Estimation of substitution and indel rates via k-mer statistics

Mahmudur Rahman Hera
School of Electrical Engineering and Computer Science, Pennsylvania State University, USA

Paul Medvedev
Department of Computer Science and Engineering, The Pennsylvania State University, University Park, USA
Department of Biochemistry and Molecular Biology, The Pennsylvania State University, University Park, USA
Huck Institutes of the Life Sciences, The Pennsylvania State University, University Park, USA

David Koslicki
School of Electrical Engineering and Computer Science, Pennsylvania State University, USA
Department of Biology, Pennsylvania State University, USA
Huck Institutes of the Life Sciences, Pennsylvania State University, USA

Antonio Blanca
School of Electrical Engineering and Computer Science, Pennsylvania State University, USA

Abstract

Methods utilizing k-mers are widely used in bioinformatics, yet our understanding of their statistical properties under realistic mutation models remains incomplete. Previously, substitution-only mutation models have been considered to derive precise expectations and variances for mutated k-mers and intervals of mutated and non-mutated sequences. In this work, we consider a mutation model that incorporates insertions and deletions in addition to single-nucleotide substitutions. Within this framework, we derive closed-form k-mer-based estimators for the three fundamental mutation parameters: substitution, deletion rate, and insertion rates. We provide theoretical guarantees in the form of concentration inequalities, ensuring accuracy of our estimators under reasonable model assumptions. Empirical evaluations on simulated evolution of genomic sequences confirm our theoretical findings, demonstrating that accounting for insertions and deletions signals allows for accurate estimation of mutation rates and improves upon the results obtained by considering a substitution-only model. An implementation of estimating the mutation parameters from a pair of fasta files is available here: github.com/KoslickiLab/estimate_rates_using_mutation_model.git. The results presented in this manuscript can be reproduced using the code available here: github.com/KoslickiLab/est_rates_experiments.git.

2012 ACM Subject Classification Applied computing → Computational biology; Theory of computation → Theory and algorithms for application domains; Mathematics of computing → Probabilistic inference problems

Keywords and phrases k-mers, mutation rate, indel, alignment-free, estimation, substitution, insertion, deletion

Supplementary Material The full version of this manuscript, including all proofs, is available on the BioRxiv: `https://doi.org/10.1101/2025.05.14.653858`

Funding This material is based upon work supported by the National Science Foundation under Grant No. DBI2138585. Research reported in this publication was supported by the National Institute Of General Medical Sciences of the National Institutes of Health under Award Number R01GM146462. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

1 Introduction

Estimating the mutation rate between two evolutionarily related sequences is a classical question in molecular evolution, with roots that pre-date the genomics era [24]. Early quantitative efforts focused on amino-acid substitution: the seminal PAM matrices of Dayhoff et al. converted curated alignments of close homologues into an evolutionary time-scale [5], while the BLOSUM series by Henikoff and Henikoff mined ungapped blocks of conserved proteins to improve sensitivity for more diverged sequences [8]. These approaches and later profile-based HMM models [6] were derived from pairwise or multiple alignments and remain the gold standard when accurate alignments are available.

Over the last decade, however, high-throughput sequencing has shifted the scale of comparative genomics from dozens to millions of genomes, rendering high computational cost pipelines (e.g. quadratic-time) increasingly impractical. Consequently, alignment-free techniques that summarize sequences by inexpensive statistics have become indispensable [20, 25]. These approaches most commonly utilize k-mer sets and sketches thereof. Popular tools such as Mash [13], Skmer [16] and more recent sketch-corrected frameworks like Sylph [19] and FracMinHash-based methods [9, 10, 14, 18] can build whole-genome phylogenies, screen metagenomic samples, and compute millions of pairwise point-mutation rate estimates in minutes rather than days.

Despite their empirical success, theoretical understanding of alignment-free estimators has lagged behind practice. Nearly all existing models treat evolution as a pure-substitution process, ignoring insertions and deletions (indels), or else their performance in the presence of indels is often not thoroughly evaluated [15]. When indels are frequent, substitution-only estimators systematically inflate divergence and can misplace taxa—even in otherwise well-resolved trees of primates constructed from k-mer Jaccard similarities [13]. Recent work has quantified how k-mer-based statistics are affected by substitutions and are also used to estimate substitution-only mutation rates [10], yet a principled treatment that jointly infers substitution and indel parameters from k-mer statistics is still absent. This omission is particularly significant because indels represent a substantial fraction of genomic variation and play crucial roles in evolution [22]. Such indels cause substitution-only k-mer-based methods to underperform, as just like with substitutions, disruption of k-mer content by indel events affects at least k k-mers, often leading to overestimates of mutation rates [3, 10].

In this paper we introduce the first closed-form, alignment-free estimators for the three fundamental mutation parameters: substitution rate p_s, deletion rate p_d, and mean insertion length d under a model that explicitly incorporates single-nucleotide substitutions, deletions, and geometrically-distributed insertions. Starting from elementary counts of unmutated and single-deletion k-mers, we derive algebraic expressions for p_s, p_d, and d and prove a sub-exponential concentration bound that guarantees a strong form of consistency as the sequence length grows as detailed in our main contribution: Theorem 9. Simulations on synthetic and real bacterial genomes demonstrate that modeling indels yields markedly more accurate distance estimates than substitution-only approaches. The remainder of this paper is organized as follows: Section 2 defines our mutation model, Section 3 derives basic statistical properties, Section 4 presents our estimators, Section 5 provides theoretical guarantees about these estimators, Section 6 details our implementation, Section 7 presents experimental results, and Section 8 concludes with a discussion of implications and future directions.

2 The mutation model

We first describe the mutation model under consideration. We use the model from [17]. This model is a type of indel channel and various variations of it have been used to model sequence evolution (e.g. [7]). Let S be a string over the alphabet $\{A, C, G, T\}$, and let L be the number of characters in S. Let S_i denote the i-th character in S where $1 \leq i \leq L$. We define four operations on a string S as

follows:

- Substitute(i): select a character $c \in \{A, C, G, T\} \setminus \{S_i\}$ uniformly at random and replace S_i by c.
- Delete(i): remove the character S_i.
- Stay(i): do nothing.
- Insert(i, s): insert string s between S_{i-1} and S_i if $i > 0$. If $i = 0$, prepend s to the start of S instead.

The mutation process takes as input a string S and three parameters p_s, p_d, and d, where $0 \leq p_s, p_d < 1$, $p_s + p_d < 1$, and $d \geq 0$. Then,

1. For each i, let a_i be the operation to be performed at position i. Then, $a_i =$ Sub with probability p_s, $a_i =$ Del with probability p_d, or $a_i =$ Stay with the remaining probability $1 - p_s - p_d$. (These operations are not performed at this point, but only recorded.)
2. Let *track* be a function mapping from a position in the original string S to its position in the modified S. Initially, $track(i) = i$ for all i. We assume that *track* is updated accordingly whenever an insertion or deletion operation is performed.
3. For each i that is a substitute action, apply Substitute(i).
4. For each i, let $l_i \geq 0$ be a sample from a geometric distribution with mean d. Then generate a random string Q_i of length exactly l_i by drawing each character in Q_i from $\{A, C, G, T\}$ independently and uniformly at random. This is equivalent to sampling a uniformly random string among all strings of length l_i with characters in $\{A, C, G, T\}$. If $l_i = 0$, Q_i is the empty string.
5. For every i, execute Insert($track(i)$, Q_i).
6. For every i such that $a_i =$ Del, execute Delete($track(i)$).
7. Return the resulting S'.

3 Preliminary statistics of the mutation process

In this section, we define several quantities whose statistics we use to design an estimator of the parameters of the mutation process. For the remainder of the manuscript, we assume that the mutation process described in Section 2 is applied to string S of length L, with the unknown parameters p_s, p_d, and d, and a mutated string S' is returned. The theoretical results we present are centered around the concept of k-spans. We define $\mathcal{K}_i$, the k-span at position $1 \leq i \leq L - k + 1$ as the range of integers $[i, i + k - 1]$ (inclusive of the endpoints of the range). In simpler terms, a k-span captures the interval of a k-mer. We assume that the string S has at least k nucleotides, and therefore, at least one k-mer (and at least one k-span). For the sake of theoretical rigor, we use k-spans to develop our results, and later discuss how our practical implementation uses k-mers to estimate the substitution and indel rates.

▶ **Lemma 1.** *Let L' be the length of S'. Then*

$$\mathbf{E}[L'] = L\,(1 + d - p_d), \tag{1}$$
$$\mathbf{Var}[L'] = L\Big(d(d+1) + p_d(1 - p_d)\Big).$$

▶ **Lemma 2.** *Let f_A and $f_A{}'$ be the number of 'A's in S and S', respectively. Then,*

$$\mathbf{E}[f_A{}'] = f_A(1 - p_s - p_d) + \frac{p_s(L - f_A)}{3} + \frac{dL}{4}. \tag{2}$$

We can analogously obtain expressions for the expectations of the corresponding counts of 'C's, 'G's, and 'T's in S', denoted by $f_C{}'$, $f_G{}'$, and $f_T{}'$, respectively.

Next, we say that a k-span has no mutations if (a) for all k positions in the k-span, the mutation process picks a Stay$(\cdot)$ action, and if (b) for all $k-1$ intermediate positions in the k-span, the mutation process does not insert anything. Let $\mathcal{N}$ be the number of such k-spans. Let $\mathcal{K}$ be the number of all k-spans in S.

▶ **Lemma 3.** $\mathbf{E}[\mathcal{N}] = \mathcal{K}\,(1-p_s-p_d)^{\,k}\,\frac{1}{(d+1)^{\,k-1}}$.

The last quantities we will use correspond to the number of k-spans with a single kind of mutation. Let $\mathcal{S}$ be the number of k-spans with a single substitution, and no other mutations, let $\mathcal{D}$ be the number of k-spans with a single deletion and no other mutations, and let $\mathcal{I}$ be the number of k-spans with a single insertion and no other mutations.

▶ **Lemma 4.** *The expectations of $\mathcal{S}$, $\mathcal{D}$, and $\mathcal{I}$ are given by:*

$$\mathbf{E}[\mathcal{S}] = \mathcal{K}\,k\,(1-p_s-p_d)^{k-1}\,p_s\,\frac{1}{(d+1)^{\,k-1}}. \tag{3}$$

$$\mathbf{E}[\mathcal{D}] = \mathcal{K}\,k\,(1-p_s-p_d)^{k-1}\,p_d\,\frac{1}{(d+1)^{\,k-1}}. \tag{4}$$

$$\mathbf{E}[\mathcal{I}] = \mathcal{K}\,(k-1)\,(1-p_s-p_d)^{k}\,\frac{d}{(d+1)^{\,k}}. \tag{5}$$

4 Estimating p_s, p_d and d

We describe next our estimators for the parameters of the mutation model p_s, p_d, and d. We derive estimators for these rates based on the statistics defined in Section 3. Recall that S is the known input string to the mutation process, and S' is the resulting random string or observation.

Our mutation model assumes $p_s + p_d < 1$, and $d \geq 0$. In addition, we assume that $L \geq k$ which implies that $\mathbf{E}[\mathcal{N}] > 0$. In reality, strings are typically much longer than k-mer sizes and therefore this is a reasonable assumption.

