relative
20 importance of the features.

21 LSTM is a recurrent neural network commonly used for the predictive modeling of written text
22 data, which has sequential dependencies. Here we used an LSTM to compute a set of hidden
23 features given a set of nucleotide sequences. These hidden features are learned by the LSTM in a
24 supervised way for the purpose of relating the nucleotide sequence to the corresponding
25 ribozyme activity (fraction cleaved). The LSTM network has an architecture where each cell C
26 outputs the next state h_t ($1 \leq t \leq n$) by taking in input from the previous state h_{t-1} and the
27 embedding x_t of the current nucleotide in the sequence. The output h_n of the last cell of the
28 LSTM is then used as input to a Random Forest regressor to predict the sequence functional
29 activity rate. The LSTM model was built using PyTorch’s open-source machine learning

1 framework. Sequences were trained using an LSTM layer with 32 hidden dimensions and a
2 dropout rate of 0.2. Each sequence was embedded in a 69 by 4 tensor (where 4 is the size of the
3 nucleotide embedding) and then batched in groups of 64 sequences for input to the model. The
4 gradient descent was performed using PyTorch's built-in Adam optimizer and MSELoss
5 criterion. Twenty-five training epochs were performed on each training set.

6 ***Training and Test Data***

7 The data set containing the fraction cleaved data from the 27,647 phylogenetically derived
8 sequences was binned based on the number of mutations relative to the wild-type ribozyme. For
9 each bin, a portion of the data (20%) was chosen at random and set aside as test data. This
10 resulted in test data sets that were also separated by the number of mutations relative to the wild-
11 type sequence. Training data sets were created from the 80% of data in each mutational bin that
12 was not set aside for testing. Training data sets were created by combining bins at a given
13 number of mutations to all the bins with lower numbers of mutations. Training data included
14 100% of the single and double mutant data. For reduced training sets were created by randomly
15 sampling different numbers of sequences from the original full training data sets.

16 **Data Availability**

17 Sequencing reads in FastQ format are available at ENA (PRJEB51631). Sequence and activity
18 data and computer code is available at GitLab (<https://gitlab.com/bsu/biocompute-public/ml-ribo-predict.git>).
19

20 **Author Contributions**

21 JB – involved in conceptualizing the project, managed data, performed computational work for
22 formal analysis and visualization, reviewed and edited the manuscript; JR – involved in
23 conceptualizing the project, performed experiments, managed the project, supervised and
24 facilitated computational work, prepared figures, reviewed and edited the manuscript, JK –
25 involved in conceptualizing the project, performed computational work for formal analysis and
26 validation; AT– was involved in conceptualizing the project, helped with experimental
27 validation, supervised and facilitated computational work; ES – Involved in conceptualizing the
28 project, supervised computational work, reviewed and edited the manuscript; FS – Involved in

1 conceptualizing the project, supervised computational work, reviewed and edited the manuscript;
2 EH – involved in conceptualizing the project, supervised experimental work, supervised
3 computational work, wrote the original draft, and reviewed and edited the manuscript.

4 **Funding**

5 The authors acknowledge funding from the National Science Foundation (EH, grant number
6 OIA-1738865, OIA-1826801, REU site grant #1950599), National Aeronautics and Space
7 Administration (EH, grant number 80NSSC17K0738), and the Human Frontier Science Program
8 (EH, Ref.-No: RGY0077/2019).

9 **Acknowledgments**

10 We thank Devin Bendixsen for producing the fitness landscape image used in the manuscript and
11 for valuable discussions.

