Estimating repeat spectra and genome length from low-coverage genome skims with RESPECT

Shahab Sarmashghi¹, Metin Balaban², Eleonora Rachtman², Behrouz Touri¹, Siavash Mirarab¹, and Vineet Bafna³

¹Department of Electrical & Computer Engineering, University of California, San Diego, La Jolla, CA, USA

²Bioinformatics & Systems Biology Graduate Program, University of California, San Diego, La Jolla, CA, USA

³Department of Computer Science & Engineering, University of California, San Diego, La Jolla, CA, USA

January 28, 2021

Abstract

The cost of sequencing the genome is dropping at a much faster rate compared to assembling and finishing the genome. The use of lightly sampled genomes (genome-skims) could be transformative for genomic ecology, and results using k -mers have shown the advantage of this approach in identification and phylogenetic placement of eukaryotic species. Here, we revisit the basic question of estimating genomic parameters such as genome length, coverage, and repeat structure, focusing specifically on estimating the k -mer repeat spectrum. We show using a mix of theoretical and empirical analysis that there are fundamental limitations to estimating the k -mer spectra due to ill-conditioned systems, and that has implications for other genomic parameters. We get around this problem using a novel constrained optimization approach (Spline Linear Programming), where the constraints are learned empirically. On reads simulated at 1X coverage from 66 genomes, our method, REPEAT SPECTRA Estimation (RESPECT), had $< 1.5\%$ error in length estimation compared to 34% error previously achieved. In shotgun sequenced read samples with contaminants, RESPECT length estimates had median error 4%, in contrast to other methods that had median error 80%. Together, the results suggest that low-pass genomic sequencing can yield reliable estimates of the length and repeat content of the genome. The RESPECT software will be publicly available at <https://github.com/shahab-sarmashghi/RESPECT.git>

Introduction

Anthropogenic pressure and other natural causes have resulted in severe disruption of global ecosystems in recent years, including loss of biodiversity¹. In North America alone, the bird population has declined by over a quarter since 1970². Simply understanding the scope and extent of biodiversity changes remains a challenging problem. Genomic sequence based biodiversity sampling provides an attractive alternative to physical sampling and cataloging, as falling costs have made it possible to shotgun sequence a reference specimen sample for at most \$10 per Gb (with another \$60 for sample prep). *However, the analysis typically requires assembling and finishing a reference genome, which can still be prohibitively costly.* Despite the many projects aimed at high quality genome sequencing of eukaryotic species³, it could be many decades before we have acquired high-quality data so that biodiversity measurements for each population can be acquired on an ongoing, routine basis.

While (meta)barcoding^{4–6} methods can be used for species identification and biodiversity measurements, they have many drawbacks including limited phylogenetic resolution^{7,8}. Organelle assembly based methods^{9–11} similarly cannot be used for populations and often require whole genome sequences but discard the nuclear reads (the vast majority of data). Therefore, there is renewed interest in the development of methods that use all nuclear DNA from *genome-skims*—low-coverage (0.5–2Gb) sequencing, providing 0.2–4 $\times$ coverage¹². The low coverage of skims makes them cost-effective, but insufficient for assembling, and calls for assembly-free methods. Such methods, based on analysis of k -mers are being actively developed¹³, and have been used for species identification (Skmer¹⁴); for phylogenetic placement of a new species not in the library (APPLES¹⁵), and contaminant filtering (unpublished). While k -mer analysis works well for species identification, it cannot be applied easily for the analysis of populations (individuals from the same species) using genome-skims, a key component of genomic ecology. Specifically, it ignores the effect of repeats, and uses heuristics to estimate sequencing error and coverage, neither of which is known.

In this manuscript, we revisit the problem of estimating genomic parameters from genome-skin data: specifically, genome length L , sequence-coverage c , and repeat content. From genome-skin data, we have as input, abundance values of k -mers denoted by $\mathbf{o}$, where o_h denotes the number of distinct k -mers of multiplicity h . A key latent variable is the k -mer-repeat-spectrum (denoted hereafter as the k -mer spectrum) of the genome described by $\mathbf{r}$, where r_h denotes the number of distinct k -mers that appear exactly h times in the genome. As the value of o_h depends upon $\mathbf{r}$, c , L , and also on sequencing error, we consider the inverse problem of estimating genomic parameters given $\mathbf{o}$ as input. The problem was studied in a seminal paper by Li and Waterman¹⁶ who mostly considered the case of high coverage and no sequencing errors. Williams et al¹⁷ improved upon this model by ignoring o_1 assuming that a large proportion of unique k -mers can be attributed to sequencing errors. This assumption works better for high coverage because at low coverage, many informative k -mers are also seen only once. Hozza *et al*¹⁸ point this out, and focus attention on k -mer spectra. Their method, CovEst, models spectra using a geometric distribution of unknown parameters, uses that parameterized model to estimate both parameters and r_1, r_2, r_3 , and improves

estimates even for low coverage and high error.