From Lemmas 3 and 4 we get

$$\frac{\mathbf{E}[\mathcal{D}]}{\mathbf{E}[\mathcal{N}]} = \frac{k\,p_d}{1-p_s-p_d}. \tag{6}$$

From (1), (2) and (6), we obtain the following system of linear equations with variables p_s, p_d, and d.

$$-p_d + d = \frac{\mathbf{E}[L']}{L} - 1 \tag{7}$$

$$\frac{L-4f_A}{3}\,p_s - f_A\,p_d + \frac{L}{4}d = \mathbf{E}[f_A{}'] - f_A \tag{8}$$

$$\mathbf{E}[\mathcal{D}]p_s + (k\mathbf{E}[\mathcal{N}] + \mathbf{E}[\mathcal{D}])\,p_d = \mathbf{E}[\mathcal{D}]. \tag{9}$$

Solving this system of equations, we obtain that:

$$p_s = 3\frac{k\,\mathbf{E}[\mathcal{N}]\Big(\mathbf{E}[L'] - L + 4f_A - 4\mathbf{E}[f_A{}']\Big) + \mathbf{E}[\mathcal{D}]\Big(\mathbf{E}[L'] - 4\mathbf{E}[f_A{}']\Big)}{\Big(4f_A - L\Big)\Big(\mathbf{E}[\mathcal{D}] + 4k\mathbf{E}[\mathcal{N}]\Big)},$$

$$p_d = \frac{\mathbf{E}[\mathcal{D}](4f_A + 12\mathbf{E}[f_A{}'] - L - 3\mathbf{E}[L'])}{(4f_A - L)(\mathbf{E}[\mathcal{D}] + 4k\mathbf{E}[\mathcal{N}])},$$

$$d = \frac{\mathbf{E}[L']}{L} - 1 + \frac{\mathbf{E}[\mathcal{D}](-4f_A - 12\mathbf{E}[f_A{}'] + L + 3\mathbf{E}[L'])}{(-4f_A + L)(\mathbf{E}[\mathcal{D}] + 4k\mathbf{E}[\mathcal{N}])}.$$

Given S and S', we can compute L' and $f_A{}'$. Let us also assume that we know $\mathcal{N}$ and $\mathcal{D}$ (we will discuss how to find these in Section 6). By replacing the expectations above with these observations, we obtain our estimators for the parameters of the mutation model. That is,

$$\hat{p}_s = 3\frac{k\,\mathcal{N}\Big(L' - L + 4f_A - 4f_A{}'\Big) + \mathcal{D}\Big(L' - 4f_A{}'\Big)}{\Big(4f_A - L\Big)\Big(\mathcal{D} + 4k\mathcal{N}\Big)}, \tag{10}$$

$$\hat{p}_d = \frac{\mathcal{D}\left[3(L' - 4f_A{}') + (L - 4f_A)\right]}{(L - 4f_A)\,(\mathcal{D} + 4k\mathcal{N})}, \tag{11}$$

$$\hat{d} = \frac{L'}{L} - 1 + \frac{\mathcal{D}\left[3(L' - 4f_A{}') + (L - 4f_A)\right]}{(L - 4f_A)\,(\mathcal{D} + 4k\mathcal{N})}. \tag{12}$$

We briefly comment on our choice of estimators for the mutation model parameters, as various statistical approaches based on a different set of observables could yield a different set of estimators. For example, we considered a variant of the estimators based on the counts of k-spans with a single insertion, deletion, or substitution (i.e., $\mathcal{I}$, $\mathcal{S}$, and $\mathcal{D}$). These quantities contain enough information to estimate the mutation parameters, specifically, by solving the non-linear system of equations given by (3), (4), and (5); the resulting estimators performed quite well in real data. However, establishing theoretical guarantees for these estimators proved challenging, as they were defined as roots of degree-k polynomials . Our current estimators address this theoretical limitation as they involve solely linear equations. As we shall see in Section 7, the performance of our estimators in real data is strong, and they strike a more favorable balance by offering both reasonable accuracy and rigorous theoretical guarantees.

5 Concentration results

In this section, we provide a theoretical guarantee for the estimator (10) of the substitution rate. In particular, we show that our estimate of p_s is not only a consistent estimator, but also that it is tightly concentrated in a symmetric interval around the true value of p_s. Similar techniques can likely be used to prove the consistency and concentration of p_d and d, though we do not do so in this paper. We provide the bias analysis for $\hat{p}_s$ by proving asymptotically tight concentration bounds for $\mathcal{N}$, $\mathcal{D}$, and $L' - 4f_A{}'$.

The concentration bounds for $\mathcal{N}$ and $\mathcal{D}$ stem from the fact that, from the perspective of the mutation process k-spans located more than k positions apart are independent of each other.

▶ **Lemma 5.** *For any $\delta \in (0,1)$:*

$$\Pr\Big[\Big|\mathcal{N} - \mathbf{E}[\mathcal{N}]\Big| \geq \delta\mathbf{E}[\mathcal{N}]\Big] \;\leq\; 3k\exp\Big\{-\frac{\delta^2\mathbf{E}[\mathcal{N}]}{3k}\Big\}.$$

▶ **Lemma 6.** *For any $\delta \in (0,1)$:*

$$\Pr\Big[\Big|\mathcal{D} - \mathbf{E}[\mathcal{D}]\Big| \geq \delta\mathbf{E}[\mathcal{D}]\Big] \;\leq\; 3k\exp\Big\{-\frac{\delta^2\mathbf{E}[\mathcal{D}]}{3k}\Big\}.$$

We prove a similar result for $P = L' - 4f_A{}'$ by noting that P can expressed as a sum of independent sub-exponential random variables. Our concentration bound for P is as follows.

▶ **Lemma 7.** *Let $J_1 = \ln 2/\min\{d+1, 8\}$. Then, there exist absolute constants $c_1, c_2 > 0$ such that the following holds for any $\delta > 0$:*

$$\Pr\Big[|P - \mathbf{E}[P]| \geq 3\delta\Big] \leq 2\exp\Big\{-\frac{\delta^2}{8f_A}\Big\} + 2\exp\Big\{-\frac{\delta^2}{8(L - f_A)}\Big\} + 2\exp\Big\{-c_1\min\Big(\frac{\delta^2}{c_2{}^2J_1{}^2}, \frac{\delta}{c_2J_1}\Big)\Big\}.$$

Finally, we prove that L' is also strongly concentrated around $\mathbf{E}[L']$.

▶ **Lemma 8.** *For any $\delta \in (0,1)$:*

$$\Pr\left[\left|L' - \mathbf{E}[L']\right| \le \delta(Ld + Lp_d)\right] \;\ge\; 2\exp\left\{-\frac{L\delta^2 d^2}{2(d+1-\delta d)(d+1)}\right\} + 2\exp\left\{-\frac{Lp_d\delta^2}{3}\right\}.$$

We can piece the results of these lemmas together and prove the following result.

▶ **Theorem 9.** *Suppose $4f_A < L$ and $\frac{4}{3}p_s + p_d < 1$. Then, for sufficiently small $\delta > 0$, there exists constants $c_1, c_2 > 0$ such that*

$$\Pr[|\hat{p}_s - p_s| \ge 12\delta] \le 8k\exp\left\{-\frac{\delta^2\mathbf{E}[\mathcal{N}]}{3k}\right\} + 6k\exp\left\{-\frac{\delta^2\mathbf{E}[\mathcal{D}]}{3k}\right\} + 2\exp\left\{-\frac{\delta^2\mathbf{E}[P]^2}{72f_A}\right\}$$
$$+ 2\exp\left\{-\frac{\delta^2\mathbf{E}[P]^2}{72(L-f_A)}\right\} + 2\exp\left\{-c_1\min\left(\frac{\delta^2\mathbf{E}[P]^2}{{c_2}^2}, \frac{\delta\mathbf{E}[P]}{c_2}\right)\right\}.$$

The requirement that $4f_A < L$ in this theorem does not restrict generality: aside from equal nucleotide frequency (where the estimators are naturally invalid), at least one character $c \in \{A, C, T, G\}$ must satisfy $4f_c \le L$. In addition, the assumption that $\frac{4}{3}p_s + p_d < 1$ holds when p_s and p_d are small (e.g., $p_s < 1/4$ and $p_d < 1/4$ suffices) which is the case most frequently encountered in practice.

Theorem 9 is our central theoretical result, establishing that the estimators developed in Section 4 are sound under the reasonable conditions. Before discussing our implementation and presenting the experimental results, we comment on the error probability in Theorem 9. This probability is small when each of the terms in the sum are small. Since each of these terms decays (at least) exponentially with δ times an expectation that grows linearly with the length of the string (assuming the mutation parameters are fixed), they will all generally be small.

6 Implementation details

Our estimators for the three rates require counting the number of k-spans with single deletion, $\mathcal{D}$, and the number of k-spans with no mutation, $\mathcal{N}$. Counting $\mathcal{D}$ and $\mathcal{N}$ can be challenging, particularly because k-mers do not contain the contextual information, and so we do not have access to their corresponding k-spans. An additional layer of complexity comes into play from the fact that identifying a k-mer with no mutation (or a single kind of mutation) is more difficult, considering many edge cases that may arise from inserting the same character that has been deleted. These challenges are cimcumvented when we use k-spans, and therefore, counting $\mathcal{D}$ and $\mathcal{N}$ solely from k-mers is not trivial and can be considered an interesting problem in and of itself. We therefore implemented an ad hoc solution to estimate $\mathcal{D}$ and $\mathcal{N}$ given the two strings S and S'. The steps for estimating $\mathcal{D}$ and $\mathcal{N}$ are as follows.

We start by extracting all k-mers in S, and building a de Bruijn graph using these k-mers using the cuttlefish tool [12]. We then extract the unitigs from this graph. Let the set of unitigs computed from S be U. We also compute the unitigs in S' in a similar manner and call this U'. We next take an arbitrary unitig u from U, and align every unitig in U' with u. To allow for partial overlap, we use semi-global alignment by using the infix option in edlib [21], which makes sure that gaps at the beginning and at the end of the alignment are not penalized. For a particular $u \in U$, we align every $u' \in U'$ to make sure all relevant alignments are considered. We use these alignments to look at all windows of length k, and count $\mathcal{D}$ and $\mathcal{N}$ accordingly for u. We repeat this for all $u \in U$, and accumulate the measurements from individual u's into a single global count.

The motivation behind using unitigs is that if there is an isolated mutation, and if the mutation is in the first or the last position of a k-mer, then there is no way to understand if the mutation is a

substitution, an insertion, or a deletion only from the k-mers. The only way to resolve this ambiguity (and other similar ambiguities) is to scan beyond the context of k characters – and unitigs are a natural way to do this. The core goal of our implementation of these steps described above was not to make it efficient, but rather to obtain a working solution. We found that executing these steps estimates $\mathcal{D}$ and $\mathcal{N}$ reasonably well, and the estimated rates are also acceptable. As such, we leave finding an efficient way to compute $\mathcal{D}$ and $\mathcal{N}$ as an open research question.

7 Experiments and results

In this section, we present a series of experiments to evaluate the performance of the estimators detailed in Section 4. As discussed earlier, these estimators are sensitive to several input parameters, including k-mer size, sequence length, and the fraction of 'A' characters in the sequence. Sections 7.1 through 7.3 explore the sensitivity of the estimators with respect to these parameters. In Sections 7.4 and 7.5, we estimate mutation rates across a wide range of known rate combinations. And finally, in Section 7.6, we demonstrate that our estimated substitution rate outperforms estimates obtained under a substitution-only mutation model. For the experiments in Sections 7.1 through 7.4, the original sequence is a randomly generated synthetic sequence. In these cases, we compute the number of k-spans containing a single deletion and the number of k-spans with no mutation directly from the known mutation process. For the experiments in Sections 7.5 and 7.6, we use real reference genomes as the original sequences. In these cases, the two types of k-span counts are estimated using the steps described in Section 6.

7.1 Sensitivity of the estimators to k-mer lengths

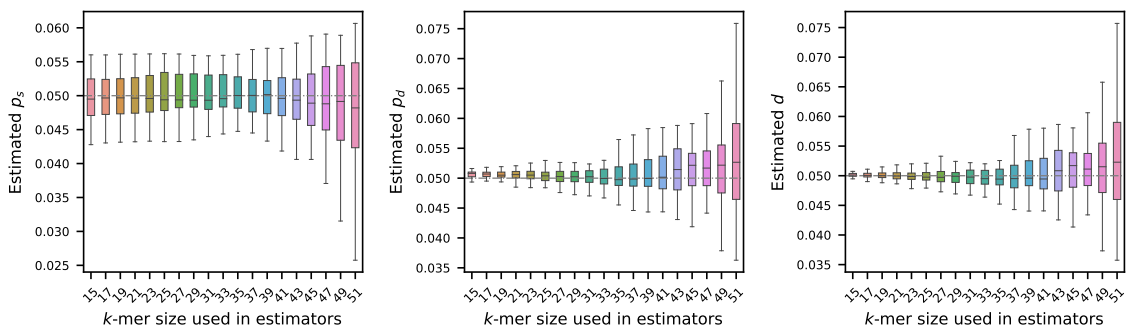

Figure 1 *Effect of k-mer length on mutation rate estimation (true rates set to 0.05).* A synthetic genome of 1 million nucleotides, mutated genomes were generated by setting p_s, p_d, and $d = 0.05$ – shown by the gray dashed horizontal line. Estimated rates were then computed using a range of k-mer sizes. Each boxplot shows the variability in estimation across 20 simulations, with error bars showing one standard deviation. The plots show that the estimators become more accurate and more precise for shorter k-mers.

We begin our analysis by examining how the choice of k-mer length affects the estimation of the mutation rates. To investigate this effect, we first generated a synthetic reference sequence of 1 million nucleotides, randomly sampling bases with fixed frequencies: 30% 'A', and equal proportions of 'C', 'G', and 'T' – making sure total frequency is 100%. From this reference, we simulated 20 mutated sequences, independently from each other, using the mutation model described in Section 2. For each of these mutated sequences, we estimated mutation rates using the estimators defined in Section 4 using a range of values for k. The results of this analysis are summarized in Figure 1.