12 **References**

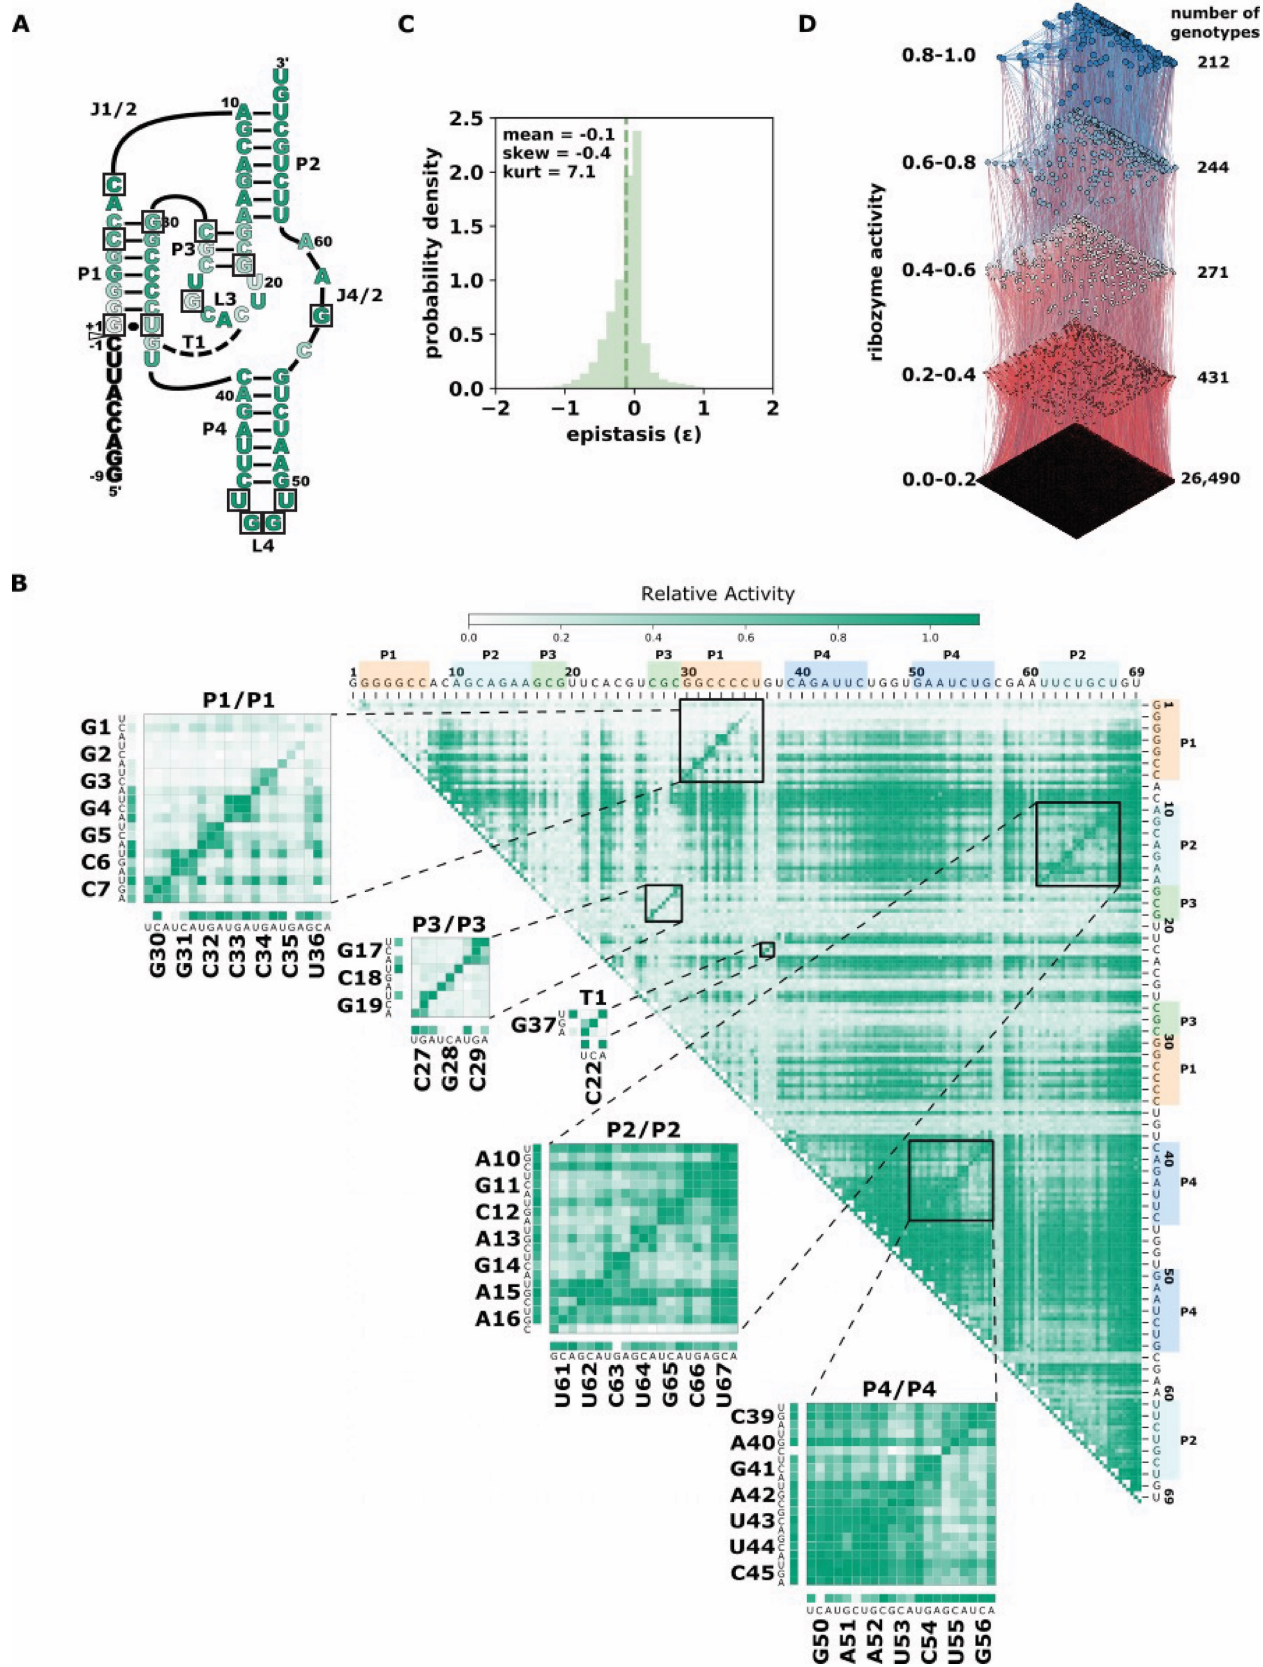
- 13 Athavale, S. S., Spicer, B., and Chen, I. A. (2014). Experimental fitness landscapes to understand
14 the molecular evolution of RNA-based life. *Current Opinion in Chemical Biology* 22,
15 35–39. doi: 10.1016/j.cbpa.2014.09.008.
- 16 Bendixsen, D. P., Østman, B., and Hayden, E. J. (2017). Negative Epistasis in Experimental
17 RNA Fitness Landscapes. *J. Mol. Evol.* doi: 10.1007/s00239-017-9817-5.
- 18 Bendixsen, D. P., Pollock, T. B., Peri, G., and Hayden, E. J. (2021). Experimental resurrection of
19 ancestral mammalian CPEB3 ribozymes reveals deep functional conservation. *Mol Biol*
20 *Evol.* doi: 10.1093/molbev/msab074.
- 21 Biswas, S., Khimulya, G., Alley, E. C., Esvelt, K. M., and Church, G. M. (2021). Low-N protein
22 engineering with data-efficient deep learning. *Nat Methods* 18, 389–396. doi:
23 10.1038/s41592-021-01100-y.
- 24 Blanco, C., Janzen, E., Pressman, A., Saha, R., and Chen, I. A. (2019). Molecular Fitness
25 Landscapes from High-Coverage Sequence Profiling. *Annu Rev Biophys.* doi:
26 10.1146/annurev-biophys-052118-115333.
- 27 Calonaci, N., Jones, A., Cuturello, F., Sattler, M., and Bussi, G. (2020). Machine learning a
28 model for RNA structure prediction. *NAR Genomics and Bioinformatics* 2, lqaa090. doi:
29 10.1093/nargab/lqaa090.

