

evSeq: Cost-Effective Amplicon Sequencing of Every Variant in a Protein Library

Bruce J. Wittmann^a, Kadina E. Johnston^a, Patrick J. Almhjell^b, and Frances H. Arnold^{a,b,*}

^aDivision of Biology and Biological Engineering,

^bDivision of Chemistry and Chemical Engineering,

California Institute of Technology, MC 210-41

1200 E. California Boulevard, Pasadena, CA 91125 (USA)

*The person to whom correspondence should be addressed

KEYWORDS: *directed evolution, protein engineering, machine learning, next-generation sequencing*

ABSTRACT: Widespread availability of protein sequence-fitness data would revolutionize both our biochemical understanding of proteins and our ability to engineer them. Unfortunately, even though thousands of protein variants are generated and evaluated for fitness during a typical protein engineering campaign, most are never sequenced, leaving a wealth of potential sequence-fitness information untapped. Primarily, this is because sequencing is unnecessary for many protein engineering strategies; the added cost and effort of sequencing is thus unjustified. It also results from the fact that, even though many lower cost sequencing strategies have been developed, they often require at least some sequencing or computational resources, both of which can be barriers to access. Here, we present every variant sequencing (evSeq), a method and collection of tools/standardized components for sequencing a variable region within every variant gene produced during a protein engineering campaign at a cost of cents per variant. evSeq was designed to democratize low-cost sequencing for protein engineers and, indeed, anyone interested in engineering biological systems. Execution of its wet-lab component is simple, requires no sequencing experience to perform, relies only on resources and services typically available to biology labs, and slots neatly into existing protein engineering workflows. Analysis of evSeq data is likewise made simple by its accompanying software (found at github.com/fhalab/evSeq, documentation at fhalab.github.io/evSeq), which can be run on a personal laptop and was designed to be accessible to users with no computational experience. Low-cost and easy to use, evSeq makes collection of extensive protein variant sequence-fitness data practical.

INTRODUCTION

Engineered proteins are valuable tools across the biological and chemical sciences and have revolutionized industries ranging from food to fuels, pharmaceuticals, and textiles by providing green and efficient protein solutions to challenging chemical problems.¹ Over the course of a protein engineering campaign, hundreds to thousands or more protein variants will be constructed and have their fitnesses (level of, e.g., thermostability, catalytic activity, substrate binding, etc.) evaluated. Notably, sequence information is typically not gathered alongside the functional information, even though it could provide useful biochemical insight.²⁻⁴ This is largely because many engineering strategies can be applied without sequencing. For example, during a typical directed evolution (DE) experiment, often only the best-performing variant or variants are sequenced in each round of mutagenesis and screening; sequencing every variant is viewed as an unnecessary expense. Given the massive amount of functional data gathered during a typical DE campaign, however, if sequencing *were* performed for the variants generated during these experiments, the resultant large datasets of sequence-fitness information could be revolutionary for biological,

biochemical, and biocatalytic research. This is especially true for data-driven protein engineering strategies such as machine learning (ML), the development of which has benefitted tremendously from large sequence-fitness datasets made available by strategies like deep mutational scanning (DMS) and in databases like ProtaBank.⁵⁻¹⁶

Unfortunately, the standard sequencing strategy employed during DE—Sanger sequencing—is too expensive for sequencing all variants tested during a round of evolution.¹⁷ Sanger sequencing is ubiquitous due to ease of sample preparation and ready availability of sequencing providers. However, the cost of Sanger sequencing scales linearly with the number of samples (Supplemental Figure S1). Thus, while the cost of sequencing just the top variants in a round of DE is minor, sequencing the hundreds or thousands of variants generated over the full engineering endeavor is not. Ideally, any new approach to sequencing during a protein engineering campaign would be comparable in cost, effort, and accessibility to that of sequencing just the top variants by Sanger sequencing. Here we present a collection of standardized and accessible protocols, components, and software that accomplishes this goal. This collection, which we call every variant sequencing

(evSeq), democratizes barcode sequencing strategies and expands on services made available by multiplexed next-generation sequencing (NGS) providers to allow amplicon sequencing of a region of interest within every variant produced during a round of DE at a cost of cents per variant.^{18,19} Sample preparation for evSeq is simple, and the method requires no experience with NGS to perform, relies only on resources and services typically available to biology labs, and slots neatly into existing protein engineering experimental workflows. The accompanying software for analysis of evSeq data (found at github.com/fhalab/evSeq, documentation at fhalab.github.io/evSeq) was designed to be accessible to users with no computational experience and can be run on a personal laptop.

In this paper, we detail the underlying strategies, protocol, and potential applications of evSeq. We begin by describing the strategies employed by evSeq to extend multiplexed NGS for sequencing protein variant libraries in a way that reduces both cost and effort. We then describe the wet-lab protocol of evSeq sample preparation, focusing on how it can be completed without disrupting an existing protein engineering workflow. Next, we discuss the features of the evSeq software before finally presenting two case studies that highlight potential applications of evSeq. In particular, we highlight how (1) the sequence-fitness data from evSeq can provide valuable information about the quality of variant libraries and the functional screen as well as how mutations modulate protein activity, and how (2) the data generated from evSeq can be used to implement ML for protein engineering. We designed evSeq for use as a routine procedure in many protein/enzyme assays (especially DE and protein engineering experiments leveraging mutagenesis strategies that target specific sites or a segment of the sequence). This tool brings cost-effective, easy-to-use sequencing to all protein engineers, regardless of experience with NGS and access to sequencing and computational resources. We believe that widespread adoption of evSeq—and the resultant datasets generated—will be invaluable for future ML-guided protein engineering and will help us better understand protein sequence-fitness relationships.

RESULTS AND DISCUSSION

evSeq uses inline barcoding to expand on commercially available multiplexed next-generation sequencing. Unlike Sanger sequencing, which outputs a single chromatogram that represents the population of DNA in a sequenced sample, NGS outputs millions of individual DNA reads that represent a random draw from the population of DNA in the sequenced sample.¹⁸ Confidence in NGS sequencing results is largely determined by the sequencing “coverage”, which for the purposes of this paper is defined as the number of returned reads that map to a specific nucleotide on a reference sequence. Higher coverage enables more confident identification of mutations relative to a reference sequence as the increased redundancy allows distinguishing between true sequence mutations and errors that arise during library preparation, clustering, or sequencing.

A single NGS run is roughly three orders of magnitude more expensive than a Sanger sequencing run, but because the run outputs millions of reads, this cost can be spread over multiple samples using a technique known as “multiplexed NGS” (Supplemental Figure S1). In multiplexed NGS, each submitted sample is tagged with a “molecular barcode”—a unique piece of DNA that encodes the sample’s original identity—before all

samples are sequenced together in the same NGS run.^{19–25} Post sequencing, the barcodes are used to assign individual reads to individual samples. For instance, barcodes can be used to distinguish reads coming from samples belonging to specific plates and wells.²⁶ Importantly, multiplexed NGS can be outsourced just like Sanger sequencing (making it accessible to all laboratories regardless of sequencing experience), and sequencing providers typically charge tens of dollars per sample in a multiplexed sequencing run, yielding on the order of 10^4 – 10^5 individual sequences per sample (assuming the run is performed on an Illumina MiSeq instrument).

The level of coverage granted by a set number of reads depends on the length of the DNA sample being sequenced, the length of the NGS read used to sequence it, and whether those reads are paired-end. NGS reads are short (300 bp or less on Illumina systems), and so reads must be spread across a longer sample to sequence it in full. The expected coverage (average coverage per nucleotide) obtained for a DNA sample thus depends both on its length and the read length used for sequencing. For example, with the $\sim 10^5$ reads returned by a commercial MiSeq multiplexed sequencing run, a 3 Mb genome could be sequenced with 150 bp paired-end reads to an expected coverage of $\sim 10\times$, whereas a 20 kb plasmid could be sequenced to an expected coverage of $\sim 1500\times$.