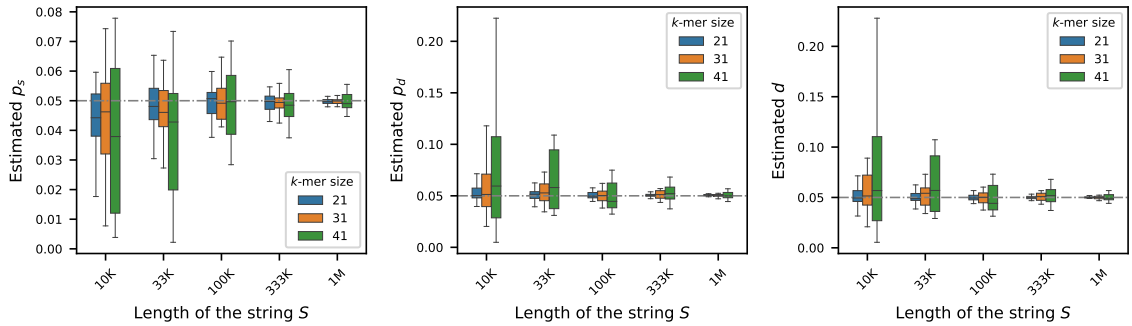

Figure 2 *Effect of sequence length on mutation rate estimation (true rates set to 0.05).* For synthetic genomes of varying lengths, mutated genomes were generated by setting p_s, p_d, and $d = 0.05$ – shown by the gray dashed horizontal line. Estimated rates were computed using three k-mer sizes: 21, 31, and 41. Each boxplot shows the variability in estimation across 20 simulations. The plots show that the estimators become more accurate and more precise for longer sequences.

As illustrated, the choice of k has a substantial impact on the stability of the estimators. In particular, longer k-mers tend to produce estimates with higher variability. This behavior is consistent with the known sensitivity of k-mers to mutations: since a single mutation can disrupt up to k consecutive k-mers, the longer the k-mer, the more susceptible it becomes to such perturbations. Our theoretical result in Theorem 9 also captures this: with a larger k, the error probabilities become larger, and the probabilistic guarantee for the estimators' performances decreases accordingly.

Interestingly, for the estimator of substitution probability p_s, we observe that the variability in the estimated values does not change significantly from 15 to 39. The reason behind this behavior is not immediately clear and warrants further investigation. It is possible that incorporating the number of k-spans with a single substitution $\mathcal{S}$ into the estimators may correct this behavior, but additional analyses are required to substantiate this hypothesis.

7.2 Sensitivity of the estimators to sequence length

To investigate how the length of the original sequence S influences estimation of the mutation rates, we simulated synthetic genomes ranging from 10K to 1M nucleotides in length. For each genome length, we generated 10 independent synthetic sequences to capture variability due to random sampling. The nucleotide composition of each sequence was fixed, with the frequency of 'A' set to 30% and frequencies of 'C', 'G', and 'T' set equally – making sure total frequency is 100%. For each synthetic sequence, we generated its mutated version by running the mutation process described in Section 2, setting each of p_s, p_d, and d to 0.05. We then estimated the mutation rates using the estimators outlined in Section 4 for three k-mer sizes: $k = 21, 31, 41$.

Figure 2 displays the estimated rates across the varying sequence lengths. As shown, the estimators are less stable for shorter sequences. However, with longer sequences, the estimators yield more accurate results – a trend expected from our core theoretical result in Theorem 9, which states that the associated error is asymptotically vanishing in L, the length of the string S: as L increases, the number of k-spans $\mathcal{K}$ increases, and therefore the probability of error decreases, leading to a more precise estimation.

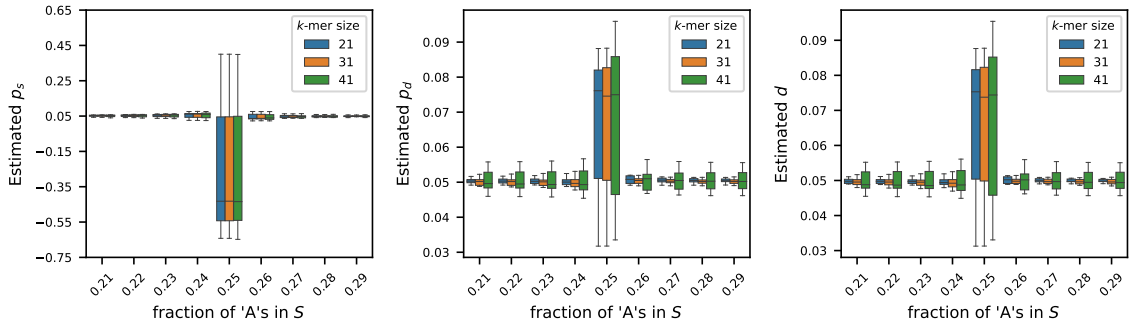

Figure 3 *Effect of nucleotide composition on mutation rate estimation (true rates set to 0.05).* For synthetic genomes of length $L = 1$M, the fraction of 'A's is varied from 0.21 to 0.29, and the frequencies of 'G's, 'C's, and 'T's are set equally. For each setting, the mutated string S' was generated by setting $p_s, p_d, d = 0.05$. Estimated rates were computed for three k-mer sizes: 21, 31, and 41. Each boxplot shows the variability in estimation across 20 simulations. The results show that the estimators generally work well for all three k-mer sizes, except when $f_A \approx L/4$, in which case the estimators become unstable – as predicted by Theorem 9.

7.3 Sensitivity of the estimators to base composition

We generated synthetic genomes of 1M nucleotides to investigate how the fraction of 'A' characters affects the estimation of the mutation rates. We varied the fraction of 'A's from 21% to 29% in increments of 1%. For each fraction of 'A's, we set the frequency of 'C's, 'G's, and 'T's equally. For each preset fraction of 'A's, we generated 10 random genomes to capture stochastic variation. For each of these genomes, we generated its mutated version using the mutation process described in Section 2, setting each of p_s, p_d, and d to 0.05. We then estimated the mutation rates using the estimators described in Section 4 using three k-mer sizes: $k = 21, 31$, and 41. Figure 3 shows the sensitivity of the estimators to the fraction of 'A's in the original string S.

We observe that the estimators $\hat{p_s}$, $\hat{p_d}$, and $\hat{d}$ work reasonably well to estimate the true rates when the frequency of 'A' characters, f_A is not $L/4$. On the contrary, when the fraction of 'A's is exactly 25% in the original string S, the estimator gives inaccurate values, some of which are even negative (see estimated values of p_s). This behavior is captured in Theorem 9: when $L - 4f_A = 0$, $\mu_P = \mathbf{E}[P] = 0$, and the probabilistic guarantees become unbounded. We only get meaningful guarantees of consistency when $4f_A$ is strictly smaller than L. While Theorem 9 does not guarantee consistency when $4f_A > L$, this does not restrict generality, as explained in Section 4, and as demonstrated by the estimators' performances when the fraction of 'A's is larger than 1/4.

7.4 Estimating rates from a randomly generated synthetic sequence

After testing our estimators for varying k-mer lengths, sequence lengths, and base compositions, we next turn to estimating mutation rates by varying the true rates across a range of values. To do this, we generated a synthetic reference genome of 1 million base pairs, fixing the base composition at 30% 'A', and equal proportions of 'C', 'G', and 'T' – making sure total frequency is 100%. Using the mutation model described in Section 2, we then simulated mutated genomes from the synthetic reference by varying the mutation rates p_s, p_d, and d across the values $\{0.01, 0.02, 0.03, 0.04, 0.05\}$. For every parameter combination, we generated 10 independent mutated genomes to capture stochastic variability. We then estimated the mutation rates using the estimators detailed in Section 4 for each of these mutated genomes.

In Figure 4, we show two sets of results:

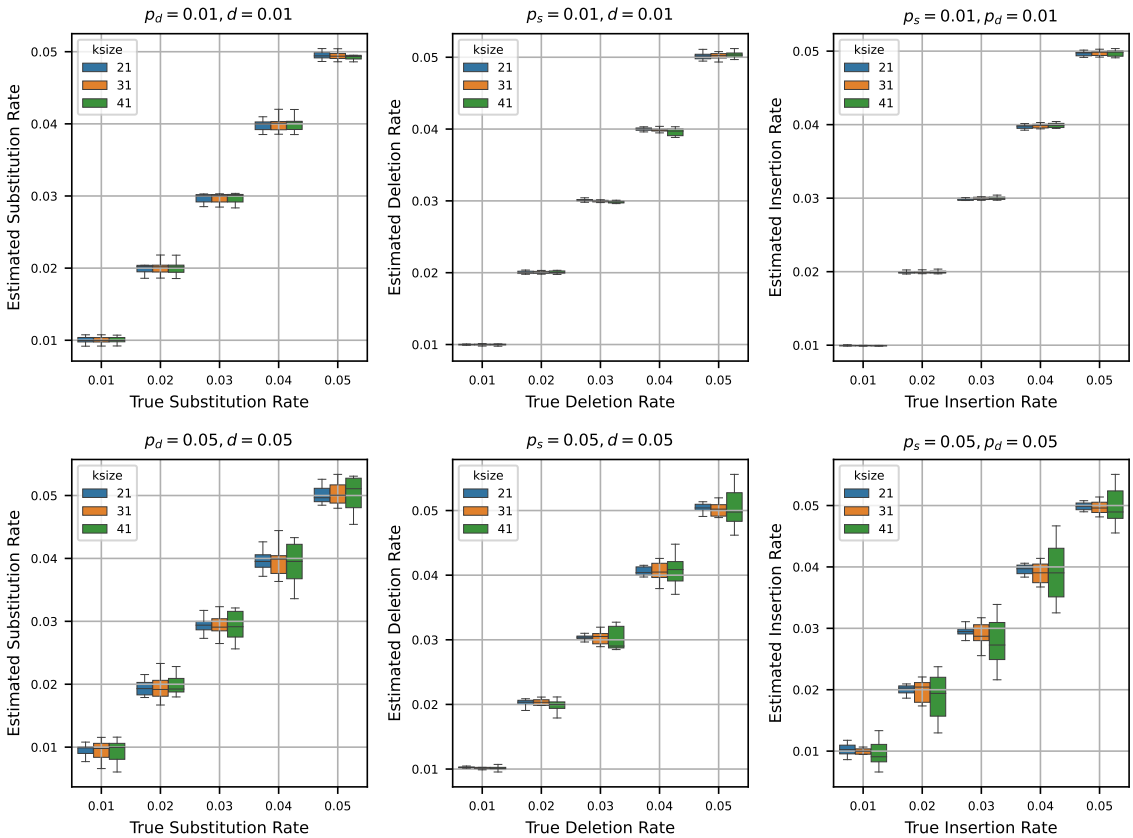

Figure 4 *Estimated mutation rates versus true values, where the original string is a synthetic sequence*. Each subplot corresponds to a case where two mutation rates are fixed (either at 0.01 or 0.05) and the third is varied from 0.01 to 0.05. Each boxplot shows the variability in estimation across 10 simulations. The results show that the rate estimation is very accurate when the other two rates are small, and is reasonably accurate when the other two rates are larger.

- **Fixed low rates (0.01):** Fixing two of the rates at 0.01, we show the estimates of the third rate as the true rate varies from 0.01 to 0.05. We repeat this process independently for p_s, p_d, and d.
- **Fixed high rates (0.05):** Fixing two of the rates at a higher value of 0.05, we also show the estimates of the third rate as the true rate varies from 0.01 to 0.05. Again, we repeat this for all three rates.

When we set the other two rates to 0.01 and estimate the third rate, we observe that the estimated rates are highly accurate across all trials. In many cases, the boxplots of the estimates nearly vanish due to minimal variance, indicating tight clustering around the true values. This trend remains consistent across multiple k-mer sizes, suggesting that the estimators are robust at low rates of mutation.

In contrast, when we fix the other two rates at 0.05 and estimate the third rate, the accuracy of the estimation decreases slightly. While the estimates still track the true values reasonably well, the variance increases, and the boxplots become more prominent. Notably, the median estimate remains close to the true rate in most settings, which indicates that the estimators retain their central tendency even under higher mutation rates. However, for larger k-mer sizes, we observe increased variability in the estimates – an effect that mirrors our earlier observations in Section 7.1, where longer k-mers

resulted in decreased precision of the estimators.