A distinct but related line of research relates to estimating $\mathbf{o}$ itself by sub-sampling or streaming reads. Melsted and colleagues^{19,20} describe streaming algorithms to estimate o_1 as well as moments $F_k = \sum_i i^k o_i$. Interestingly, these moments can also be used to estimate genome parameters. For example, $\mathbb{E}[F_1] = \lambda L$, where $\lambda = (1 - (k - 1)/\ell)c$ denotes the k -mer coverage, or the average number of k -mers covering a position derived from reads of length ℓ . We note that streaming is akin to low-coverage sampling and consider the case of estimating parameters over a range of λ .

Estimating genome repetitiveness and other parameters using k-mers

While previous research has emphasized the estimation of genome length and coverage, we focus specifically on estimating the k -mer spectrum $\mathbf{r}$, defined below. Consider a genome of length L . Decompose the genome into a collection of all fixed-length (overlapping) sequences of length k , called k -mers. Let variables r_j ($j \geq 1$) denote the number of k -mers that occur exactly j times in the genome. When k is large enough ($k \geq \log_4 L$), high values of r_j , for $j \geq 2$, can be attributed to the repetitive structures in the genome rather than chance similarities. Therefore, we define $\mathbf{r} = [r_1, r_2, \dots]$ as the $(k\text{-mer})\text{-repeat-spectrum}$ of the genome.

While the repetitive sequences occur in a variety of arrangements in terms of their multiplicity, complexity and the size of repeating unit, the repeat spectrum provides a valuable summary of the extent of repetition in the genome as well as other parameters. For example, the genome length can be estimated as $L = k - 1 + \sum_j j r_j \simeq \sum_j j r_j$. Define the *uniqueness ratio* of a genome as r_1/L , or the ratio of the number of k -mers seen only once to the genome length (which is the total number of k -mers in the genome). We computed the uniqueness ratio for 622 eukaryotic genomes in RefSeq using $k = 31$. The ratio revealed a broad spectrum of values, ranging from 0.287 for *A. tauschii* (Tausch's goatgrass) to 0.995 for a mite species, *V. jacobsoni* (Fig. 1a). Expectedly, there is some phylogenetic correlation and the variation of uniqueness ratio within a genus (intra-generic) is significantly lower than inter-generic variation of uniqueness ratios (Fig. S2). At higher taxonomic ranks, we observed that plants had a significantly lower uniqueness ratio compared to other groups (Fig. 1b), consistent with a prevalence of whole genome duplication (WGD) events (see Methods). Nevertheless, the correlation is not strong enough to predict uniqueness ratios solely from taxonomy. For example, rice species *O. sativa* and *O. brachyantha* have different ratios 0.91 and 0.75, respectively.

The repeat spectrum provides other insights. In genomes composed largely of unique sequences, $r_1/L \simeq 1$ and r_j values decrease rapidly for $j \geq 2$ with $\log r_1/r_5 \geq 4.5$ (Fig. 1c). On the other hand, genomes with higher repetitive content have a smoother decrease of r_j values (Fig. 1d) with $\log \frac{r_1}{r_5} \leq 2.5$ (Supplementary Fig. S3). Additionally, a genome that has duplicated very recently will have $r_1 \simeq 0$ and a very high value of r_2 . Over time, however, r_1 increases due to the accumulation of mutations. Similarly, $r_j > 0$ for large values of j suggest the presence of interspersed repeats.