Because shorter samples can be sequenced at higher coverage for a given number of reads, it can be advantageous to sequence only the region of interest of a sample. This is exemplified by amplicon sequencing, a strategy in which a researcher sequences a PCR product (an amplicon) that targets a specific region of interest in the DNA.²⁷ For instance, continuing the example from above, with $\sim 10^5$ total 150 bp paired-end reads, a 300 bp PCR product could be sequenced to an expected coverage of $\sim 100,000\times$.

Many mutagenesis methods employed in protein engineering (e.g., site-saturation²⁸ and tile-based mutagenesis²⁹ strategies) target mutations to a specific position or region in the sequence of a protein, and thus the variants produced can be sequenced with amplicon sequencing to high coverage.²⁰ Notably, however, even though increasing coverage yields more confident results, it comes with diminishing returns, and it is generally held that coverage in the tens is more than sufficient for effective reference-based identification of mutations (Supplemental Figure S1).³⁰ Indeed, clinical sequencing of human genomes targets 30x coverage or greater to minimize false base calls. Given this reference, it is clear that the $\sim 100,000\times$ coverage that would be returned from a multiplexed sequencing run for a 300 bp amplicon is far higher than necessary for effective identification of mutations—2,000 amplicons could be sequenced in the same run and still yield clinical-grade coverage.

evSeq achieves cost-effectiveness by relying on the facts that (1) at tens of dollars per sample, the cost of sending a single sample to an outsourced multiplexed NGS run is comparable to the total cost of Sanger sequencing the top variants in a round of DE, (2) amplicon sequencing can be used to identify mutations in protein variants from many protein engineering library types, and (3) enough reads are returned for a single sample in a commercial multiplexed NGS run to sequence hundreds of amplicons. Specifically, the evSeq protocol (Figure 1, *Supplemental: evSeq Library Preparation/Data Analysis Protocol*) works by focusing all reads of a single multiplexed NGS sample to specific regions on hundreds of protein variants, achieving sequencing depths of 10^1 – 10^3 at the approximate cost

and level of accessibility of using Sanger sequencing of just the top variants in a round of DE (Supplemental Figure S1).

The evSeq library preparation protocol begins with PCR amplification of the region of interest in each variant (i.e., the position/region where mutations were made) and appending inline DNA barcodes to the resultant amplicons that encode their original plate-well position (Figure 1A).^{26,31,32} This is a one-pot, two-step, plate-based PCR procedure that uses two sets of primer pairs. Each primer in the first set of primers (“inner” primers) consists of a user-specified 3’ “seed” region that binds to the regions flanking the region of interest as well as a 5’ predefined universal adapter (*Supplemental: Inner Primer Design*). Each primer in the second set of primers (“outer” primers) consists of (1) a 3’ region that matches the adapter of the inner primers, (2) a central 7-nucleotide barcode where each barcode pair between forward and reverse outer primers is unique to a plate-well position, and (3) a 5’ sequence matching the Illumina Nextera transposase adapters (*Supplemental: Outer Primer Design, Supplemental: Barcode Design, Supplemental Tables S1–S2*). We designed 96 unique forward and 96 unique reverse outer primers for evSeq which, because both forward and reverse outer primers contain a barcode, can be combined to encode up to $96^2 = 9,216$ possible plate-well positions (*Supplemental: Preparation of evSeq Barcode Primer Mixes, Supplemental Tables S3–S10*). Note that we also provide a pre-filled IDT order form for the outer primers on the GitHub associated with this work—see *Supplemental: Ordering Barcode Primers from IDT* for details. While we recommend using these pre-tested barcodes, users can also design their own to, e.g., further expand the number of available combinations.). Importantly, this set of outer primers can be used to sequence any target region from any gene with evSeq, and so only needs to be ordered once, constituting a one-time initial setup cost in the range of a few hundred dollars (the exact cost will vary based on oligo provider and any institutional agreements set up with said provider). Once outer primers are ordered, only a new inner primer pair is needed for each new region of interest to be targeted by evSeq.

Once all barcoded amplicons have been produced, they are pooled and sent to a sequencing provider, who will then use the transposase adapters installed with the outer primers as a handle to perform a third and final PCR to barcode the *pool* of amplicons *once again* with a pair of sample-specific Illumina indices (Figure 1B). At this point each amplicon in the pool has one pair of sample-specific Illumina barcodes and one pair of plate-well-specific inline evSeq barcodes. This complete evSeq library is sequenced as a single sample in a multiplexed NGS run along with samples from other users (whether or not they are also evSeq samples). Post sequencing, the sequencing provider uses the sample-specific barcodes to identify those sequences belonging to the evSeq pool and returns them to the user (i.e., the provider “demultiplexes” the run, separating evSeq sequences from those of other users). The user then uses the evSeq software to analyze the returned sequences, assigning them to corresponding plate-well positions using the evSeq barcodes and identifying the mutations in the variants relative to a reference (Figure 1B, 1C).

evSeq library preparation fits into existing protein engineering and sequencing workflows and was designed to be resource efficient. A typical procedure for evaluating protein variants involves (1) arraying colonies of an organism (e.g., *Escherichia coli*) that harbor a plasmid encoding a protein variant into the wells of a (usually 96-well) microplate, (2) growing the resulting cultures to stationary phase (colloquially,

an “overnight culture”), (3) using the overnight culture to inoculate a fresh culture that will be used to express the protein variants, and (4) evaluating the fitnesses of expressed protein variants. The expression stage (step 3) typically involves downtime where the experimentalist must wait until the culture reaches sufficient density before inducing protein expression and then again as expression takes place. evSeq library preparation can be performed easily in either of these time windows. The evSeq library preparation protocol begins with the barcoding PCR described at the end of the previous section; this one-pot, two-step, plate-based PCR was designed to be compatible with outsourced sequencing workflows, minimize preparation time, and minimize laboratory resource usage (*Supplemental: evSeq Library Preparation/Data Analysis Protocol*). For instance, use of inline barcodes is a known, effective strategy for expanding the number of samples that can be multiplexed without having to modify the Illumina indices used during multiplexed sequencing.^{31,32} Because evSeq library preparation uses inline barcodes, it grants the outsourced sequencing provider maximal flexibility in choice of Illumina indices. In other words, evSeq library preparation is decoupled from preparation of the Illumina library that will eventually be sequenced, allowing the evSeq library to be run just as any other sample would be that is submitted to a sequencing provider.

As mentioned in the previous section, use of a two-step PCR reduces the number of primers that must be ordered per new sequencing region of interest. Because evSeq relies on 96 unique forward barcodes and 96 unique reverse barcodes, a single-primer PCR would require ordering 192 new barcoding primers for each new target region evaluated in each library. In a two-primer protocol, however, the inclusion of a universal adapter on the inner primers allows the same 192 outer primers to be used regardless of target position in the variant—only two unique primers (forward and reverse inner) must be purchased for each new target region, and only if existing inner primers from previously targeted regions are not already compatible. Additionally, the evSeq PCR directly uses liquid from the overnight culture as a source of template DNA (Figure 1B, *Supplemental: evSeq Library Preparation/Data Analysis Protocol*); the template DNA is released from lysed cells during the initial heating step of the PCR, avoiding a costly and time-intensive DNA isolation/purification step and allowing researchers to use materials already prepared as part of the protein expression workflow.³²

The remaining steps of evSeq library preparation were, like the PCR stage, also designed to be resource and time efficient. After completion of the PCR, the resulting barcoded amplicons are pooled by plate and purified via gel extraction. Pooling prior to purification goes against standard practice for multiplexed NGS library preparation, which is to purify samples individually, quantify their DNA concentration, then combine them in equimolar quantities to ensure more equal read distribution across samples after sequencing.³³ However, because individual plates in protein engineering libraries tend to contain variants from the same region of the same protein scaffold (e.g., as would be typical for variants from a comprehensive site-saturation library), it is assumed that the variation in PCR reaction yield will be minor within plates and that, as a result, the same plate can be pooled prior to quantification with only minor effects on read distribution. Using this “pooling first” strategy, only as many purifications as there are *plates* must be performed as opposed to as many as there are *variants*, thus enabling faster processing of evSeq amplicons while reducing resource usage. As will be shown in

later sections, the distribution of reads returned using pooling first is perfectly acceptable for confidently identifying variant sequences.