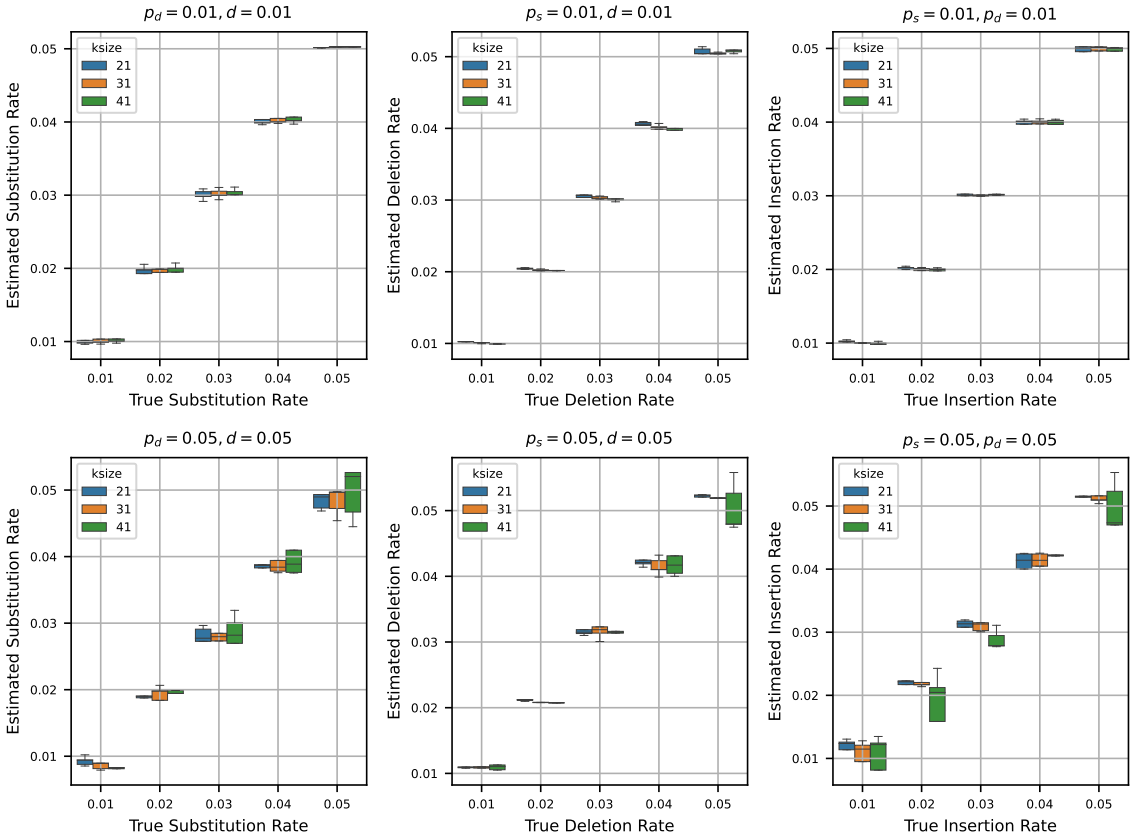

Figure 5 *Estimated mutation rates versus true values, where the original string is the reference genome of Staphylococcus (length 2.8 million).* Each subplot corresponds to a case where two mutation rates are fixed (either at 0.01 or 0.05) and the third varies from 0.01 to 0.05. Each boxplot shows the variability in estimation across 10 simulations. Estimated mutation rates closely match true values when other rates are low, but estimation becomes less precise under high total mutation rates (>10%) due to increased difficulty in real genomic contexts.

7.5 Estimating rates from real sequences

Having the estimators tested for a synthetic reference, we next estimate rates from a real genome sequence. For this set of experiments, we used the reference assembly of *Staphylococcus aureus* (subspecies: aureus USA300_TCH1516), which has 2.8 million nucleotides. We simulated mutated sequences from this reference by running the mutation process described in Section 2 by varying the mutation rates p_s, p_d, and d from 0.01 to 0.05. Similar to Section 7.4, we generated 10 independent mutated sequences for each combination of p_s, p_d, and d to capture stochastic variability. We then estimated the mutation rates using the estimators outlined in Section 4 for each of these simulated sequences.

Figure 5 shows the estimated mutation rates plotted against the true rates that were used to run the mutation process. We observe that the results shown in Figure 5 are consistent with previously discussed results. Specifically, when estimating a given mutation rate while keeping the other two rates low (0.01), the corresponding estimator performs with high precision, closely tracking the true value. On the other hand, when the estimation is carried out with the other two mutation rates set

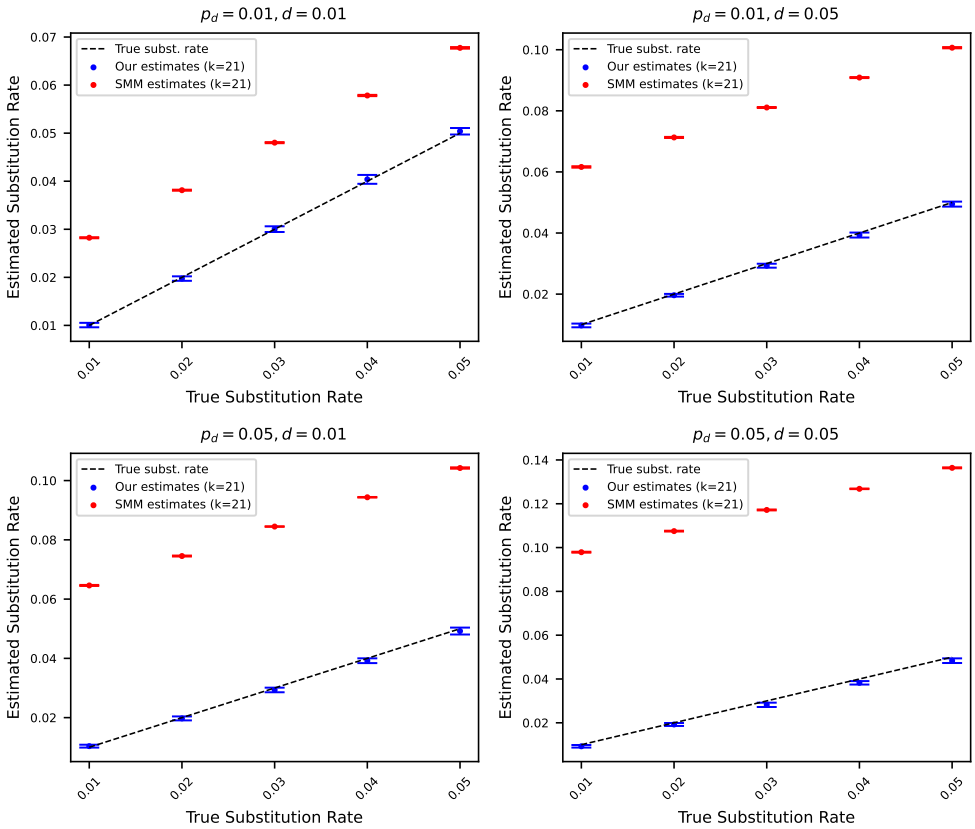

Figure 6 *Estimated substitution rates versus true substitution rates, where the original string is the reference genome of Staphylococcus (length 2.8 million). Rates were estimated using* (10) *and using a simple mutation model (SMM) that only considers substitution.* Each subplot corresponds to a case where p_d and d are fixed at either 0.01 or 0.05. The points show the average of 10 estimates, and the error bars show one standard deviation. The dashed black line corresponds to the true substitution rates. Estimated substitution rates using our method closely match true rates, whereas SMM overestimates due attributing to substitutions mismatching k-mers originating from insertions and deletions.

to higher values (0.05), the estimates appear more confounded. This is likely due to the increased difficulty of accurately estimating the number of k-spans with a single deletion ($\mathcal{D}$) or no mutation ($\mathcal{N}$) in a real genomic context. Notably, these experiments involve challenging conditions, with total mutation rates exceeding 10%. Despite this, the estimators yield reasonably accurate results, indicating potential for practical effectiveness.

7.6 Comparison with substitution rates estimated using simple mutation model

We conclude the experiments section by contrasting the substitution rates estimated using (10) with substitution rates estimated considering a simple mutation model. We use the statistics of k-mers developed in a recent work [2] to estimate substitution rates under a simple mutation model. The simple mutation model captures only substitutions, and no insertions or deletions. Consequently, we can only compute substitution rates considering this simple model. Henceforth, we refer to the simple mutation model as SMM.

We estimated the substitution rates using the SMM for the same simulated mutated sequences described in Section 7.5. In Figure 6, we show the estimated substitution rates using our estimators in (10), and using SMM. The results highlight that the substitution rates estimated using the estimator we developed track the true substitution rates accurately. On the other hand, the substitution rates estimated using SMM make a gross overestimation. This is because the SMM does not consider indels, and therefore, the effects of all three mutation rates are subsumed in the single substitution rate we get using the SMM. As such, a simple mutation model cannot disentangle the distinct contributions of substitution, insertion, and deletion rates. In contrast, the mutation model we introduce effectively decomposes these components, enabling more accurate and meaningful estimation of individual mutation rates.

8 Conclusions

We have presented a mutation model that accommodates single-nucleotide substitutions, as well as insertions and deletions while retaining enough mathematical structure to admit closed-form rate estimators derived solely from k-mer statistics. From this model, we obtained algebraic estimators for the three elementary mutation rates: p_s, p_d, and d; we also proved a relatively tight sub-exponential concentration bound on p_s that guarantees consistency as sequence length grows. We also identified regimes in which the estimation becomes ill-conditioned (i.e. large k, $p_d = 0$, or sequence composition with 25% 'A'). These results establish a bridge between sequence evolution and combinatorial word statistics, thus providing additional tools for theoretical algorithmic computational biology.

In our prototype implementation, we demonstrated that our estimates remain accurate on simulated evolution of real genomes, and outperforms a substitution-only simple mutation model by avoiding spurious attribution of indel signals. While naive counting of unmutated and single-deletion k-mers sufficed to show practical accuracy of our estimators, this raises an interesting open problem: estimating the number of these unmutated and single-deletion k-mers efficiently for large scale data sets.

Several directions invite further investigation. First, incorporating the count $\mathcal{S}$ of single-substitution k-spans may illuminate why p_s remains relatively stable even for moderately large k. Second, our framework can extend to heterogeneous or context-dependent rates by replacing global expectations with position-specific covariates. Third, coupling our estimators with sketch-based distance measures (such as in [10]) may provide a theory-backed avenue for larger scale applications such as phylogenetic placement in the presence of high indel activity. Finally, a more thorough investigation on real genomic data (where the unitig-based approach we used in the practical implementation starts to become infeasible) will be necessary to understand the utility of the mutation estimates in practice.

In summary, by utilizing probabilistic modeling and concentration inequalities, we provide a theoretical foundation and initial practical implementation for quantifying the parameters of a relatively complex mutation process directly from k-mers. We anticipate that these ideas will continue to inform new alignment-free computational biology tasks, particularly relevant as sequencing data continues to outpace traditional alignment-based paradigms.

References

1. Dana Angluin and Leslie G Valiant. Fast probabilistic algorithms for hamiltonian circuits and matchings. In *Proceedings of the ninth annual ACM symposium on Theory of computing*, pages 30–41, 1977.
2. Antonio Blanca, Robert S Harris, David Koslicki, and Paul Medvedev. The statistics of k-mers from a sequence undergoing a simple mutation process without spurious matches. *Journal of Computational Biology*, 29(2):155–168, 2022.
3. Ke Chen, Vinamratha Pattar, and Mingfu Shao. Sequence similarity estimation by random subsequence sketching. *bioRxiv*, pages 2025–02, 2025.
4. Herman Chernoff. A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *Annals of Mathematical Statistics*, 23(4):493–507, 1952.
5. Margaret Oakley Dayhoff. Computer analysis of protein evolution. *Scientific American*, 221(1):86–95, 1969.
6. Richard Durbin, editor. *Biological sequence analysis*. Cambridge Univ. Press, Cambridge [u.a.], 17. print. edition, 2013.
7. Arun Ganesh and Aaron Sy. Near-Linear Time Edit Distance for Indel Channels. In Carl Kingsford and Nadia Pisanti, editors, *20th International Workshop on Algorithms in Bioinformatics (WABI 2020)*, volume 172 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 17:1–17:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPIcs.WABI.2020.17.
8. Steven Henikoff and Jorja G Henikoff. Amino acid substitution matrices from protein blocks. *Proceedings of the National Academy of Sciences*, 89(22):10915–10919, 1992.
9. Mahmudur Rahman Hera, Shaopeng Liu, Wei Wei, Judith S Rodriguez, Chunyu Ma, and David Koslicki. Metagenomic functional profiling: to sketch or not to sketch? *Bioinformatics*, 40(Supplement_2):ii165–ii173, 2024.
10. Mahmudur Rahman Hera, N Tessa Pierce-Ward, and David Koslicki. Deriving confidence intervals for mutation rates across a wide range of evolutionary distances using FracMinHash. *Genome Research*, 33(7):1061–1068, 2023.
11. Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963.
12. Jamshed Khan and Rob Patro. Cuttlefish: fast, parallel and low-memory compaction of de bruijn graphs from large-scale genome collections. *Bioinformatics*, 37(Supplement_1):i177–i186, 2021.
13. Brian D Ondov, Gabriel J Starrett, Anna Sappington, Aleksandra Kostic, Sergey Koren, Christopher B Buck, and Adam M Phillippy. Mash Screen: high-throughput sequence containment estimation for genome discovery. *Genome Biology*, 20:1–13, 2019.
14. Mahmudur Rahman Hera and David Koslicki. Estimating similarity and distance using fracminhash. *Algorithms for Molecular Biology*, 20(1):1–13, 2025.
15. Sophie Röhling, Alexander Linne, Jendrik Schellhorn, Morteza Hosseini, Thomas Dencker, and Burkhard Morgenstern. The number of k-mer matches between two dna sequences as a function of k and applications to estimate phylogenetic distances. *Plos one*, 15(2):e0228070, 2020.
16. Shahab Sarmashghi, Kristine Bohmann, M Thomas P Gilbert, Vineet Bafna, and Siavash Mirarab. Skmer: assembly-free and alignment-free sample identification using genome skims. *Genome Biology*, 20(1):1–20, 2019.
17. Tizian Schulz and Paul Medvedev. ESKEMAP: exact sketch-based read mapping. *Algorithms for Molecular Biology*, 19(1):19, 2024.
18. Jim Shaw and Yun William Yu. Fast and robust metagenomic sequence comparison through sparse chaining with skani. *Nature Methods*, 20(11):1661–1665, 2023.
19. Jim Shaw and Yun William Yu. Rapid species-level metagenome profiling and containment estimation with sylph. *Nature Biotechnology*, pages 1–12, 2024.
20. Kai Song, Jie Ren, Gesine Reinert, Minghua Deng, Michael S Waterman, and Fengzhu Sun. New developments of alignment-free sequence comparison: measures, statistics and next-generation sequencing. *Briefings in bioinformatics*, 15(3):343–353, 2014.
21. Martin Šošić and Mile Šikić. Edlib: a c/c++ library for fast, exact sequence alignment using edit distance. *Bioinformatics*, 33(9):1394–1395, 2017.