- 1 Chadalavada, D. M., Gratton, E. A., and Bevilacqua, P. C. (2010). The Human HDV-like CPEB3
2 Ribozyme Is Intrinsically Fast-Reacting. *Biochemistry* 49, 5321–5330. doi:
3 10.1021/bi100434c.
- 4 De Leonardis, E., Lutz, B., Ratz, S., Cocco, S., Monasson, R., Schug, A., et al. (2015). Direct-
5 Coupling Analysis of nucleotide coevolution facilitates RNA secondary and tertiary
6 structure prediction. *Nucleic Acids Res* 43, 10444–10455. doi: 10.1093/nar/gkv932.
- 7 Detlefsen, N. S., Hauberg, S., and Boomsma, W. (2022). Learning meaningful representations of
8 protein sequences. *Nat Commun* 13, 1914. doi: 10.1038/s41467-022-29443-w.
- 9 Dykstra, P. B., Kaplan, M., and Smolke, C. D. (2022). Engineering synthetic RNA devices for
10 cell control. *Nature Reviews Genetics*, 1–14.
- 11 Ferré-D’Amaré, A. R., and Scott, W. G. (2010). Small self-cleaving ribozymes. *Cold Spring*
12 *Harbor perspectives in biology* 2, a003574.
- 13 Ferretti, L., Weinreich, D., Tajima, F., and Achaz, G. (2018). Evolutionary constraints in fitness
14 landscapes. *Heredity* 121, 466–481. doi: 10.1038/s41437-018-0110-1.
- 15 Groher, A.-C., Jager, S., Schneider, C., Groher, F., Hamacher, K., and Suess, B. (2019). Tuning
16 the Performance of Synthetic Riboswitches using Machine Learning. *ACS Synth. Biol.* 8,
17 34–44. doi: 10.1021/acssynbio.8b00207.
- 18 Groher, F., and Suess, B. (2014). Synthetic riboswitches—a tool comes of age. *Biochimica et*
19 *Biophysica Acta (BBA)-Gene Regulatory Mechanisms* 1839, 964–973.
- 20 Jimenez, R. M., Polanco, J. A., and Lupták, A. (2015). Chemistry and Biology of Self-Cleaving
21 Ribozymes. *Trends in Biochemical Sciences* 40, 648–661. doi:
22 10.1016/j.tibs.2015.09.001.
- 23 Kapral, G. J., Jain, S., Noeske, J., Doudna, J. A., Richardson, D. C., and Richardson, J. S. (2014).
24 New tools provide a second look at HDV ribozyme structure, dynamics and cleavage.
25 *Nucleic Acids Research* 42, 12833–12846. doi: 10.1093/nar/gku992.
- 26 Kobori, S., and Yokobayashi, Y. (2016). High-Throughput Mutational Analysis of a Twister
27 Ribozyme. *Angew. Chem. Int. Ed. Engl.* 55, 10354–10357. doi: 10.1002/anie.201605470.
- 28 Lévesque, D., Reymond, C., and Perreault, J.-P. (2012). Characterization of the Trans Watson-
29 Crick GU Base Pair Located in the Catalytic Core of the Antigenomic HDV Ribozyme.
30 *PLoS ONE* 7, e40309. doi: 10.1371/journal.pone.0040309.
- 31 Li, C., Qian, W., Maclean, C. J., and Zhang, J. (2016). The fitness landscape of a tRNA gene.
32 *Science*, aae0568. doi: 10.1126/science.aae0568.
- 33 Peri, G., Gibard, C., Shults, N. H., Crossin, K., and Hayden, E. J. (2022). Dynamic RNA fitness
34 landscapes of a group I ribozyme during changes to the experimental environment.
35 *Molecular Biology and Evolution*, msab373. doi: 10.1093/molbev/msab373.