Once all pooled plates have been purified, the concentrations of the individual purified pools are measured. The pools are then normalized by molarity and combined into a final evSeq library, which is in turn submitted as a single sample to a sequencing provider. As described in the previous section, the provider will perform a final PCR on the evSeq library to add

sample-specific barcodes before sequencing it as a single sample in a multiplexed sequencing run. Outsourcing the sequencing stage has two main benefits: First, it allows evSeq to be performed by research groups with no prior sequencing experience and no direct access to sequencing equipment—groups need only be familiar with PCR, a ubiquitous technology in protein engineering laboratories. Second, to be cost effective, multiplexed sequencing should be run with tens of samples at least (Supplemental Figure S1). By outsourcing




Figure 1. Overview of evSeq library preparation and processing. (A) In the first stage of the PCR, a region of interest is amplified with primers that include a 3' site-specific region (gray) with 5' adapter sequences (dark blue). The second PCR stage adds molecular barcodes (rainbow) with primers that bind to the adapter regions and add adapters for downstream NGS processing (light blue). (B) To avoid costly DNA isolation steps, evSeq uses liquid cultures of cells harboring mutated DNA (e.g., an “overnight culture” of *E. coli*) as template during the one-pot two-step barcoding PCR described in A. Each plate is pooled individually and gel purified. Purified pools are then adjusted for concentration differences and pooled together before being sent to a sequencing provider, who then appends another set of barcodes as well as sequence elements necessary for Illumina NGS sequencing. This sample is now pooled with those of other users and a multiplexed sequencing run is performed. After sequencing, the sequencing provider uses the barcodes that they attached to separate (“demultiplex”) the evSeq reads from reads of other users; the provider returns evSeq reads in .fastq files. (C) The .fastq files returned by the NGS provider are inputs to the evSeq software, which uses the evSeq forward/reverse barcode pair to map each read to a specific plate and well based on known barcode combinations. The software also processes the mapped reads (see the Supplemental and evSeq documentation for more details) to, among other things, assign variant identities to each well and return interactive HTML visualizations.

the sequencing stage, groups that do not frequently produce evSeq libraries need not wait until enough libraries have accumulated to run sequencing—a single outsourced submission, for instance, can be run along with those of other research groups with a variety of different sequencing needs.

The final stage of the evSeq workflow is data analysis using the evSeq software (github.com/fhalab/evSeq) (Figure 1C). Extensive documentation of the software and its capabilities is available as a website (fhalab.github.io/evSeq). The software was designed to be accessible to users with varying degrees of computational experience and can be run through either a graphical user interface (GUI), a command line application, or in a Python environment (e.g., a Jupyter notebook). Outputs from the software range from high-level overviews of data (e.g., an interactive “Platemap” graphic that displays sequencing coverage and identified mutations in each well of each plate; see Figure 1C for an example) to low-level details about the population of reads assigned to each well (e.g., in a well identified as polyclonal, the percentage of reads mapping to each of the identified variants). Functional data can also be easily associated with identified variants using the evSeq software outputs to produce sequence-fitness datasets, and we provide Jupyter notebooks and web pages that walk users through the process.

evSeq facilitates library construction, validation, and sequence-fitness pairing. To highlight the utility of evSeq for engineering and biochemical experiments, we first examined how it could be used to construct high-confidence and informative sequence-fitness data. Specifically, we constructed and screened eight single-site-saturation libraries of the enzyme Tm9D8*—an engineered β -subunit of tryptophan synthase from *Thermotoga maritima* (TmTrpB)—for tryptophan-forming activity at 30 °C (Figure 2).³⁴ In two of the screened libraries, we targeted two positions distant from the active site (A118 and S292) that have been seen to play a role in allosteric regulation of TmTrpB enzymes; in the other six libraries, we targeted active-site residues known to modulate the activity of TrpB (E105, L162, I166, F184, S228, and Y301) (Figure 2A).^{35–37} As we show below, this type of sequence-fitness data can be used to assess the quality of a protein engineering library, identify improved variants during a round of directed evolution, and give insight into the significance of a given residue in catalysis.

Many factors can introduce bias into a site-saturation mutagenesis experiment, such as annealing bias for the native nucleotides during the PCR for library construction or contamination with the template plasmid during transformation. Without sequencing all of the variants, it is impossible to know that the library is representative of the experimental design. Since evSeq reports exactly which variants are contained in a library, researchers can leverage this to implement important quality control practices as part of the standard protein screening workflow. For instance, of all 153 possible unique variants in our eight single-site-saturation libraries, we observed 149 of them (Figures 2B and C); only I166A, S292C, S292D, and S292H could not be assigned with confidence, where we define >80% abundance in a well with >10 reads as our confidence threshold. Of the variants identified, many were found in replicate (Figure 2D) due to oversampling during colony picking, which ensures that all protein variants have a chance to be found and screened (All libraries were constructed with the 22-codon trick³⁸ and 88 individual colonies were screened for each library, so we expected a 98% probability of seeing all variants assuming

perfect construction of libraries). Conveniently, this oversampling also allows us to evaluate the noise in our functional screen (Figure 2E) which further improves the confidence in the quality of data gathered.

Given just the fitness data gathered in this experiment, a protein engineer would identify 50 wells that are at least 1.2-fold improved over the parent enzyme Tm9D8*. However, with the sequence-fitness pairs constructed via evSeq, we know that these 50 wells correspond to only 16 unique variants. Depending on how conservative the engineer was as to what should be sequenced, a decision to sequence hits with Sanger sequencing could result in anywhere from 12 (2-fold improvement) to 50 (1.2-fold improvement) wells sent off for sequencing for a total cost of \$36 to \$150 (using an estimate of \$3 per sequence). It would cost ~\$2000 to sequence all eight plates via Sanger. Using evSeq, however, we obtained the sequences of *all* 625 wells of variants for only \$100, corresponding to \$0.13 per non-control well. In other words, using evSeq, we can produce far more sequence-fitness information than sequencing just the top hits using Sanger all for a similar cost. Importantly, although the evSeq defaults currently allow only eight plates to be sequenced at once, the number of variants included in this experiment could likely have been increased as the median number of reads per well was 86 (mean: 98), which is above what is needed for reliable sequencing. Assuming that doubling the number of plates would halve the number of reads seen for each well, doubling the number of plates sequenced would cause only 14 non-control well sequences to drop below the confidence threshold.

The per-variant cost of evSeq may be reduced even further using different services and sequencing platforms. For instance, in both this section and the next, the reported number of reads and ~\$100 total cost are from outsourced MiSeq runs, which returned hundreds of thousands of total reads per evSeq library. We report these numbers because outsourced multiplexed MiSeq is a standard service available to all research groups. As an alternative to outsourcing, however, our institution provides multiplexed sequencing (via the Caltech Millard and Muriel Jacobs Genetics and Genomics Laboratory) on an Illumina NextSeq platform, returning an average of ~10x more reads than the outsourced MiSeq run for a total cost of ~\$10. At 10x more reads and 10x less the total cost, the per-variant evSeq cost could decrease 100-fold to <\$0.01. Indeed, we were able to re-sequence the TrpB libraries at a per-variant cost of ~\$0.01 with ~2.2 million total reads returned for an average of thousands of reads per variant, far higher than what is needed for reliable variant calling. It must be noted, however, that analysis of the millions of evSeq reads was no longer practical on a personal laptop, requiring a desktop workstation instead. Computational power beyond a laptop will be needed when processing more than hundreds of thousands of reads with the existing evSeq software.