22 Shiheng Tao, Yanhui Fan, Wenjuan Wang, Guoji Ma, Lijing Liang, and Qi Shi. Patterns of insertion and deletion in mammalian genomes. *Current Genomics*, 8(6):370–378, 2007.

23 Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.

24 Kun Wu, Danqi Qin, Yang Qian, and Haoxuan Liu. A new era of mutation rate analyses: Concepts and methods. *Zoological Research*, 45(4):767, 2024.

25 Andrzej Zielezinski, Susana Vinga, Jonas Almeida, and Wojciech M Karlowski. Alignment-free sequence comparison: benefits, applications, and tools. *Genome biology*, 18:1–17, 2017.

A List of notations used in the manuscript

For convenience, we gather in the following table the notation used in the manuscript.

Notation	Meaning
p_s	Substitution rate
p_d	Deletion rate
d	Average insertion length
S	The original string
S'	The mutated string
L	Length of (number of characters in) S
L'	Length of (number of characters in) S'
f_A, f_C, f_G, f_T	Number of 'A', 'C', 'G', 'T' in S, respectively. Here, $f_A + f_C + f_G + f_T = L$.
f_A', f_C', f_G', f_T'	Number of 'A', 'C', 'G', 'T' in S', respectively. Here, $f_A' + f_C' + f_G' + f_T' = L'$.
P	$P = L' - 4f_A'$
$\mathcal{K}$	Number of k-spans in S. Here, $\mathcal{K} = L - k + 1$.
$\mathcal{K}'$	Number of k-spans in S'. Here, $\mathcal{K}' = L' - k + 1$.
S_i	The character in S at position i
Q_i	The string inserted after S_i by the mutation process
l_i	Length of Q_i
$\mathcal{K}_i$	The interval $[i, i + k - 1]$, inclusive of the endpoints
$\mathcal{N}$	Number of k-spans that were not mutated by the mutation process
$\mathcal{S}$	Number of k-spans where a single substitution, and no other mutation was introduced by the mutation process
$\mathcal{D}$	Number of k-spans where a single deletion, and no other mutation was introduced by the mutation process
$\mathcal{I}$	Number of k-spans where a single insertion, and no other mutation was introduced by the mutation process
$\mathcal{L}_{\text{ins}}$	Total number of characters inserted by the mutation process.
$\mathcal{L}_{\text{del}}$	Total number of characters deleted by the mutation process.

Table 1 List of notations.

B Preliminaries

This section gathers several definitions and standard concentration inequalities that we use in our proofs. For a random variable X with finite non-zero variance σ^2, Chebyshev's inequality says that for any real value $r > 0$

$$\Pr\Big[|X - \mathbf{E}[X]| \geq r\sigma\Big] \leq \frac{1}{r^2}.$$

Sums of independent random variables can exhibit sharper concentration, and this is captured by the following standard inequalities. The first one is the Chernoff bound for sums of independent and identically distributed Bernoulli random variables [1, 4].

▶ **Lemma 10.** *Let $X_1, \ldots, X_n$ be a sequence of independently and identically distributed Bernoulli random variables. Consider $X = \sum_{i=1}^{n} X_i$ and let $\mathbf{E}[X] = \mu$. Then, the following three inequalities*

hold:

$$\Pr(X \geq (1+\delta)\mu) \leq e^{-\delta^2\mu/(2+\delta)}, \qquad 0 \leq \delta.$$

$$\Pr(X \leq (1-\delta)\mu) \leq e^{-\delta^2\mu/2}, \qquad 0 < \delta < 1. \tag{13}$$

$$\Pr(|X-\mu| \geq \delta\mu) \leq 2e^{-\delta^2\mu/3}, \qquad 0 < \delta < 1. \tag{14}$$

A more general inequality that applies to sums of independent and bounded random variables is Hoeffding's inequality [11].

▶ **Lemma 11.** *Let $X_1, \ldots, X_n$ be a sequence of independently and identically distributed random variables taking values in $[a, b]$. Let $X = \sum_{i=1}^n X_i$, and $\mathbf{E}[X] = \mu$. Then, for any $\delta > 0$:*

$$\Pr\left(|X-\mu| \geq \delta\right) \leq 2e^{-\frac{2\delta^2}{n(b-a)^2}}.$$

We will also analyze sums of certain *unbounded* random variables. For this, we recall the definition of sub-exponential random variables. For a comprehensive survey about this type of random variables, the reader is referred to [23, Chapter 2.7].

▶ **Definition 12.** *A random variable X such that $\mathbf{E}[X] = 0$ is called sub-exponential if either of the following equivalent properties holds:*

1. *There exists $J_1 > 0$ such that for all $t \geq 0$:*

$$\Pr\Big[|X| \geq t\Big] \leq 2\exp\{-t/J_1\}.$$

2. *There exists $J_2 > 0$ such that*

$$\mathbf{E}\Big[e^{|X|/J_2}\Big] \leq 2. \tag{15}$$

The constants J_1 and J_2 differ by an absolute multiplicative constant. The smallest J_2 such that (15) holds is called the *sub-exponential norm of X* and is denoted by $\|X\|_{\psi_1}$. For sums of sub-exponential random variables, the following version of Bernstein's inequality is known; see [23, Theorem 2.8.1].

▶ **Lemma 13.** *Let $X_1, \ldots, X_n$ be a sequence of independent, mean-zero, sub-exponential random variables, and let $X = \sum_{i=1}^n X_i$. Then, there exists an absolute constant $c > 0$ such that for every $t \geq 0$:*

$$\Pr\left[|X| \geq t\right] \leq 2\exp\left[-c\min\left(\frac{t^2}{\sum_{i=1}^N \|X_i\|_{\psi_1}^2}, \frac{t}{\max_i \|X_i\|_{\psi_1}}\right)\right].$$

C Proofs of preliminary statistics

We prove in this section the lemmas from Section 3.

▶ **Lemma 1.** *Let L' be the length of S'. Then*

$$\mathbf{E}[L'] = L\,(1 + d - p_d), \tag{1}$$

$$\mathbf{Var}[L'] = L\Big(d(d+1) + p_d(1-p_d)\Big).$$

Proof. For $1 \leq i \leq L$, let D_i be an indicator random variable defined as:

$$D_i = \begin{cases} 1 & \text{if character } S_i \text{ is removed,} \\ 0 & \text{otherwise.} \end{cases}$$

Using the D_i's and l_i's, we can express L' as:

$$L' = L + \sum_{i=1}^{L} l_i - \sum_{i=1}^{L} D_i, \tag{16}$$

and the linearity of expectation yields that

$$\mathbf{E}[L'] = L + \sum_{i=1}^{L} \mathbf{E}[l_i] - \sum_{i=1}^{L} \mathbf{E}[D_i].$$

We note that l_i is a geometric distribution with mean of d; i.e., $\mathbf{E}[l_i] = d$. Moreover, $\mathbf{E}[D_i] = \Pr[D_i = 1] = p_d$, and thus we obtain that $\mathbf{E}[L'] = L(1 + d - p_d)$ as claimed.

To compute the variance of L', note that D_i's and l_i's are all independent. Hence,

$$\mathbf{Var}[L'] = \mathbf{Var}\Big[L + \sum_{i=1}^{L} l_i - \sum_{i=1}^{L} D_i\Big] = \sum_{i=1}^{L} \mathbf{Var}[l_i] - \sum_{i=1}^{L} \mathbf{Var}[D_i].$$

The D_i's are Bernoulli random variables, and therefore, $\mathbf{Var}[D_i] = p_d(1 - p_d)$. Moreover, the l_i's are geometrically distributed with mean d, so $\mathbf{Var}[l_i] = d(d + 1)$. This yields that $\mathbf{Var}[L'] = L(d(d + 1) + p_d(1 - p_d))$. ◀

▶ **Lemma 2.** *Let f_A and $f_A{}'$ be the number of 'A's in S and S', respectively. Then,*

$$\mathbf{E}[f_A{}'] = f_A(1 - p_s - p_d) + \frac{p_s(L - f_A)}{3} + \frac{dL}{4}. \tag{2}$$

Proof. Let A_i be an indicator defined as:

$$A_i = \begin{cases} 1 & \text{if } a_i \text{ was a stay operation,} \\ 0 & \text{otherwise.} \end{cases}$$

Note that $\mathbf{E}[A_i] = \Pr[A_i = 1] = 1 - p_s - p_d$. Let us also define B_i as :

$$B_i = \begin{cases} 1 & \text{if } a_i \text{ was a substitute operation and } S'_i = \text{'A'}, \\ 0 & \text{otherwise,} \end{cases}$$

and observe that $\mathbf{E}[B_i] = \Pr[B_i = 1] = p_s/3$.

Finally, let C_i be the number of 'A's in Q_i. Then,

$$\mathbf{E}[C_i] = \sum_{\ell \geq 0} \mathbf{E}[C_i \mid l_i = \ell] \Pr[l_i = \ell] = \sum_{\ell \geq 0} \frac{\ell}{4} \Pr[l_i = \ell] = \frac{d}{4}. \tag{17}$$

We can express $f_A{}'$ as:

$$f_A{}' = \sum_{i:S_i = \text{'A'}} A_i + \sum_{i:S_i \neq \text{'A'}} B_i + \sum_{i=1}^{L} C_i, \tag{18}$$

and thus

$$\mathbf{E}[f_A{}'] = \sum_{i:S_i = \text{'A'}} \mathbf{E}[A_i] + \sum_{i:S_i \neq \text{'A'}} \mathbf{E}[B_i] + \sum_{i=1}^{L} \mathbf{E}[C_i] = f_A(1 - p_s - p_d) + \frac{(L - f_A)p_s}{3} + \frac{dL}{4},$$

as claimed. ◀

▶ **Lemma 3.** $\mathbf{E}[\mathcal{N}] = \mathcal{K}\,(1 - p_s - p_d)^{\,k}\,\frac{1}{(d+1)^{\,k-1}}$.

Proof. Let N_i be the indicator variable that equals 1 if the k-span $\mathcal{K}_i$ has no mutations. Then, $\mathcal{N} = \sum_{i=1}^{\mathcal{K}} N_i$. For a k-span to have no mutations (i.e., $N_i = 1$), the mutation model would have to choose the stay operation for all of its k positions, the probability of which is $(1 - p_s - p_d)^k$. Additionally, for all the $k - 1$ insert positions, the mutation model would need to choose strings of length zero. Therefore, $\Pr[N_i = 1] = (1 - p_s - p_d)^k/(d + 1)^{k-1}$, and thus

$$\mathbf{E}[\mathcal{N}] = \sum_{i=1}^{\mathcal{K}} \Pr[N_i = 1] = \mathcal{K}\frac{(1 - p_s - p_d)^k}{(d + 1)^{k-1}},$$

as claimed.