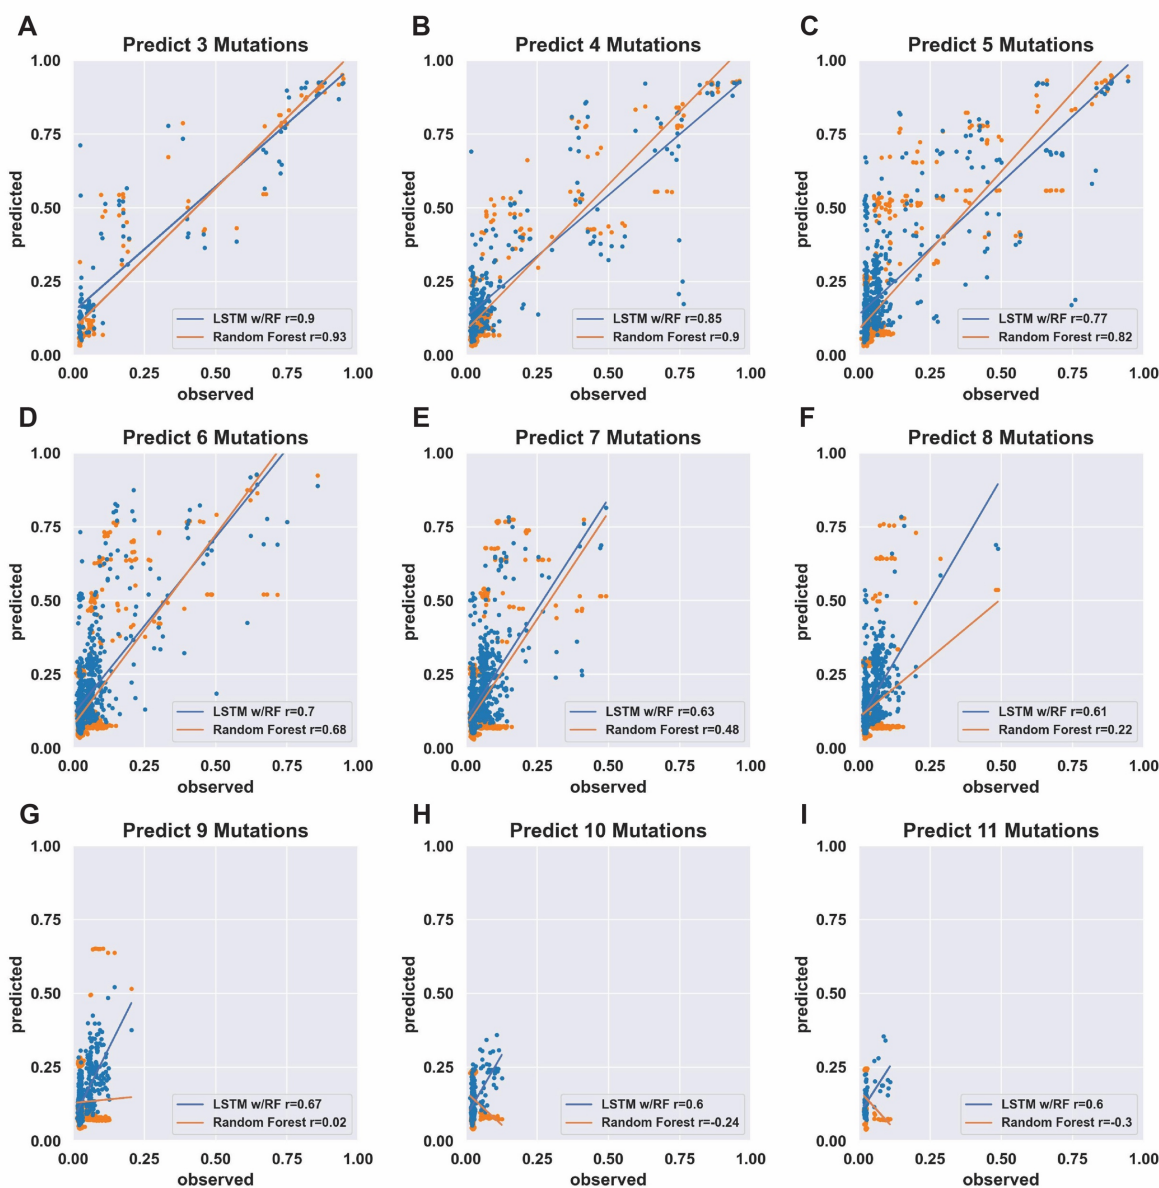
- 1 Pressman, A. D., Liu, Z., Janzen, E., Blanco, C., Müller, U. F., Joyce, G. F., et al. (2019).
2 Mapping a Systematic Ribozyme Fitness Landscape Reveals a Frustrated Evolutionary
3 Network for Self-Aminoacylating RNA. *J. Am. Chem. Soc.* 141, 6213–6223. doi:
4 10.1021/jacs.8b13298.
- 5 Riesselman, A. J., Ingraham, J. B., and Marks, D. S. (2018). Deep generative models of genetic
6 variation capture the effects of mutations. *Nat Methods* 15, 816–822. doi:
7 10.1038/s41592-018-0138-4.
- 8 Salehi-Ashtiani, K., Lupták, A., Litovchick, A., and Szostak, J. W. (2006). A Genomewide
9 Search for Ribozymes Reveals an HDV-Like Sequence in the Human CPEB3 Gene.
10 *Science* 313, 1788–1792. doi: 10.1126/science.1129308.
- 11 Schmidt, C. M., and Smolke, C. D. (2021). A convolutional neural network for the prediction
12 and forward design of ribozyme-based gene-control elements. *eLife* 10, e59697. doi:
13 10.7554/eLife.59697.
- 14 Skilandat, M., Rowinska-Zyrek, M., and Sigel, R. K. O. (2016). Secondary structure
15 confirmation and localization of Mg²⁺ ions in the mammalian CPEB3 ribozyme. *RNA*
16 22, 750–763. doi: 10.1261/rna.053843.115.
- 17 Soskine, M., and Tawfik, D. S. (2010). Mutational effects and the evolution of new protein
18 functions. *Nat. Rev. Genet* 11, 572–582. doi: 10.1038/nrg2808.
- 19 Szendro, I. G., Schenk, M. F., Franke, J., Krug, J., and Visser, J. A. G. M. de (2013).
20 Quantitative analyses of empirical fitness landscapes. *J. Stat. Mech.* 2013, P01005. doi:
21 10.1088/1742-5468/2013/01/P01005.
- 22 Townshend, R. J. L., Eismann, S., Watkins, A. M., Rangan, R., Karelina, M., Das, R., et al.
23 (2021). Geometric deep learning of RNA structure. *Science* 373, 1047–1051. doi:
24 10.1126/science.abe5650.
- 25 Weinreb, C., Riesselman, A., Ingraham, J. B., Gross, T., Sander, C., and Marks, D. S. (2016). 3D
26 RNA and functional interactions from evolutionary couplings. *Cell* 165, 963–975. doi:
27 10.1016/j.cell.2016.03.030.
- 28 Yamagami, R., Kayedkhordeh, M., Mathews, D. H., and Bevilacqua, P. C. (2019). Design of
29 highly active double-pseudoknotted ribozymes: a combined computational and
30 experimental study. *Nucleic Acids Research* 47, 29–42. doi: 10.1093/nar/gky1118.