Of final note, aside from providing valuable information for protein engineering experiments, evSeq can also facilitate investigation into the biochemical relevance of specific positions/mutations. Specifically, because all possible variants in a site-saturation library can be identified by evSeq, the sequence-fitness information generated can be used to explore the effects of mutations more fully than, for instance, an alanine scanning experiment.³⁹ Using an example from the TrpB data gathered here, an alanine scanning experiment would tell a biochemist that the mutation to the conserved catalytic residue E105A inactivates the enzyme, with no information about the effects of other amino acid changes at this position. Using site-




Figure 2. evSeq enables low-cost investigation of library quality and sequence-fitness pairing in site-saturation mutagenesis libraries. (A) Eight residues (red) known to modulate the activity of Tm9D8* were independently targeted with site-saturation mutagenesis: A118 and S292 (distal residues), E105, L162, I166, F184, S228, and Y301 (active-site residues). An active form of the pyridoxal 5'-phosphate cofactor is represented in green, and the substrate indole is shown in light blue. (B) Library quality can be investigated by plotting a heatmap of the number of times each variant/mutant was identified at each targeted position ("# in Library") from processed evSeq data. Parent amino acids are each marked with an asterisk. (C) Likewise, the effect of mutations and mutational "hotspots" can be identified by plotting a heatmap of the average activity measurements for each variant/mutation in each library, normalized to the average parent activity for that library ("Normalized Rate"), when fitness data is combined with evSeq data. (D) An example plot made possible by evSeq visualization functions shows the number of times each amino acid was found in a single TrpB library (position 105), also accounting for known controls and unidentified wells. (E) Another example output of the evSeq software shows activity for a single library (position 105), showing biological replicates. The inset displays the role of the mutated residue in this library, which is to coordinate the nitrogen of the indole substrate. Note that the circles in this plot correspond to individual measurements while the bar plot represents the mean of these measurements. If no circles are present for a bar (e.g., E105D), then this is because only a single instance of this mutation was observed. Circles are not shown in this case to allow distinguishing between a single replicate and a tight distribution of multiple replicates.

saturation with evSeq, we instead find that all mutations to E105 except for E105D inactivate the enzyme. The fact that glutamate and aspartate are the only amino acids containing a carboxylic acid suggests that this functional group is critical for activity (Figure 2E, with inset).

evSeq can be used to generate data for machine learning-assisted protein engineering. We next wanted to demonstrate the utility of evSeq for advancing and applying machine learning-assisted protein engineering (MLPE). In MLPE, models are trained to learn a function that relates protein sequence to protein fitness (i.e., they learn $f(\text{sequence}) = \text{fitness}$).^{5,6,9-11} These models are then used for rapid, low-cost *in silico* prediction of protein fitness, avoiding or greatly reducing the need for often-costly laboratory screening of variants (Figure 3).

Sequence-fitness data is critical for effective MLPE. Indeed, even though strategies exist that *can* predict protein fitness from sequence alone (e.g., those that use evolutionary data to predict protein fitness), their effectiveness is improved with the inclusion of sequence-fitness information.^{7,14,15,40} As a result, the most effective MLPE workflows require that both sequence

and fitness data be collected, unlike a DE workflow, which requires only fitness data.

The need to collect sequence data in addition to fitness data is an often-overlooked additional cost of MLPE strategies compared to standard DE. For instance, we recently developed an ML strategy known as machine learning-assisted directed evolution (MLDE) for efficient navigation of epistatic combinatorial protein variant libraries.^{41,42} Previously, we used MLDE to evolve *Rhodothermus marinus* nitric oxide dioxygenase (RmaNOD) for greater enantioselectivity in a carbon-silicon bond-forming reaction.⁴¹ Over the course of the engineering campaign, we collected six 96-well plates of sequence-fitness data for training ML models. In total, sequencing the variants in these plates by Sanger sequencing cost ~\$1700. High additional sequencing costs like these can make MLPE methods far less attractive, even if they are more effective than traditional DE at finding high-fitness protein variants.⁴² However, given that evSeq enables sequencing all variants for a cost similar to standard DE methods, it enables use of MLPE without added cost. In essence, evSeq eliminates the sequencing burden of MLPE.




Figure 3. evSeq eliminates the sequencing burden of MLPE. Traditional DE only collects sequence information for top variants, essentially “throwing away” fitness data from inferior variants and learning nothing about the underlying fitness landscape. If, instead, evSeq is used to collect sequence information for all variants, MLPE methods, which require sequence-fitness pairs for supervised model training, can be implemented. Sampling from a fitness landscape, an ML model can be trained to predict the fitnesses of missing sequences and reconstruct the missing regions of this landscape.

To demonstrate the application of evSeq to MLPE, we used it to sequence five plates of *RmaNOD* variants from a four-site combinatorial library. Coupled with fitness data, the sequences resulting from this run could be used to drive a round of MLDE. Notably, sequencing these plates by Sanger sequencing would have cost ~\$1400; in contrast, sequencing by evSeq using an outsourced multiplexed MiSeq run cost ~\$100 for a per-variant cost of ~\$0.21. The median read depth per variant in this run was 463 (mean: 506), much higher than is required for accurate sequencing, and so more plates—from either the same or a different library—could have reasonably been added to this evSeq run to decrease the per-variant sequencing cost even further (Figure 3B). Of course, as discussed in the previous section, in-house sequencing could have cut sequencing costs an additional ten-fold.

The cost of sequencing is most notably a barrier for MLPE strategies that focus on developing models for a single protein with a well-defined fitness (e.g., MLDE); however, the applicability of evSeq to MLPE is not limited solely to cost-reduction. For instance, ML strategies have been developed that, rather than focusing on a specific protein, train models on sequence-fitness information across multiple different protein scaffolds.^{16,43} The goal is for these models to learn global

determinants of protein fitness, then to use the models as general-purpose protein fitness predictors. By enabling the collection of sequence-fitness pairs across a wider array of proteins and fitness definitions, evSeq opens these approaches to new and more diverse data sources. Generally speaking, the more sequence-fitness data available to train and benchmark these strategies, the better we expect them to perform and the more rapidly we expect improvements to be developed.¹⁶ It is no coincidence that large leaps forward in other ML disciplines have followed increased availability of large, diverse datasets, with the rapid advance in computer vision sparked by ImageNet being perhaps the most prominent example.⁴⁴ Widespread adoption of evSeq—and commitment to depositing sequence-fitness data in resources such as ProtaBank—would provide such a dataset for protein engineering.⁸ This dataset would span the range of all engineered proteins and all target fitnesses, capture examples of sequences with both higher and lower/zero fitness relative to a parent (the latter of which is effectively never recorded with current DE sequencing practices), and overall enable rapid advancement in MLPE.

evSeq detects all variability in the sequenced amplicons.

Although we focused here on demonstrating applications involving targeted mutagenesis strategies, evSeq is also applicable to other mutagenesis methods as the associated software can identify both user-specified and unspecified positions of variability (Figure 4A). This feature not only informs the user of potential unexpected mutations in the sequenced amplicon (Supplemental Table S11), but also allows it to work effectively with tile-based mutagenesis strategies and other semi-targeted mutagenesis strategies (e.g., error-prone PCR of specific regions or small genes). All that is required is that the amplicon length and read length be able to capture the full region containing mutations.