◀

▶ **Lemma 4.** *The expectations of $\mathcal{S}$, $\mathcal{D}$, and $\mathcal{I}$ are given by:*

$$\mathbf{E}[\mathcal{S}] = \mathcal{K}\,k\,(1 - p_s - p_d)^{k-1}\,p_s\,\frac{1}{(d+1)^{\,k-1}}. \tag{3}$$

$$\mathbf{E}[\mathcal{D}] = \mathcal{K}\,k\,(1 - p_s - p_d)^{k-1}\,p_d\,\frac{1}{(d+1)^{\,k-1}}. \tag{4}$$

$$\mathbf{E}[\mathcal{I}] = \mathcal{K}\,(k-1)\,(1 - p_s - p_d)^{k}\,\frac{d}{(d+1)^{\,k}}. \tag{5}$$

Proof. Let E_i be an indicator random variable defined as follows for $1 \le i \le \mathcal{K}$:

$$E_i = \begin{cases} 1 & \text{if the } k\text{-span } \mathcal{K}_i \text{ has a single substitution and no other mutation,} \\ 0 & \text{otherwise.} \end{cases}$$

We note that for a k-span to have a single substitution and no other mutations, the mutation process needs to insert 0 length strings in all the $k - 1$ insert positions, and other than a single substitution position, for all other $k-1$ positions in the k-span, the mutation process needs to select a stay operation. The single substitution operation could be at any of the k positions. Then,

$$\mathbf{E}[E_i] = \frac{k(1 - p_s - p_d)^{k-1} p_s}{(d + 1)^{k-1}}.$$

Since $\mathcal{S} = \sum_{i=1}^{\mathcal{K}} E_i$, we obtain that

$$\mathbf{E}[\mathcal{S}] = \sum_{i=1}^{\mathcal{K}} \mathbf{E}[E_i] = \mathcal{K}k(1 - p_s - p_d)^{k-1} p_s/(d + 1)^{k-1}.$$

The expectations of $\mathcal{D}$ and $\mathcal{I}$ can be derived following an analogous argument; the details are omitted. ◀

D Proofs of concentration results

We prove in this section the results from Section 5.

▶ **Lemma 5.** *For any $\delta \in (0, 1)$:*

$$\Pr\Big[\big|\mathcal{N} - \mathbf{E}[\mathcal{N}]\big| \ge \delta\mathbf{E}[\mathcal{N}]\Big] \;\le\; 3k\exp\Big\{-\frac{\delta^2\mathbf{E}[\mathcal{N}]}{3k}\Big\}.$$

Proof. As in the proof of Lemma 3, let N_i be the indicator variable that equals 1 if the k-span $\mathcal{K}_i$ has no mutations so that $\mathcal{N} = \sum_{i=1}^{\mathcal{K}} N_i$. Recall that $\Pr[N_i = 1] = (1 - p_s - p_d)^k/(d+1)^{k-1} = q$. The N_i's are not independent, but N_i and N_j are independent if $|i - j| \geq k$. Then, for $\ell \in \{1, \ldots, k\}$ let

$$M_\ell = \sum_{j\,:\,\ell+kj \leq L-k+1} N_{\ell+kj}.$$

Observe that the M_ℓ's are defined such that each is a sum of independent Bernoulli random variables with success probability q. The Chernoff bound in (14) implies that for $\delta \in (0, 1)$:

$$\Pr\Big[\big|M_\ell - \mathbf{E}[M_\ell]\big| \geq \delta\mathbf{E}[M_\ell]\Big] \;\leq\; 2\exp\Big\{-\frac{\delta^2\mathbf{E}[M_\ell]}{3}\Big\},$$

and this holds for $\ell \in \{1, 2, 3, \ldots, k\}$. By the triangle inequality:

$$\big|\mathcal{N} - \mathbf{E}[\mathcal{N}]\big| = \Big|\sum_{\ell=1}^{k} M_\ell - \mathbf{E}[M_\ell]\Big| \leq \sum_{\ell=1}^{k}\big|M_\ell - \mathbf{E}[M_\ell]\big|,$$

so if $\big|M_\ell - \mathbf{E}[M_\ell]\big| < \delta\mathbf{E}[M_\ell]$ for each ℓ, we obtain that $\big|\mathcal{N} - \mathbf{E}[\mathcal{N}]\big| < \delta\mathbf{E}[\mathcal{N}]$. A union bound then implies that

$$\Pr\Big[\big|\mathcal{N} - \mathbf{E}[\mathcal{N}]\big| \geq \delta\mathbf{E}[\mathcal{N}]\Big] \;\leq\; 2k\exp\Big\{-\frac{\delta^2\min_\ell \mathbf{E}[M_\ell]}{3}\Big\}.$$

Finally, observe that $\mathbf{E}[M_\ell] \geq \lfloor (L-k+1)/k \rfloor q \geq \mathbf{E}[\mathcal{N}] - q$. Since $e^{q\delta^2/3} \leq 1.5$, we obtain that

$$\Pr\Big[\big|\mathcal{N} - \mathbf{E}[\mathcal{N}]\big| \geq \delta\mathbf{E}[\mathcal{N}]\Big] \;\leq\; 3k\exp\Big\{-\frac{\delta^2\mathbf{E}[\mathcal{N}]}{3}\Big\},$$

as claimed. ◀

▶ **Lemma 6.** *For any $\delta \in (0, 1)$:*

$$\Pr\Big[\big|\mathcal{D} - \mathbf{E}[\mathcal{D}]\big| \geq \delta\mathbf{E}[\mathcal{D}]\Big] \;\leq\; 3k\exp\Big\{-\frac{\delta^2\mathbf{E}[\mathcal{D}]}{3k}\Big\}.$$

Proof. Observe that $\mathcal{D} = \sum_{i=1}^{L-k+1} D_i$, where each D_i is an indicator random variable such that

$$\Pr[D_i = 1] = k\Big(1 - p_s - p_d\Big)^{k-1} p_d \frac{1}{(d+1)^{k-1}} = \hat{q}.$$

In simpler terms, $D_i = 1$ indicates that the k-span $\mathcal{K}_i$ has been introduced with only a single deletion, and no other mutation during the mutation process. The D_i's have the same dependence structure as the N_i's in the proof of Lemma 5, so proceeding in exactly the same fashion (simply replacing q by $\hat{q}$), the result follows. ◀

▶ **Lemma 7.** *Let $J_1 = \ln 2/\min\{d+1, 8\}$. Then, there exist absolute constants $c_1, c_2 > 0$ such that the following holds for any $\delta > 0$:*

$$\Pr\Big[|P - \mathbf{E}[P]| \geq 3\delta\Big] \leq 2\exp\Big\{-\frac{\delta^2}{8f_A}\Big\} + 2\exp\Big\{-\frac{\delta^2}{8(L - f_A)}\Big\} + 2\exp\Big\{-c_1\min\Big(\frac{\delta^2}{c_2{}^2 J_1{}^2}, \frac{\delta}{c_2 J_1}\Big)\Big\}.$$

Proof. We can characterize $P = L' - 4f_A{}'$ by three separate groups of characters: the 'A' characters in S, the non-'A' characters in S, and the characters that are inserted into S by the mutation process.

First, we observe that an 'A' in S contributes -3 to P if it was untouched by the mutation process, contributes 1 to P if it was substituted into a non-'A' character by the mutation process, and contributes 0 to P if it was deleted by the mutation process. To capture these cases, we define X_i as follows:

$$X_i = \begin{cases} -3 & \text{with probability } 1 - p_s - p_d, \\ 1 & \text{with probability } p_s, \text{ and} \\ 0 & \text{with probability } p_d; \end{cases}$$

Next, we observe that a non-'A' character in S contributes 1 to P if it was untouched by the mutation process or substituted into another non-'A' character, contributes -3 to P if it is substituted into an 'A' character by the mutation process, and contributes 0 to P if it was deleted by the mutation process. To capture these cases, let

$$Y_i = \begin{cases} 1 & \text{with probability } 1 - p_s/3 - p_d, \\ -3 & \text{with probability } p_s/3, \text{ and} \\ 0 & \text{with probability } p_d; \end{cases}$$

Recall that C_i denotes the number of 'A's in Q_i, and $Z_i = l_i - 4C_i$. We can therefore express P in terms of the X_i's, Y_i's, and Z_i's as follows:

$$P = \sum_{i:S_i=\text{'A'}} X_i + \sum_{i:S_i\neq\text{'A'}} Y_i + \sum_{i=1}^{L} Z_i = X + Y + Z,$$

where for notational simplicity, we set $X = \sum_{i:S_i=\text{'A'}} X_i$, $Y = \sum_{i:S_i\neq\text{'A'}} Y_i$, and $Z = \sum_{i=1}^{L} Z_i$.

We first observe that

$$\mathbf{E}[X_i] = p_s - 3(1 - p_s - p_d) = -3 + 4p_s + 3p_d,$$
$$\mathbf{E}[Y_i] = 1 - \frac{p_s}{3} - p_d - p_s = 1 - \frac{4}{3}p_s - p_d, \quad \text{and}$$
$$\mathbf{E}[Z_i] = 0;$$

in the last equality we use the fact that $\mathbf{E}[l_i] = d$ and $\mathbf{E}[C_i] = d/4$; see (17). From these, we obtain

$$\mathbf{E}[X] = f_A(-3 + 4p_s + 3p_d),$$
$$\mathbf{E}[Y] = (L - f_A)\Big(1 - \frac{4}{3}p_s - p_d\Big),$$
$$\mathbf{E}[Z] = 0, \quad \text{and}$$
$$\mathbf{E}[P] = (L - 4f_A)\Big(1 - \frac{4}{3}p_s - p_d\Big).$$

By Hoeffding's inequality (see Lemma 11), for any $\delta > 0$:

$$\Pr(|X - \mathbf{E}[X]| \geq \delta) \leq 2\exp\left\{-\frac{\delta^2}{8f_A}\right\}. \tag{19}$$

and

$$\Pr(|Y - \mathbf{E}[Y]| \geq \delta) \leq 2\exp\Big\{-\frac{\delta^2}{8(L - f_A)}\Big\}. \tag{20}$$

To show that Z is strongly concentrated around its mean, we first prove that the Z_i's are sub-exponential random variables and then use Lemma 13.

▶ **Lemma 14.** *Let $J_1 = \ln 2/\min\{\frac{1}{d+1}, \frac{1}{8}\}$. The following holds for each Z_i for all $t \geq 0$:*

$$\Pr[|Z_i| \geq t] \leq 4\exp\Big\{-t/J_1\Big\}.$$

We prove this lemma shortly, but first we finish the proof of Lemma 7.

As Z_i's are mean-zero sub-exponential random variables, it follows from Bernstein's inequality (Lemma 13) that there exist absolute constants $c_1, c_2 > 0$ such that the following holds for any $\delta > 0$

$$\Pr\Big[|Z| \geq \delta\Big] \leq 2\exp\left\{-c_1 \min\left(\frac{\delta^2}{{c_2}^2 {J_1}^2}, \frac{\delta}{c_2 J_1}\right)\right\} \tag{21}$$

The triangle inequality and a union bound over (19), (20) and (21) completes the proof of Lemma 7. ◀

Proof of Lemma 14. Recall that $Z_i = l_i - 4C_i$, where l_i is a Geometric random variable with parameter $p = 1/(d+1)$ and mean d, and C_i is a Binomial random variable with parameters l_i and $1/4$. Let B_y be a Binomial random variable with parameters y and $1/4$, For $t \geq 0$, conditioning on the value of l_i and using Hoeffding's inequality, we obtain:

$$\begin{aligned}
\Pr[|Z_i| \geq t] &= \sum_{y\geq 0} \Pr[l_i = y]\Pr[|B_y - y/4| \geq t/4] \\
&\leq 2p\sum_{y\geq 0}(1-p)^y \exp\Big\{\frac{-t^2}{8y}\Big\} \\
&\leq 2p\left(\sum_{y=0}^{\lfloor t\rfloor}(1-p)^y\exp\Big\{-\frac{t^2}{8y}\Big\} + \sum_{y=\lfloor t\rfloor+1}^{\infty}(1-p)^y\exp\Big\{-\frac{t^2}{8y}\Big\}\right).
\end{aligned}$$

For $y \leq t$, we have $\frac{t^2}{8y} \geq \frac{t}{8}$, so $\exp\{-\frac{t^2}{8y}\} \leq \exp\{-t/8\}$ and we have

$$\sum_{y=0}^{\lfloor t\rfloor}(1-p)^y\exp\{-\frac{t^2}{8y}\} \leq \exp\{-t/8\}\sum_{y=0}^{\lfloor t\rfloor}(1-p)^y \leq \frac{e^{-t/8}}{p}.$$

When $y > t$,

$$\sum_{y=\lfloor t\rfloor+1}^{\infty}(1-p)^y\exp\{-\frac{t^2}{8y}\} \leq \sum_{y=\lfloor t\rfloor+1}^{\infty}(1-p)^y \leq \frac{(1-p)^t}{p} \leq \frac{e^{-pt}}{p}.$$

Putting these bounds together:

$$\Pr[|Z_i| \geq t] \leq 2\left(e^{-t/8} + e^{-pt}\right) \leq 4e^{-t/J_1}.$$

where $J_1 = \ln 2/\min\{p, 1/8\}$. ◀

▶ **Lemma 8.** *For any $\delta \in (0,1)$:*

$$\Pr\Big[\Big|L' - \mathbf{E}[L']\Big| \leq \delta(Ld + Lp_d)\Big] \;\geq\; 2\exp\left\{-\frac{L\delta^2 d^2}{2(d+1-\delta d)(d+1)}\right\} + 2\exp\left\{-\frac{Lp_d\delta^2}{3}\right\}.$$

Proof. We prove that L' is strongly concentrated around its expected value by proving concentration bounds for the individual pieces of (16). Specifically, we prove that $\sum l_i$ and $\sum D_i$ are strongly concentrated around their respective expected values.