31

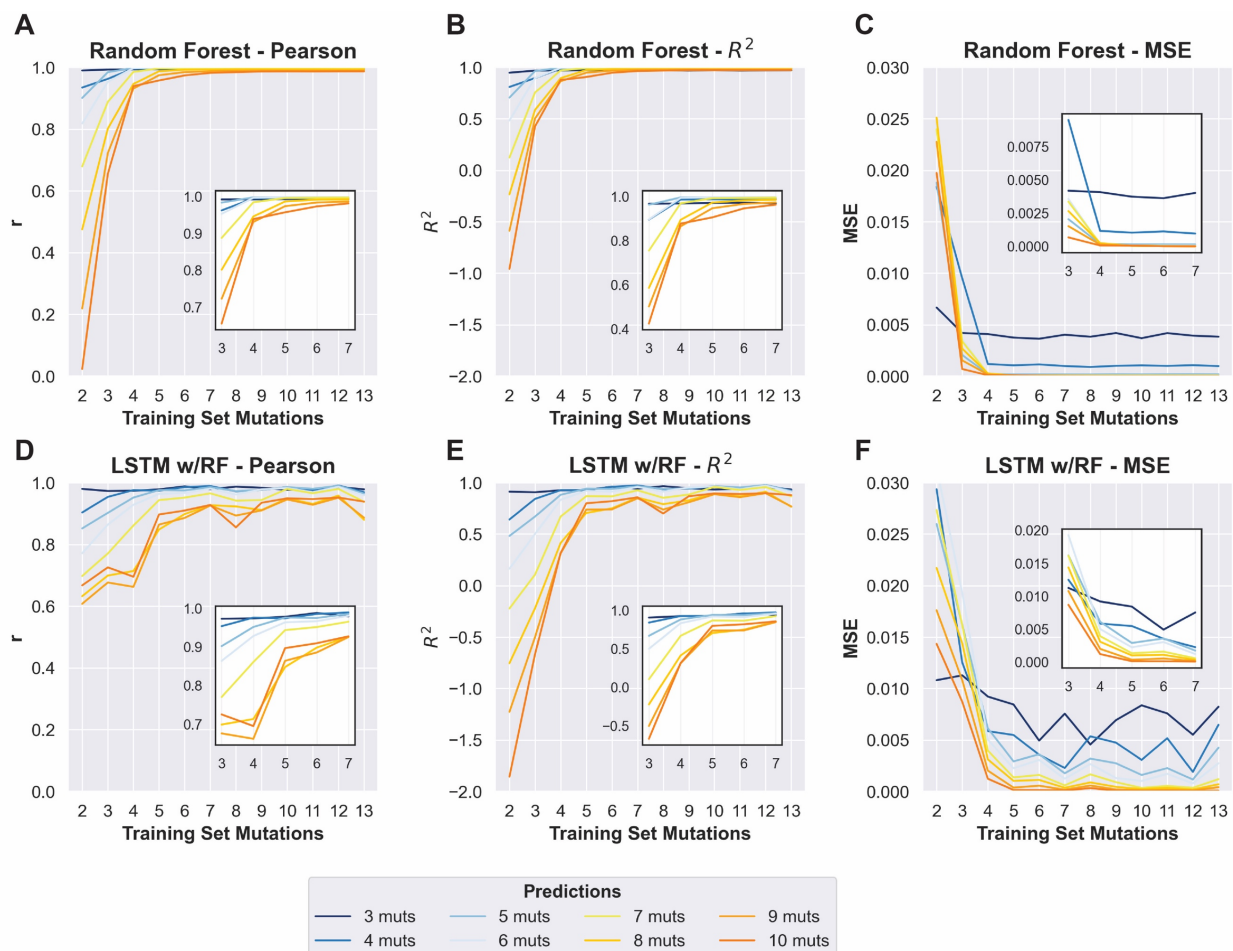
32 **Figure legends**



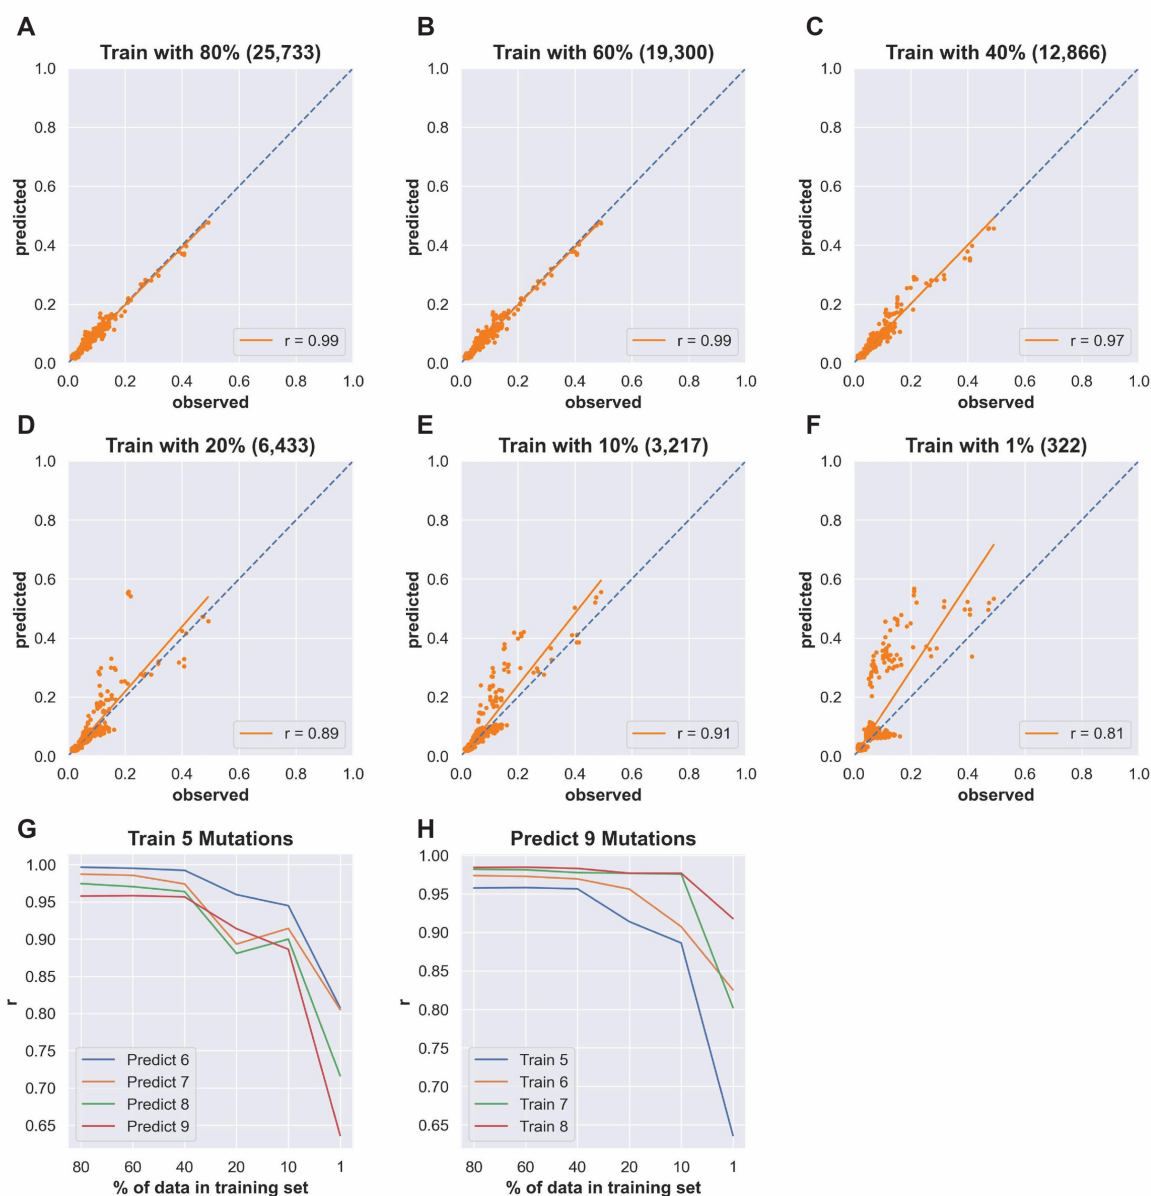
1 Figure 1. The CPEB3 ribozyme and data prediction challenge. (A) Secondary structure diagram
2 of the CPEB3 ribozyme. The white arrow indicates the site of self-cleavage. Nucleotide color
3 indicates the average relative activity of the three possible point mutations at each position.
4 Boxes indicate nucleotide positions mutated in the phylogenetically derived higher-order mutants
5 (B) Heatmap representation of comprehensive single and double-mutant data. Each pixel in the
6 heatmap shows the ribozyme activity for a specific double mutant indicated by the nucleotide
7 positions on the top and right of the heatmap. Insets show base paired regions and specific
8 mutations. Ribozyme activity is determined as the fraction of total reads that map to each
9 sequence that are in the cleaved form (fraction cleaved) relative to the wildtype fraction cleaved.
10 (C) Distribution of pairwise epistasis from double mutant data. Epistasis was calculated as $\varepsilon =$
11 $\log_{10}(W_{AB} * W_{wt} / W_A * W_B)$, where W_{wt} is the fraction cleaved of the wild-type ribozyme, W_A and
12 W_B are fraction cleaved of sequences with individual mutations and W_{AB} is the fraction cleaved
13 of the sequence with both individual mutations. (D) Higher mutational distance variants of the
14 CPEB3 ribozyme represented as a fitness landscape. *Ribozyme activity* (fraction cleaved) is
15 shown for 27,647 sequence variants derived from permutations of naturally occurring mutations.
16 Each node represents a different sequence and the size and color of the node is scaled to the
17 ribozyme activity. Edges connect nodes that differ by a single mutation. Sequences are binned
18 into quintiles of ribozyme activity and the *number of genotypes* reports the number of sequences
19 in each quintile.