It should be noted that evSeq will not detect off-target mutations outside of the constructed amplicon as these regions are not sequenced, meaning that it is unable to identify other mutations in a larger DNA element that may be contributing to activity. Due to this fact, for exceedingly unexpected mutational effects that are not seen in replicate, we suggest sequencing the rest of the DNA element to confirm the presence or absence of any off-target mutations. However, this limitation is mitigated by the fact that off-target mutations are rare and,




Figure 4. evSeq detects variability and can be expanded for random mutagenesis. (A) evSeq does not require that the user specify which position in the amplicon was targeted. Instead, the software can identify variable regions by comparing to a reference (B) evSeq can be used to sequence entire genes by designing a set of inner primer pairs which together capture the entire gene. Different evSeq barcodes can then be used for each region, and the user can reconstruct the entire sequence.

importantly, evSeq is agnostic to read length and will work with any length of paired-end sequencing.⁴⁵

While the current software version is not yet suited for other, long-read sequencing technologies (e.g., PacBio or Oxford Nanopore), future versions could be updated and validated with these data formats and make full gene-length evSeq experiments more straightforward and cost effective. Given this, evSeq is currently best suited and most cost effective when all expected mutations exist in the sequenced amplicon, though sequencing of multiple overlapping amplicons can readily allow evSeq to be expanded to sequence entire genes of variants arrayed in microplates (Figure 4B). Care must be taken in such an application, however, to account for the fact that aggressive mutation rates could compromise the annealing efficiency of inner primers binding in the variable region, as could mutations to positions closer to the binding region of the 3' end of the inner primer. Such situations would lead to a higher proportion of wells failing sequencing.

CONCLUSION

Hundreds to thousands of protein variants (or more) are constructed and their fitnesses evaluated over the course of a standard protein engineering campaign. Without sequencing, these fitnesses are next to useless—the time, effort, and resources expended to produce them are largely wasted. Comparable in cost to existing protocols, accessible to scientists with no or minimal sequencing and computational experience, and easy to implement with existing technology, evSeq rescues these fitness data by making the collection of sequence data for every variant a practical and highly useful step of the protein engineering pipeline. Given the number of research groups working on DE and other protein engineering projects, widespread adoption of evSeq would lead to an explosion in the availability of sequence-fitness information. By sequencing every variant, no laboratory screening effort is wasted, and we open the door to advances in both our biochemical understanding of proteins and our ability to engineer them with data-driven methods.

MATERIALS AND METHODS

Single-Site-Saturation Library Generation for TrpB.

Saturation mutagenesis libraries were prepared using a modification of the “22-codon trick” described by Kille et al.³⁸ We first designed primers using the templates given in Supplemental Table S12. For the forward primers, each sequence of “NNN” in these templates was replaced with “NDT”, “VHG”, and “TGG”, resulting in a total of three degenerate primers which could then be mixed at a ratio of 12:9:1, respectively. The reverse primers were used without changes.

We also designed primers that bind within the ampicillin resistance (AmpR) gene in pET22b(+) with sequences as given in Supplemental Table S13. These primers were designed such that, when used in combination with the site-specific primers to run a PCR, two medium-length fragments would be created with a break in the AmpR gene. For the forward site-saturation primers, a PCR was performed using the reverse AmpR primer, resulting in a fragment from ~1500–2000 bp long. For the reverse site-saturation primers, a PCR was performed using the forward AmpR primer, resulting in a fragment ~4500–5000 bp long.

Once PCRs finished, 1 μ L of DpnI (NEB R0176S) was added to each of the reactions, which were then incubated at 37 °C for 1 h to digest the unmutated template plasmid. The presence of correctly sized fragments was confirmed via gel electrophoresis and each fragment was then excised from the gel and purified with the Zymoclean Gel DNA Recovery Kit (Zymo Research D4002).

Purified fragments were then assembled following the standard Gibson assembly method.⁴⁶ After 1 h at 50 °C, the reaction mixtures were desalted with a DNA Clean & Concentrator-5 kit (Zymo Research D4013) and used to transform electrocompetent E. cloni® cells (Lucigen 60051-1). Libraries were spread onto solid agar selection medium consisting of Luria Broth (RPI L24040-5000.0) supplemented with 100 μ g/mL carbenicillin (LB_{carb}) and incubated at 37 °C until single colonies were observed. Individual colonies were then transferred into the wells of 96-well 2-mL deep-well plates containing 300 μ L of LB_{carb} to isolate monoclonal enzyme variants, with 8 wells being reserved for control conditions, giving 4-fold oversampling of the 22-codon library. These cultures were grown overnight at 37 °C, 220 rpm, and 80% humidity in an Infors Multitron HT until they reached stationary phase, at which point 100 μ L from each well were mixed with an equal volume of 50% glycerol and stored at –80 °C for future use.

For protein expression, 20 μ L of the remaining culture were used to inoculate 630 μ L of Terrific Broth with 100 μ g/mL carbenicillin (TB_{carb}). These were then grown at 37 °C, 220 rpm, and 80% humidity for 3 hours in an Infors Multitron HT, at which point they were placed on ice for 30 minutes. Following this, 50 μ L of a 14 mM solution of isopropyl- β -D-thiogalactoside (IPTG; GoldBio #I2481C100) in TB_{carb} were added to each well to induce protein expression at a final concentration of 1 mM IPTG. Expression proceeded in the same Infors Multitron HT shaker as before at 22 °C, 220 rpm for roughly 18 hours. Cells were harvested via centrifugation at 4500g for 10 minutes, the supernatant was removed, and the plates (now containing pelleted, expressed cells) were placed at –20 °C until needed.

Once cells had been harvested, cultures for evSeq were prepared. These cultures were started from the 96-well plate glycerol stocks prepared prior to moving into the cell expression protocol; the cultures were grown overnight (~18hrs) in an Infors Multitron HT (220 rpm, 37 °C) to saturation in 96-well deep-well plates in 300 μ L of LB_{carb}. These cultures were then frozen and stored at –20 °C to be used for sequencing with evSeq.

A GenBank file detailing the plasmid and primers used in this section is available on the evSeq GitHub at https://github.com/fhalab/evSeq/tree/master/genbank_files/tm9d8s.gb.

Sequencing TrpB Libraries with evSeq. Frozen overnight cultures (preparation detailed in the previous section) were thawed at room temperature. Libraries were then sequenced with the process described in *Supplemental: evSeq Library Preparation/Data Analysis Protocol*; the evSeq software was run using all default parameters (average_q_cutoff = 25, bp_q_cutoff = 30, length_cutoff = 0.9, match_score = 1, mismatch_penalty = 0, gap_open_penalty = 3, gap_extension_penalty = 1, variable_thresh = 0.2, variable_count = 10) with the “return_alignments” flag thrown. The inner primers used for library preparation are in Supplemental Table S14. The barcode plates (Supplemental Tables S3 – S10) were paired to positions as given in Supplemental Table S15.

Measuring the Rate of Tryptophan Formation. Rate of tryptophan formation data was collected with the same procedure described in Rix et al. for non-heat-treated lysate preparation in the section “Indole rate measurements” with a few modifications: lysis occurred in 300 μ L KPi buffer with 100 μ M pyridoxal 5'-phosphate (PLP) supplemented with 1 mg/mL lysozyme, 0.02 mg/mL bovine pancreas DNase I, and 0.1x BugBuster; lysis occurred at 37 °C for 1 h.³⁵

Four-Site-Saturation Library Generation for *RmaNOD*. Positions S28, M31, Q52, and L56 of a variant of *RmaNOD* (*RmaNOD* Y32G) were targeted for comprehensive site-saturation mutagenesis using a variant of the 22-codon trick originally described by Kille et al.³⁸ Due to the proximity of positions S28 and M31, it was easiest to use the same mutagenesis primers to target them; the same was done for positions Q52 and L56. Because the 22-codon trick requires three degenerate codons per position targeted, nine individual primers capturing all combinations (3 codons ² 2 positions/per primer = 9 primers) of the degenerate codons had to be ordered for each of the two mutagenic primers. Sequences of these primers are given in Supplemental Table S16.