For the sake of notational simplicity, let us define $\mathcal{L}_{\text{ins}} = \sum_{i=1}^{L} l_i$, and $\mathcal{L}_{\text{del}} = \sum_{i=1}^{L} D_i$. Then, L' is simply

$$L' = L + \mathcal{L}_{\text{ins}} - \mathcal{L}_{\text{del}}. \tag{22}$$

We first prove a strong concentration of $\mathcal{L}_{\text{ins}}$ around its expected value, and later prove the same result for $\mathcal{L}_{\text{del}}$. We finally use these results to show that L' is strongly concentrated around its mean.

We observe that every l_i is sampled from a geometric distribution with mean $d \geq 0$. This geometric distribution has support for values 0, 1, 2, 3, etc. (opposed to supporting positive values only). Another interpretation of this Geometric distribution is that it assigns probability to the number of failures before encountering a success – where each trial is a Bernoulli trial with success probability $1/(d+1)$. We also observe that the expected value of l_i is d, and therefore, the expected value of $\mathcal{L}_{\text{ins}}$ is Ld.

We aim to determine $\Pr[\mathcal{L}_{\text{ins}} \geq (1+\delta)Ld]$. We notice that $\mathcal{L}_{\text{ins}} \geq (1+\delta)Ld$ is equivalent to observing at least $(1+\delta)Ld$ failures before observing L successes from a series of Bernoulli trials with success probability $1/(d+1)$. Using the same argument, this is equivalent to observing at most L successes in $(1+\delta)Ld + L$ Bernoulli trials with success probability $1/(d+1)$. Let Z_i be a Bernoulli random variable with success probability $1/(d+1)$ for $i \geq 1$. Then, we have the following.

$$\Pr[\mathcal{L}_{\text{ins}} \geq (1+\delta)Ld] = \Pr[\sum_{i=1}^{(1+\delta)Ld+L} Z_i \leq L] \tag{23}$$

The expected number of success events in this collection of Bernoulli trials is $\mu = \frac{(1+\delta)Ld+L}{d+1}$. By reorganizing (23), we get the following result.

$$\Pr[\mathcal{L}_{\text{ins}} \geq (1+\delta)Ld] = \Pr[\sum_{i=1}^{(1+\delta)Ld+L} Z_i \leq (1-\epsilon)\mu] \tag{24}$$

where $\epsilon = \frac{\delta d}{\delta d+d+1}$. Using Chernoff inequality for the sum of independent Bernoulli random variables from (13), the probability in (24) is at most $\exp\{-\epsilon^2\mu/2\}$. Using the values of μ and ϵ as defined before, we have the following result.

$$\Pr[\mathcal{L}_{\text{ins}} \geq (1+\delta)Ld] \leq \exp\{-\frac{L\delta^2 d^2}{2(d+1+\delta d)(d+1)}\} \tag{25}$$

Proceeding in an analogous manner, we can prove the probability for the other tail.

$$\Pr[\mathcal{L}_{\text{ins}} \leq (1-\delta)Ld] \leq \exp\{-\frac{L\delta^2 d^2}{2(d+1-\delta d)(d+1)}\} \tag{26}$$

Taking a union bound using the probabilities in (25) and (26), we have the following.

$$\Pr[|\mathcal{L}_{\text{ins}} - Ld| \leq \delta Ld] \geq 1 - 2\exp\{-\frac{L\delta^2 d^2}{2(d+1-\delta d)(d+1)}\} \tag{27}$$

We next prove that $\mathcal{L}_{\text{del}}$ is strongly concentrated around its mean. We recall that $\mathcal{L}_{\text{del}} = \sum_{i=1}^{L} D_i$, where each D_i for $i = 1$ to L is a Bernoulli trial with success probability p_d. Here, all D_i's are independent and identically distributed, and $\mathbf{E}[\mathcal{L}_{\text{del}}] = Lp_d$. Therefore, we can use the Chernoff bound in (14) for $\mathcal{L}_{\text{del}}$, which gives us the following result.

$$\Pr[|\mathcal{L}_{\text{del}} - Lp_d| \le \delta L p_d] \quad \ge \quad 1 - 2\exp\{-\frac{Lp_d\delta^2}{3}\}. \tag{28}$$

Recall that $L' = L + \mathcal{L}_{\text{ins}} - \mathcal{L}_{\text{del}}$, as defined in (22). Therefore, by using the strong concentration of $\mathcal{L}_{\text{ins}}$ and $\mathcal{L}_{\text{del}}$ around their respective expected values in (27) and (28), the proof is completed. ◀

▶ **Theorem 9.** *Suppose $4f_A < L$ and $\frac{4}{3}p_s + p_d < 1$. Then, for sufficiently small $\delta > 0$, there exists constants $c_1, c_2 > 0$ such that*

$$\Pr[|\hat{p}_s - p_s| \ge 12\delta] \le 8k\exp\Big\{-\frac{\delta^2\mathbf{E}[\mathcal{N}]}{3k}\Big\} + 6k\exp\Big\{-\frac{\delta^2\mathbf{E}[\mathcal{D}]}{3k}\Big\} + 2\exp\Big\{-\frac{\delta^2\mathbf{E}[P]^2}{72f_A}\Big\}$$
$$+ 2\exp\Big\{-\frac{\delta^2\mathbf{E}[P]^2}{72(L-f_A)}\Big\} + 2\exp\Big\{-c_1\min\Big(\frac{\delta^2\mathbf{E}[P]^2}{{c_2}^2}, \frac{\delta\mathbf{E}[P]}{c_2}\Big)\Big\}.$$

Proof. Assume $4f_A < L$ and that $\frac{4}{3}p_s + p_d \le 1$. Recall that $P = L' - 4f'_A$ and let $Q = k\mathcal{N} + \mathcal{D}$, $R = k(L - 4f_A)\mathcal{N}$, and $T = 4k\mathcal{N} + \mathcal{D}$. Observe that

$$p_s = -3 \cdot \frac{\mathbf{E}[P]\mathbf{E}[Q] - \mathbf{E}[R]}{(L-4f_A)\mathbf{E}[T]}, \text{ and } \quad \hat{p}_s = -3 \cdot \frac{PQ - R}{(L-4f_A)T};$$

our proof will derive a concentration bound for $\hat{p}_s$ from those for the random variables P, Q, R, and T. First, by Lemma 5

$$\Pr\left[|\mathcal{N} - \mathbf{E}[\mathcal{N}]| \ge \delta\mathbf{E}[\mathcal{N}]\right] \le 2k\exp\Big\{-\frac{\delta^2\mathbf{E}[\mathcal{N}]}{3k}\Big\}.$$

So, since $4f_A < L$, we can multiply the inequality inside the probability by $k(L - 4f_A)$ to obtain that:

$$\Pr\left[|R - \mathbf{E}[R]| \ge \delta\mathbf{E}[R]\right] \le 2k\exp\Big\{-\frac{\delta^2\mathbf{E}[\mathcal{N}]}{3k}\Big\}.$$

From Lemmas 5 and 6, the triangle inequality and a union bound, we have that for any $\delta \ge 0$

$$\Pr\left[|Q - \mathbf{E}[Q]| \ge \delta\mathbf{E}[Q]\right] \le 3k\exp\Big\{-\frac{\delta^2\mathbf{E}[\mathcal{N}]}{3k}\Big\} + 3k\exp\Big\{-\frac{\delta^2\mathbf{E}[\mathcal{D}]}{3k}\Big\},$$
$$\Pr\left[|T - \mathbf{E}[T]| \ge \delta\mathbf{E}[T]\right] \le 3k\exp\Big\{-\frac{\delta^2\mathbf{E}[\mathcal{N}]}{3k}\Big\} + 3k\exp\Big\{-\frac{\delta^2\mathbf{E}[\mathcal{D}]}{3k}\Big\}.$$

Moreover, since under the assumption that $\frac{4}{3}p_s - p_d < 1$

$$\mathbf{E}[P] = (L - 4f_A)\Big(1 - \frac{4}{3}p_s - p_d\Big) > 0,$$

Lemma 7 implies there exists constants $c_1 > 0$ and $c_2 = c_2(d) > 0$ such that

$$\Pr\Big[|P - \mathbf{E}[P]| \ge \delta\mathbf{E}[P]\Big] \le 2\exp\Big\{-\frac{\delta^2\mathbf{E}[P]^2}{72f_A}\Big\} + 2\exp\Big\{-\frac{\delta^2\mathbf{E}[P]^2}{72(L-f_A)}\Big\} + 2\exp\Big\{-c_1\min\Big(\frac{\delta^2\mathbf{E}[P]^2}{{c_2}^2}, \frac{\delta\mathbf{E}[P]}{c_2}\Big)\Big\}.$$

Now, let us assume that the concentration bounds above do hold and that indeed:

$$(1-\delta)\,\mathbf{E}[P] \le P \le (1+\delta)\,\mathbf{E}[P], \tag{29}$$
$$(1-\delta)\,\mathbf{E}[Q] \le Q \le (1+\delta)\,\mathbf{E}[Q],$$
$$(1-\delta)\,\mathbf{E}[R] \le R \le (1+\delta)\,\mathbf{E}[R], \tag{30}$$
$$(1-\delta)\,\mathbf{E}[T] \le T \le (1+\delta)\,\mathbf{E}[T]. \tag{31}$$

Note that $\mathbf{E}[P], \mathbf{E}[Q], \mathbf{E}[R], \mathbf{E}[T] > 0$. Then,

$$\frac{(1-\delta)^2}{1+\delta}\frac{\mathbf{E}[P]\mathbf{E}[Q]}{\mathbf{E}[T]} \leq \frac{PQ}{T} \leq \frac{(1+\delta)^2}{1-\delta}\frac{\mathbf{E}[P]\mathbf{E}[Q]}{\mathbf{E}[T]}.$$

Observe that $\frac{(1-\delta)^2}{1+\delta} \geq 1-4\delta$ and $\frac{(1+\delta)^2}{1-\delta} \leq 1+4\delta$ for small enough δ. Then, we have the following:

$$(1-4\delta)\frac{\mathbf{E}[P]\mathbf{E}[Q]}{\mathbf{E}[T]} \leq \frac{PQ}{T} \leq (1+4\delta)\frac{\mathbf{E}[P]\mathbf{E}[Q]}{\mathbf{E}[T]}. \tag{32}$$

Similarly, by dividing (30) by (31), and noting that that $\frac{1-\delta}{1+\delta} \geq 1-3\delta$, and $\frac{1+\delta}{1-\delta} \leq 1+3\delta$ for small enough δ, we obtain:

$$(1-3\delta)\frac{\mathbf{E}[R]}{\mathbf{E}[T]} \leq \frac{1-\delta}{1+\delta}\cdot\frac{\mathbf{E}[R]}{\mathbf{E}[T]} \leq \frac{R}{T} \leq \frac{1+\delta}{1-\delta}\cdot\frac{\mathbf{E}[R]}{\mathbf{E}[T]} \leq (1+3\delta)\frac{\mathbf{E}[R]}{\mathbf{E}[T]}.$$

Combined with (32) this yields that

$$\frac{PQ-R}{T} \leq (1+4\delta)\frac{\mathbf{E}[P]\mathbf{E}[Q]}{\mathbf{E}[T]} - (1-3\delta)\frac{\mathbf{E}[R]}{\mathbf{E}[T]} \quad \leq \frac{\mathbf{E}[P]\mathbf{E}[Q]-\mathbf{E}[R]}{\mathbf{E}[T]} + 4\delta\frac{\mathbf{E}[P]\mathbf{E}[Q]+\mathbf{E}[R]}{\mathbf{E}[T]}$$

and

$$\frac{PQ-R}{T} \geq \frac{\mathbf{E}[P]\mathbf{E}[Q]-\mathbf{E}[R]}{\mathbf{E}[T]} - 4\delta\frac{\mathbf{E}[P]\mathbf{E}[Q]+\mathbf{E}[R]}{\mathbf{E}[T]}.$$

Multiplying by $\frac{-3}{L-4f_A}$ we obtain that $p_s - 4\delta W \leq \hat{p}_s \leq p_s + 4\delta W$ with $W = 3\frac{\mathbf{E}[P]\mathbf{E}[Q]+\mathbf{E}[R]}{(L-4f_A)E[T]}$. It can be checked that $W \leq 3$, and the result then follows from a union bound so that (29)-(31) all hold. ◀

E Interesting but unused results

Herein, we describe some interesting observations and results that were uncovered in the process of finding our main results, but which ultimately were not used in any proofs or the like. These include additional statistics, concentration inequalities, and an alternate, polynomial-based mutation parameter estimation approach.