1
2 Figure 2. Prediction accuracy of models trained on comprehensive individual and pairs of
3 mutations. (A-I) Scatter plots of *Predicted* (fraction cleaved from the models) and *Observed*
4 (fraction cleaved from experiments). The models were trained on the experimentally determined
5 fraction cleaved for the wild-type and all possible sequences with one mutation (207 sequences)
6 or two mutations (21,114 sequences). Insets report Pearson correlation coefficients r for the
7 model trained by the Random Forest approach (orange) and the LSTM-RF approach (blue). The
8 sequences used to compare prediction vs. observed were separated by the number of mutations
9 relative to the wild-type, as indicated by the title of each graph.



1
2 Figure 3. Improvement in prediction accuracy when including sequences with increased
3 mutational distances in the training data. Changes in Pearson r , R^2 , and mean squared error
4 (MSE) of prediction-observed correlation (y-axis) with increasing numbers of max mutations
5 within the training data (x-axis). Training sets included all sequences up to and including the y-
6 axis value. (A-C) Results obtained for the random forest model. (D-F) Results from the LSTM
7 model. For each plot, colors indicate the numbers of mutations in sequences in the test data (see
8 key). Insets show changes to the same prediction accuracy measurement with the 3-7 mutation
9 training data, to allow more visual resolution.



1
2 Figure 4. Effects of reducing the number of sequences in the training data. (A-F) Scatter plots of
3 *Predicted* (fraction cleaved from the models) and *Observed* (fraction cleaved from experiments)
4 for models trained with decreasing amounts of sequences with 5 or fewer mutations using the
5 random forest approach and predicting the fraction cleaved of sequences with 7 mutations. The
6 percent of the total sequence used in the training data is indicated in the title of each plot, and the
7 number of unique sequences in the training data is reported in parentheses. Pearson correlation
8 coefficients r are indicated as insets. (G) The correlation between predicted and observed for a
9 model trained with decreasing amounts data from sequences with 5 or fewer mutations (“Train

1 5”) and predicting the activity of sequences with increasing numbers of mutations (Predict 6-9).
2 (H) Predicting the activity of sequences with 9 mutations (“Predict 9 Mutations”) with models
3 trained on different reduced data sets.

4

5

6 Table 1. Counts of sequences in training and testing data sets.

No. of mutations	Training	Testing
1	207	---
2	21114	---
3	414	104
4	1240	310
5	2650	662
6	4162	1040
7	4867	1217
8	4241	1060
9	2720	680
10	1249	312
11	389	97
12	74	18
13	6	2

7