The primers from Supplemental Table S16 were all ordered from IDT at 100 μ M. Both a “forward” and a “reverse” primer mixture were prepared by combining individual forward and reverse primers in proportion to the number of individual codons they encoded. A 10 μ M forward-reverse primer mixture was then prepared by adding 10 μ L of both the forward and reverse primer mixtures to 80 μ L ddH₂O. Once the forward-reverse primer mixture was prepared, it was used in a PCR to build a pool of DNA fragments containing the four-site combinatorial libraries. Two fragments that captured the remainder of the *RmaNOD* gene and host plasmid (pET22b(+)) were also produced by PCR. The primers used for these flanking fragments are given in Supplemental Table S17.

After PCR completed 1 μ L DpnI (NEB R0176S) was added to each reaction. The reactions were then held at 37 °C in a thermocycler for 1 h. The PCR fragments were then gel-extracted using a Zymoclean Gel DNA Recovery Kit (D4002).

Fragments were to eventually be assembled using Gibson assembly.⁴⁶ Because the efficiency of Gibson assembly increases with decreasing numbers of fragments, an assembly PCR was performed to combine flanking fragment 1 (see Supplemental Table S17 for details) and the variant fragment. The resultant assembled fragment was then gel-extracted, again using a Zymoclean Gel DNA Recovery Kit (D4002).

To complete construction of the library of variant plasmids, a Gibson assembly was performed to combine the assembled PCR fragment and flanking fragment 0. After Gibson assembly, the Gibson reaction was cleaned using a Monarch PCR & DNA Cleanup Kit (NEB CAT T1030L). The cleaned Gibson product was next used to transform electrocompetent *E. coli* BL21 DE3. Transformed cells were spread onto solid agar selection medium consisting of Luria Broth (RPI L24040-5000.0) supplemented with 100 μ g/mL ampicillin (LB_{amp}) and incubated at 37 °C until single colonies were observed.

To build the 96-well plates of *RmaNOD* variants used to demonstrate evSeq, 400 μ L LB + 100 μ g/mL ampicillin were first added to each well of 5x 96-well deepwell plates. Colonies from the agar plates grown overnight were then picked into the wells of the deepwell plates. The plates were placed in an Infors Multitron HT at 240 rpm, 37 °C for ~16 h. To glycerol stock the now-stationary-phase culture, 100 μ L overnight culture were added to 100 μ L 50% glycerol before being stored at -80 °C until its use in evSeq library preparation.

A GenBank file detailing the plasmid and primers used in this section is available on the evSeq GitHub (https://github.com/fhalab/evSeq/tree/master/genbank_files/rmanod_v32g.gb).

Sequencing *RmaNOD* Libraries with evSeq. To begin preparation of culture for evSeq with the *RmaNOD* variants, cultures in 96-well deep-well plates (with 300 μ L of LB_{carb}) were started from the 96-well plate glycerol stocks prepared in the previous section. The plates were placed in an Infors Multitron HT at 240 rpm; the cultures were grown overnight (~18hrs) before being frozen and stored at -20 °C.

To start the evSeq protocol, frozen overnight cultures were thawed in a room temperature water bath. Libraries were then sequenced with the process described in *Supplemental: evSeq Library Preparation/Data Analysis Protocol*; the evSeq software was run using the same parameters as for the TrpB data analysis (see *Sequencing TrpB Libraries with evSeq*, above). The inner primers used for evSeq library preparation are given in Supplemental Table S18. The barcode plates (Supplemental Tables S3 – S10) were paired to positions as given in Supplemental Table S19.

AUTHOR INFORMATION

Corresponding Author

Frances H. Arnold – Division of Biology and Biological Engineering, Division of Chemistry and Chemical Engineering, California Institute of Technology, MC 210-41, 1200 E. California Boulevard, Pasadena, CA 91125, USA;
Email: frances@cheme.caltech.edu

Authors

Bruce J. Wittmann – Division of Biology and Biological Engineering, California Institute of Technology, MC 210-41, 1200 E. California Boulevard, Pasadena, CA 91125, USA;

Kadina E. Johnston – Division of Biology and Biological Engineering, California Institute of Technology, MC 210-41, 1200 E. California Boulevard, Pasadena, CA 91125, USA;

Patrick J. Almhjell – Division of Chemistry and Chemical Engineering, California Institute of Technology, MC 210-41, 1200 E. California Boulevard, Pasadena, CA 91125, USA;

Author Contributions

Author contributions are provided using the CRediT taxonomy:

BJW: Conceptualization, methodology, software, validation, investigation, writing – original draft, writing – review and editing, funding acquisition. **KEJ:** Methodology, data-collection, software, investigation, writing – original draft, writing – review and editing, visualization, funding acquisition. **PJA:** Methodology, data-collection, software, validation, investigation, writing – original draft, writing – review and editing, visualization. **FHA:** Resources, writing – original draft, writing – review and editing, funding acquisition.

Conflict of Interest Statement

Conflicts of Interest: None

Funding Sources

This work was supported by an Amgen Chem-Bio-Engineering Award (CBEA). This work was supported by the NSF Division of Chemical, Bioengineering, Environmental and Transport Systems (CBET 1937902). This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences, under Award Number DE-SC0022218. This

report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

ASSOCIATED CONTENT

Supporting Information: Figures showing cost comparisons between sequencing methods and additional information on TrpB and RmaNOD sequencing libraries, additional details on the design of evSeq primers and barcodes, detailed protocols for implementation of evSeq, and tables giving DNA sequences used in this work.

Data Availability: All raw and processed data generated by this study can be found at CaltechData (DOI: 10.22002/D1.2140). The software version used to analyze all data in this study is tagged as v1.0.0 on the associated GitHub repository.

ACKNOWLEDGMENTS

The authors thank Shan Li, Adrienne Rollie, and Eric Brustad at Illumina, Inc: Shan Li and Adrienne Rollie for helping us troubleshoot the evSeq method and Eric Brustad for critical reading of the manuscript. We also thank fellow Arnold laboratory members Nathaniel Goldberg and Nicholas Porter for implementing evSeq (which pointed us to necessary improvements), Anders Knight for suggesting and prototyping evSeq software features, Ella Watkins-Dulaney for assistance in building the TrpB libraries, and Sabine Brinkmann-Chen for critical reading of the manuscript.

ABBREVIATIONS

evSeq, every variant sequencing; DE, directed evolution; ML, machine learning; DMS, deep mutational scanning; NGS, next-generation sequencing; GUI, graphic user interface; *Tm*TrpB, *Thermotoga maritima* β -subunit of tryptophan synthase; MLPE, machine learning-assisted protein engineering; MLDE, machine learning-assisted directed evolution; RmaNOD, *Rhodothermus marinus* nitric oxide dioxygenase.

REFERENCES

- (1) BCC Research Staff (2018) Global Markets for Enzymes in Industrial Applications. BCC Research LLC. <https://www.bccresearch.com/market-research/biotechnology/global-markets-for-enzymes-in-industrial-applications.html>.
- (2) Podgornaia, A. I., and Laub, M. T. (2015) Pervasive degeneracy and epistasis in a protein-protein interface. *Science*. 347, 673–677. DOI: [10.1126/science.1257360](https://doi.org/10.1126/science.1257360).
- (3) Wu, N. C., Dai, L., Olson, C. A., Lloyd-Smith, J. O., and Sun, R. (2016) Adaptation in protein fitness landscapes is facilitated by indirect paths. *Elife* 5. DOI: [10.7554/eLife.16965](https://doi.org/10.7554/eLife.16965).
- (4) Faure, A. J., Domingo, J., Schmiedel, J. M., Hidalgo-Carcedo, C., Diss, G., and Lehner, B. (2021) Global mapping of the energetic and allosteric landscapes of protein binding domains. *bioRxiv*. DOI: [10.1101/2021.09.14.460249](https://doi.org/10.1101/2021.09.14.460249).
- (5) Yang, K. K., Wu, Z., and Arnold, F. H. (2019) Machine-learning-guided

directed evolution for protein engineering. *Nat. Methods* 16, 687–694. DOI: [10.1038/s41592-019-0496-6](https://doi.org/10.1038/s41592-019-0496-6).