E.1 Estimating the mutation rates using a polynomial

Unlike the estimators shown in Section 4, we can estimate the mutation rates by using a different combination of observations. In this section, we assume that $\mathbf{E}[\mathcal{S}]$, $\mathbf{E}[\mathcal{D}]$, and $\mathbf{E}[\mathcal{I}]$ are all positive. This is equivalent to assuming that p_s, p_d, and d are all non-zero. We later handle the cases for $p_s = 0$, and/or $p_d = 0$, and/or $d = 0$ separately.

For the ease of algebraic manipulations, we first define the following.

$$\mathcal{S}_{\text{norm}} = \frac{\mathcal{S}}{k\,\mathcal{K}}$$
$$\mathcal{I}_{\text{norm}} = \frac{\mathcal{I}}{(k-1)\mathcal{K}}$$
$$\mathcal{D}_{\text{norm}} = \frac{\mathcal{D}}{k\,\mathcal{K}}$$

We first take the ratio of (3) and (4) to get the following.

$$p_s = \frac{\mathbf{E}[\mathcal{S}_{\text{norm}}]}{\mathbf{E}[\mathcal{D}_{\text{norm}}]} p_d$$

We also take the ratio of (4) and (5) to derive the following.

$$\frac{\mathbf{E}[\mathcal{D}_{\text{norm}}]}{\mathbf{E}[\mathcal{I}_{\text{norm}}]} = \frac{d+1}{d} \times \frac{p_d}{1 - \mathbf{E}[\mathcal{S}_{\text{norm}}]p_d/\mathbf{E}[\mathcal{D}_{\text{norm}}] - p_d}$$

By reorganizing, we get the following.

$$d + 1 = \frac{\mathbf{E}[\mathcal{D}_{\text{norm}}] - (\mathbf{E}[\mathcal{S}_{\text{norm}}] + \mathbf{E}[\mathcal{D}_{\text{norm}}])p_d}{\mathbf{E}[\mathcal{D}_{\text{norm}}] - (\mathbf{E}[\mathcal{S}_{\text{norm}}] + \mathbf{E}[\mathcal{D}_{\text{norm}}] + \mathbf{E}[\mathcal{I}_{\text{norm}}])p_d}$$

We next plug the values of $d + 1$ and p_s into (4), and get the following polynomial after reorganizing.

$$\mathbf{E}[\mathcal{D}_{\text{norm}}]^k = p_d\Big(\mathbf{E}[\mathcal{D}_{\text{norm}}] - (\mathbf{E}[\mathcal{S}_{\text{norm}}] + \mathbf{E}[\mathcal{D}_{\text{norm}}] + \mathbf{E}[\mathcal{I}_{\text{norm}}])p_d\Big)^{k-1} \tag{33}$$

Given S and S', we can only observe $\mathcal{I}_{\text{norm}}$, $\mathcal{D}_{\text{norm}}$, and $\mathcal{I}_{\text{norm}}$. Using these observations, we can estimate p_d from (33).

$$\mathcal{D}_{\text{norm}}{}^k = \hat{p}_d\Big(\mathcal{D}_{\text{norm}} - (\mathcal{S}_{\text{norm}} + \mathcal{D}_{\text{norm}} + \mathcal{I}_{\text{norm}})\hat{p}_d\Big)^{k-1}$$

Let $y = \mathcal{D}_{\text{norm}} - (\mathcal{S}_{\text{norm}} + \mathcal{D}_{\text{norm}} + \mathcal{I}_{\text{norm}})\hat{p}_d$. This results in the following polynomial:

$$y^k - \mathcal{D}_{\text{norm}}\, y^{k-1} + (\mathcal{S}_{\text{norm}} + \mathcal{D}_{\text{norm}} + \mathcal{I}_{\text{norm}})\mathcal{D}_{\text{norm}}{}^k = 0,$$

where

$$\hat{p_d} = \frac{\mathcal{D}_{\text{norm}} - y}{\mathcal{S}_{\text{norm}} + \mathcal{D}_{\text{norm}} + \mathcal{I}_{\text{norm}}}.$$

After solving for $\hat{p_d}$, we estimate the other two rates as follows.

$$\hat{p_s} = \frac{\mathcal{S}_{\text{norm}}}{\mathcal{D}_{\text{norm}}}\hat{p_d}$$

$$\hat{d} = \frac{(\mathcal{D}_{\text{norm}})\hat{p_d}}{\mathcal{D}_{\text{norm}} - (\mathcal{S}_{\text{norm}} + \mathcal{D}_{\text{norm}} + \mathcal{I}_{\text{norm}})\hat{p_d}}$$

We next handle the edge cases in which these equations do not work. Note that we assume that $1 - p_s - p_d > 0$. From (4), we see that $p_d = 0$ if and only if $\mathbf{E}[\mathcal{D}] = 0$. Therefore, if we observe $\mathcal{D} = 0$, we estimate $\hat{p_d} = 0$. Using a similar argument, if we observe $\mathcal{S} = 0$, then we estimate $\hat{p_s} = 0$. And finally, if we observe $\mathcal{I} = 0$, then we estimate $\hat{d} = 0$.

E.2 Concentration of $f_A{}'$ around $\mathbf{E}[f_A{}']$

We next prove that $f_A{}'$ is strongly concentrated around its expected value. Recall the definition of $f_A{}'$ defined in (18). For the sake of notational simplicity, we define the following.

$$\mathcal{A} = \sum_{i:S_i=\text{'A'}} A_i\ , \tag{34}$$

$$\mathcal{B} = \sum_{i:S_i\neq\text{'A'}} B_i\ , \text{and}$$

$$C = \sum_{i=1}^{L} C_i,$$

where A_i, B_i, and C_i are the same as introduced in Section 3. Using these notations, $f_A{}'$ is simply

$$f_A{}' = \mathcal{A} + \mathcal{B} + C.$$

E.2.1 Strong concentration of $\mathcal{A}$ around $\mathbf{E}[\mathcal{A}]$

We note that the A_i's in (34) are iid Bernoulli random variables with success probability $1 - p_s - p_d$. We also note that number of 'A' characters in S is f_A. Therefore, we have the following.

$$\mathbf{E}[\mathcal{A}] = f_A(1 - p_s - p_d)$$

As the $\mathcal{A}$ is a sum of iid Bernoulli random variables, we can use the Chernoff bound in (14) for $\mathcal{A}$, which gives us the following result.

$$\Pr\left[\left|\mathcal{A} - \mathbf{E}[\mathcal{A}]\right| \leq \delta f_A(1 - p_s - p_d)\right] \ \geq\ 1 - 2\exp\{-\frac{\delta^2 f_A(1 - p_s - p_d)}{3}\} \tag{35}$$

E.2.2 Strong concentration of $\mathcal{B}$ around $\mathbf{E}[\mathcal{B}]$

We next prove a concentration bound for the sum of B_i's, as defined in (18). We note that B_i's here are iid Bernoulli random variables, each with success probability of $p_s/3$. We also note that the number of characters in S that are not 'A' is $L - f_A$. Therefore, we have the following result.

$$\mathbf{E}[\mathcal{B}] = (L - f_A)\frac{p_s}{3}$$

As $\mathcal{B}$ is simply a sum of iid Bernoulli random variables, we can use the same Chernoff inequality in (14) for $\mathcal{B}$, which gives us the following result.

$$\Pr\left[\left|\mathcal{B} - \mathbf{E}[\mathcal{B}]\right| \leq \delta(L - f_A)\frac{p_s}{3}\right] \geq 1 - 2\exp\{-\frac{\delta^2(L - f_A)p_s}{9}\} \tag{36}$$

E.2.3 Strong concentration of C around $\mathbf{E}[C]$

We finally prove a concentration inequality for $\sum_{i=1}^{L} C_i$, as defined in (18). Recall the result in (27) – developed for a sum of iid geometric random variables, each with mean d. We prove a similar result for C by showing that each C_i for $i = 1$ to L is geometrically distributed with mean $d/4$.

We first recall that the probability mass function (PMF) for a geometric distribution with parameter p. If a random variable X follows a geometric distribution with probability parameter p where $0 < p \leq 1$, then the PMF is as follows:

$$\Pr[X = j] = p(1 - p)^j,$$

where $j = 0, 1, 2, 3, \cdots$, etc. The expected value of X is given by

$$\mathbf{E}[X] = \frac{1 - p}{p}.$$

We can express the PMF of X using its mean. If the mean is m, then we need $p = 1/(m + 1)$, and therefore, the PMF of a geometric random variable with mean m is given by the following:

$$\Pr[X = j] = \frac{m^j}{(m + 1)^{j+1}}, \tag{37}$$

where $j = 0, 1, 2, 3, \cdots$, etc.

We next observe that C_i is the number of 'A' characters in Q_i, where Q_i is the string inserted after S_i and before S_{i+1}. The number of characters in Q_i is l_i, which follows a geometric distribution with mean d, and each character can be an 'A' with probability 1/4. Note that it immediately follows that $\mathbf{E}[C_i] = d/4$. We can therefore compute the probability that $C_i = j$ as follows:

$$\begin{aligned}
\Pr[C_i = j] &= \sum_{x=j}^{\infty} \Pr[C_i = j \mid l_i = x] \Pr[l_i = x] \\
&= \sum_{x=j}^{\infty} \binom{x}{j} \left(\frac{1}{4}\right)^j \left(\frac{3}{4}\right)^{x-j} \Pr[l_i = x] \\
&= \sum_{x=j}^{\infty} \binom{x}{j} \left(\frac{1}{4}\right)^j \left(\frac{3}{4}\right)^{x-j} \frac{d^x}{(d+1)^{x+1}} \\
&= \left(\frac{1}{4}\right)^j \left(\frac{4}{3}\right)^j \frac{1}{d+1} \sum_{x=j}^{\infty} \binom{x}{j} \left(\frac{3}{4}\right)^x \frac{d^x}{(d+1)^x} \\
&= \left(\frac{1}{4}\right)^j \left(\frac{4}{3}\right)^j \frac{1}{d+1} \sum_{x=j}^{\infty} \binom{x}{j} \left(\frac{3d}{4d+4}\right)^x \qquad (38)
\end{aligned}$$

We can now use the following summation for Binoamial coefficients:

$$\sum_{x=0}^{\infty} \binom{x}{j} r^x = \frac{r^j}{(1-r)^{j+1}},$$

where $|r| < 1$. In our case, $r = 3d/(4d+4)$. Our mutation process assumes $d \geq 0$, which means $4d + 4 > 3d$, and therefore, this summation can be used to simplify (38).

$$\begin{aligned}
\Pr[C_i = j] &= \left(\frac{1}{4}\right)^j \left(\frac{4}{3}\right)^j \frac{1}{d+1} \left(\frac{3d}{4d+4}\right)^j \frac{1}{\left(1 - \frac{3d}{4d+4}\right)^{j+1}} \\
&= \left(\frac{1}{4}\right)^j \left(\frac{4}{3}\right)^j \frac{1}{d+1} \times \frac{3^j d^j}{4^j (d+1)^j} \times \frac{4^{j+1}(d+1)^{j+1}}{(d+4)^{j+1}} \\
&= \frac{4d^j}{(d+1)^{j+1}} \\
&= \frac{(d/4)^j}{\left(d/4+1\right)^{j+1}}
\end{aligned}$$

This is the same form as the PMF of a geometric distribution shown in (37). Since $\mathbf{E}[C_i] = d/4$, it follows that each C_i is a geometric random variable with mean $d/4$.

We can now use the result proved for sum of geometric random variables with mean d, as shown in (27). Note that $C = \sum_{i=1}^{L} C_i$, which means $\mathbf{E}[C] = Ld/4$, where each C_i follows a geometric distribution with mean $d/4$. Using the result in (27) gives us the following:

$$\Pr\left[\left|C - \frac{Ld}{4}\right| \leq \delta \frac{Ld}{4}\right] \geq 1 - 2\exp\left\{-\frac{L\delta^2 d^2}{2(d+4-\delta d)(d+4)}\right\}. \qquad (39)$$

E.2.4 Strong concentration of $f_A{}'$ around $\mathbf{E}[f_A{}']$

We can now use the concentration bounds in (35), (36), and (39) to take a union bound and prove the following.

$$\Pr\left[\left|f_A' - \mathbf{E}[f_A']\right| \leq \delta \mathbf{E}[f_A']\right] \leq$$
$$1 - 2\exp\left\{-\frac{\delta^2 f_A(1-p_s-p_d)}{3}\right\} - 2\exp\left\{-\frac{\delta^2(L-f_A)p_s}{9}\right\} - 2\exp\left\{-\frac{L\delta^2 d^2}{2(d+4-\delta d)(d+4)}\right\}.$$