- (6) Li, G., Dong, Y., and Reetz, M. T. (2019) Can Machine Learning Revolutionize Directed Evolution of Selective Enzymes? *Adv. Synth. Catal.* 361, 2377–2386. DOI: [10.1002/adsc.201900149](https://doi.org/10.1002/adsc.201900149).
- (7) Hopf, T. A., Ingraham, J. B., Poelwijk, F. J., Schärfe, C. P. I., Springer, M., Sander, C., and Marks, D. S. (2017) Mutation effects predicted from sequence co-variation. *Nat. Biotechnol.* 35, 128–135. DOI: [10.1038/nbt.3769](https://doi.org/10.1038/nbt.3769).
- (8) Wang, C. Y., Chang, P. M., Ary, M. L., Allen, B. D., Chica, R. A., Mayo, S. L., and Olafson, B. D. (2018) ProtaBank: A repository for protein design and engineering data. *Protein Sci.* 27, 1113–1124. DOI: [10.1002/pro.3406](https://doi.org/10.1002/pro.3406).
- (9) Mazurenko, S., Prokop, Z., (2020) Machine Learning in Enzyme Engineering. *ACS Catal.* 10, 1210–1223. DOI: [10.1021/acscatal.9b04321](https://doi.org/10.1021/acscatal.9b04321).
- (10) Siedhoff, N. E., Schwaneberg, U., and Davari, M. D. (2020) Machine learning-assisted enzyme engineering, in *Methods in Enzymology* 1st ed., pp 281–315. Elsevier Inc. DOI: [10.1016/bs.mie.2020.05.005](https://doi.org/10.1016/bs.mie.2020.05.005).
- (11) Wittmann, B. J., Johnston, K. E., Wu, Z., and Arnold, F. H. (2021) Advances in machine learning for directed evolution. *Curr. Opin. Struct. Biol.* 69, 11–18. DOI: [10.1016/j.sbi.2021.01.008](https://doi.org/10.1016/j.sbi.2021.01.008).
- (12) Fowler, D. M., and Fields, S. (2014) Deep mutational scanning: a new style of protein science. *Nat. Methods* 11, 801–807. DOI: [10.1038/nmeth.3027](https://doi.org/10.1038/nmeth.3027).
- (13) Livesey, B. J., and Marsh, J. A. (2020) Using deep mutational scanning to benchmark variant effect predictors and identify disease mutations. *Mol. Syst. Biol.* 16. DOI: [10.15252/msb.20199380](https://doi.org/10.15252/msb.20199380).
- (14) Riesselman, A. J., Ingraham, J. B., and Marks, D. S. (2018) Deep generative models of genetic variation capture the effects of mutations. *Nat. Methods* 15, 816–822. DOI: [10.1038/s41592-018-0138-4](https://doi.org/10.1038/s41592-018-0138-4).
- (15) Meier, J., Rao, R., Verkuil, R., Liu, J., Sercu, T., and Rives, A. (2021) Language models enable zero-shot prediction of the effects of mutations on protein function. *bioRxiv*. DOI: [10.1101/2021.07.09.450648](https://doi.org/10.1101/2021.07.09.450648).
- (16) Gray, V. E., Hause, R. J., Luebeck, J., Shendure, J., and Fowler, D. M. (2018) Quantitative Missense Variant Effect Prediction Using Large-Scale Mutagenesis Data. *Cell Syst.* 6, 116–124. DOI: [10.1016/j.cels.2017.11.003](https://doi.org/10.1016/j.cels.2017.11.003).
- (17) Sanger, F., and Coulson, A. R. (1975) A Rapid Method for Determining Sequences in DNA by Primed Synthesis with DNA Polymerase. *J. Mol. Biol.* 94, 441–448. DOI: [10.1016/0022-2836\(75\)90213-2](https://doi.org/10.1016/0022-2836(75)90213-2).
- (18) Metzker, M. L. (2010) Sequencing technologies — the next generation. *Nat. Rev. Genet.* 11, 31–46. DOI: [10.1038/nrg2626](https://doi.org/10.1038/nrg2626).
- (19) Smith, A. M., Heisler, L. E., StOnge, R. P., Farias-Hesson, E., Wallace, I. M., Bodeau, J., Harris, A. N., Perry, K. M., Giaever, G., Pourmand, N., and Nislow, C. (2010) Highly-multiplexed barcode sequencing: an efficient method for parallel analysis of pooled samples. *Nucleic Acids Res.* 38. DOI: [10.1093/nar/gkq368](https://doi.org/10.1093/nar/gkq368).
- (20) Appel, M. J., Longwell, S. A., Morri, M., Neff, N., Herschlag, D., and Fordyce, P. M. (2021) uPIC-M: efficient and scalable preparation of clonal single mutant libraries for high-throughput protein biochemistry. *ACS Omega* 6, 30542–30554. DOI: [10.1021/acsomega.1c04180](https://doi.org/10.1021/acsomega.1c04180).
- (21) Srivathsan, A., Lee, L., Katoh, K., Hartop, E., Kutty, S. N., Wong, J., Yeo, D., and Meier, R. (2021) ONTbarcoder and MinION barcodes aid biodiversity discovery and identification by everyone, for everyone. *BMC Biol.* 19. DOI: [10.1186/s12915-021-01141-x](https://doi.org/10.1186/s12915-021-01141-x).
- (22) Glenn, T. C., Nilsen, R. A., Kieran, T. J., Sanders, J. G., Bayona-Vázquez, N. J., Finger, J. W., Pierson, T. W., Bentley, K. E., Hoffberg, S. L., Louha, S., Garcia-De Leon, F. J., del Rio Portilla, M. A., Reed, K. D., Anderson, J. L., Meece, J. K., Aggrey, S. E., Rekaya, R., Alabady, M., Belanger, M., Winker, K., and Faircloth, B. C. (2019) Adapterama I:

- universal stubs and primers for 384 unique dual-indexed or 147,456 combinatorially-indexed Illumina libraries (iTru & iNext). *PeerJ* 7, e7755. DOI: [10.7717/peerj.7755](https://doi.org/10.7717/peerj.7755).
- (23) Wierbowski, S. D., Vo, T. V., Falter-Braun, P., Jobe, T. O., Kruse, L. H., Wei, X., Liang, J., Meyer, M. J., Akturk, N., Rivera-Erick, C. A., Cordero, N. A., Paramo, M. I., Shayhidin, E. E., Bertolotti, M., Tippens, N. D., Akther, K., Sharma, R., Katayose, Y., Salehi-Ashtiani, K., Hao, T., Ronald, P. C., Ecker, J. R., Schweitzer, P. A., Kikuchi, S., Mizuno, H., Hill, D. E., Vidal, M., Moghe, G. D., McCouch, S. R., and Yu, H. (2020) A massively parallel barcoded sequencing pipeline enables generation of the first ORFeome and interactome map for rice. *Proc. Natl. Acad. Sci.* 117, 11836–11842. DOI: [10.1073/pnas.1918068117](https://doi.org/10.1073/pnas.1918068117).
- (24) Chubiz, L. M., Lee, M.-C., Delaney, N. F., and Marx, C. J. (2012) FREQ-Seq: A Rapid, Cost-Effective, Sequencing-Based Method to Determine Allele Frequencies Directly from Mixed Populations. *PLoS One* 7, e47959. DOI: [10.1371/journal.pone.0047959](https://doi.org/10.1371/journal.pone.0047959).
- (25) Weile, J., Sun, S., Cote, A. G., Knapp, J., Verby, M., Mellor, J. C., Wu, Y., Pons, C., Wong, C., Lieshout, N., Yang, F., Tasan, M., Tan, G., Yang, S., Fowler, D. M., Nussbaum, R., Bloom, J. D., Vidal, M., Hill, D. E., Aloy, P., and Roth, F. P. (2017) A framework for exhaustively mapping functional missense variants. *Mol. Syst. Biol.* 13, 957. DOI: [10.1525/msb.20177908](https://doi.org/10.1525/msb.20177908).
- (26) Campbell, N. R., Harmon, S. A., and Narum, S. R. (2015) Genotyping-in-Thousands by sequencing (GT-seq): A cost effective SNP genotyping method based on custom amplicon sequencing. *Mol. Ecol. Resour.* 15, 855–867. DOI: [10.1111/1755-0998.12357](https://doi.org/10.1111/1755-0998.12357).
- (27) Wen, C., Wu, L., Qin, Y., Van Nostrand, J. D., Ning, D., Sun, B., Xue, K., Liu, F., Deng, Y., Liang, Y., and Zhou, J. (2017) Evaluation of the reproducibility of amplicon sequencing with Illumina MiSeq platform. *PLoS One* 12, e0176716. DOI: [10.1371/journal.pone.0176716](https://doi.org/10.1371/journal.pone.0176716).
- (28) Siloto, R. M. P., and Weselake, R. J. (2012) Site saturation mutagenesis: Methods and applications in protein engineering. *Biocatal. Agric. Biotechnol.* 1, 181–189. DOI: [10.1016/j.bcab.2012.03.010](https://doi.org/10.1016/j.bcab.2012.03.010).
- (29) Melnikov, A., Rogov, P., Wang, L., Gnirke, A., and Mikkelsen, T. S. (2014) Comprehensive mutational scanning of a kinase in vivo reveals substrate-dependent fitness landscapes. *Nucleic Acids Res.* 42, e112. DOI: [10.1093/nar/gku511](https://doi.org/10.1093/nar/gku511).
- (30) Sims, D., Sudbery, I., Iltott, N. E., Heger, A., and Ponting, C. P. (2014) Sequencing depth and coverage: Key considerations in genomic analyses. *Nat. Rev. Genet.* 15, 121–132. DOI: [10.1038/nrg3642](https://doi.org/10.1038/nrg3642).
- (31) de Muinck, E. J., Trosvik, P., Gilfillan, G. D., Hov, J. R., and Sundaram, A. Y. M. (2017) A novel ultra high-throughput 16S rRNA gene amplicon sequencing library preparation method for the Illumina HiSeq platform. *Microbiome* 5, 68. DOI: [10.1186/s40168-017-0279-1](https://doi.org/10.1186/s40168-017-0279-1).
- (32) Tresnak, D. T., and Hackel, B. J. (2020) Mining and Statistical Modeling of Natural and Variant Class IIa Bacteriocins Elucidate Activity and Selectivity Profiles across Species. *Appl. Environ. Microbiol.* 86, e01646-20. DOI: [10.1128/AEM.01646-20](https://doi.org/10.1128/AEM.01646-20).
- (33) Illumina. (2019) Nextera XT DNA Library Prep Reference Guide. https://emea.support.illumina.com/content/dam/illumina-support/documents/documentation/chemistry_documentation/samplepreps_nextera/nextera-xt/nextera-xt-library-prep-reference-guide-15031942-05.pdf.
- (34) Boville, C. E., Romney, D. K., Almhjell, P. J., Sieben, M., and Arnold, F. H. (2018) Improved Synthesis of 4-Cyanotryptophan and Other Tryptophan Analogues in Aqueous Solvent Using Variants of TrpB from *Thermotoga maritima*. *J. Org. Chem.* 83, 7447–7452. DOI: [10.1021/acs.joc.8b00517](https://doi.org/10.1021/acs.joc.8b00517).
- (35) Rix, G., Watkins-Dulaney, E. J., Almhjell, P. J., Boville, C. E., Arnold, F. H., and Liu, C. C. (2020) Scalable continuous evolution for the generation of diverse enzyme variants encompassing promiscuous activities. *Nat. Commun.* 11, 5644. DOI: [10.1038/s41467-020-19539-6](https://doi.org/10.1038/s41467-020-19539-6).
- (36) Romney, D. K., Murciano-Calles, J., Wehrmüller, J. E., and Arnold, F. H. (2017) Unlocking Reactivity of TrpB: A General Biocatalytic Platform for Synthesis of Tryptophan Analogues. *J. Am. Chem. Soc.* 139, 10769–10776. DOI: [10.1021/jacs.7b05007](https://doi.org/10.1021/jacs.7b05007).
- (37) Buller, A. R., Brinkmann-Chen, S., Romney, D. K., Herger, M., Murciano-Calles, J., and Arnold, F. H. (2015) Directed evolution of the tryptophan synthase β -subunit for stand-alone function recapitulates allosteric activation. *Proc. Natl. Acad. Sci.* 112, 14599–14604. DOI: [10.1073/pnas.1516401112](https://doi.org/10.1073/pnas.1516401112).
- (38) Kille, S., Acevedo-Rocha, C. G., Parra, L. P., Zhang, Z. G., Opperman, D. J., Reetz, M. T., and Acevedo, J. P. (2013) Reducing codon redundancy and screening effort of combinatorial protein libraries created by saturation mutagenesis. *ACS Synth. Biol.* 2, 83–92. DOI: [10.1021/sb300037w](https://doi.org/10.1021/sb300037w).
- (39) Morrison, K. L., and Weiss, G. A. (2001) Combinatorial alanine-scanning. *Curr. Opin. Chem. Biol.* 5, 302–307. DOI: [10.1016/s1367-5931\(00\)00206-4](https://doi.org/10.1016/s1367-5931(00)00206-4).
- (40) Hsu, C., Nisonoff, H., Fannjiang, C., and Listgarten, J. (2022) Learning protein fitness models from evolutionary and assay-labeled data. *Nat. Biotechnol.* DOI: [10.1038/s41587-021-01146-5](https://doi.org/10.1038/s41587-021-01146-5).
- (41) Wu, Z., Kan, S. B. J., Lewis, R. D., Wittmann, B. J., and Arnold, F. H. (2019) Machine learning-assisted directed protein evolution with combinatorial libraries. *Proc. Natl. Acad. Sci.* 116, 8852–8858. DOI: [10.1073/pnas.1901979116](https://doi.org/10.1073/pnas.1901979116).
- (42) Wittmann, B. J., Yue, Y., and Arnold, F. H. (2021) Informed training set design enables efficient machine learning-assisted directed protein evolution. *Cell Syst.* 12, 1026–1045.e7. DOI: [10.1016/j.cels.2021.07.008](https://doi.org/10.1016/j.cels.2021.07.008).
- (43) Alieva, A., Aceves, A., Song, J., Mayo, S., Yue, Y., and Chen, Y. (2021) Learning to Make Decisions via Submodular Regularization. *ICLR*.
- (44) Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A. C., and Fei-Fei, L. (2015) ImageNet Large Scale Visual Recognition Challenge. *Int. J. Comput. Vis.* 115, 211–252. DOI: [10.1007/s11263-015-0816-y](https://doi.org/10.1007/s11263-015-0816-y).
- (45) McInerney, P., Adams, P., and Hadi, M. Z. (2014) Error Rate Comparison during Polymerase Chain Reaction by DNA Polymerase. *Mol. Biol. Int.* 2014. DOI: [10.1155/2014/287430](https://doi.org/10.1155/2014/287430).
- (46) Gibson, D. G., Young, L., Chuang, R. Y., Venter, J. C., Hutchison, C. A., and Smith, H. O. (2009) Enzymatic assembly of DNA molecules up to several hundred kilobases. *Nat. Methods* 6, 343–345. DOI: [10.1038/nmeth.1318](https://doi.org/10.1038/nmeth.1318).