Our method RESPECT (Repeat Spectrum identification) derives genomic length and coverage from low-coverage genome skims, while also providing insight into the repeat structure. We showed,

through a mix of theoretical reasoning and empirical evidence, that the k -mer repeat spectra estimation problem is fundamentally difficult because of severe ill-conditioning of the system. In fact, the spectra are hard to estimate even when the coverage and sequencing error rate are known. We resolve this problem for the case of known coverage and sequencing error by imposing constraints on r_h and solving a constrained optimization problem. This approach provides greatly improved estimates of $\mathbf{r}$, which in turn lead to even better estimation of coverage, genome length and sequencing error through a stochastic iteration method. Results on genomes sampled from different parts of the tree of life and with differing repeat structures illustrate the validity of our approach. RESPECT is available at <https://github.com/shahab-sarmashghi/RESPECT.git>

Results

A simple model for estimating repeat spectra from unassembled data performs poorly

Assume that reads in the genome-skim are sequenced with a fixed mean error rate of ϵ per bp, and that the read start positions follow a Poisson distribution with a mean coverage of λ per bp. Denote the *observed* k -mer data as the vector $\mathbf{o} = [o_1, o_2, \dots]$, where o_h denotes the number of k -mers observed exactly h times in the genome-skim input. The value o_h is the outcome of a random variable O_h that depends upon the parameter set $\Phi = \{\lambda, \epsilon, \mathbf{r}\}$ (See Methods: ‘Modeling genomic parameters’). Specifically, we assume that each k -mer with copy number j in the genome is sampled h times according to a Poisson distribution with rate dependent upon k, Φ . Let P_{hj} represent the probability of h observances of a k -mer with copy number j . Then, in expectation,

$$\mathbb{E}[\mathbf{O}] = \mathbf{r}\mathbf{P}^T + \mathbf{1}_{h=1}E \quad (1)$$

where E is the expected number of erroneous k -mers that in turn depends upon Φ . Φ could be estimated using:

$$\Phi = \arg \min_{\Phi} \|\mathbf{o} - \mathbb{E}[\mathbf{O}]\| = \arg \min_{\Phi} \|\mathbf{o} - (\mathbf{r}\mathbf{P}^T + \mathbf{1}_{h=1}E)\| \quad (2)$$

In principle, an iterative procedure could be used to solve the optimization; we start with initial estimates of λ and ϵ , and use them to compute $\mathbf{P}$ and E . Then, we can use the least-square (LS) method to find $\mathbf{r}$ which minimizes $\|\mathbf{o} - (\mathbf{r}\mathbf{P}^T + \mathbf{1}_{h=1}E)\|$ (Eqn. 2) (See Methods: ‘Least-squares estimate of repeat spectrum’).

To study the accuracy of this model for repeat spectra estimation, we simulated genome skims at 1X coverage with no sequencing errors ($E = 0$) for all 622 genomes in RefSeq in four major taxonomic groups of eukaryotes. A subset of 66 species was selected as the test set. The test genomes were sampled such that their uniqueness-ratio (r_1/L) values matched the distribution of uniqueness-ratios of all 622 RefSeq genomes (Supplementary Fig. S4, see Methods: ‘Comparing r_1/L distribution over different sets’). In the following text, all parameters were trained on the 556

Figure 2: **Repeat spectra estimation** (a) The relative error in estimating repeat spectra using Least-Squares (LS), constrained Linear Programming (LP), and Spline Linear Programming (SLP). The genome-skims are simulated at 1X with no sequencing error. (b) Correlation between true r_2/r_1 ratios, and our estimates of $r_1/\sum_{i=1} r_i$ for each genome. (c) Similar correlation plot between true r_3/r_2 and estimated $r_2/\sum_{i=2} r_i$. In both (b) and (c), true spectral ratios on Y axis are computed from the assemblies, and the estimated indices on X axis are obtained by applying the LP method to the simulated skims described in (a).

the estimation set $r_j = 0$ for many j 's, contrary to its true value in the genome.

Empirical and theoretical results showed that the poor performance could be attributed to severe ill-conditioning. We proved that the condition number of $\mathbf{P}$ grows exponentially with the number of spectra (see Supplementary Methods). Therefore, small changes in $\mathbf{o}$ relative to $\mathbb{E}[\mathbf{O}]$ (Eqn. 1), for example due to the sampling variability or the simplifying assumptions of model, led to very large errors in estimates of $\mathbf{r}$.

Overview of RESPECT algorithm.

The negative result suggested a fundamental limitation to the use of k -mer based methods for estimating repeat spectra. Regularization is a proposed remedy for ill-conditioned matrices. However, most regularization methods enforce sparsity and $\mathbf{r}$ is known to be not sparse. A second challenge is that both observed counts and k -mer spectra are very skewed towards lower indices. Thus, a small (even 1%) relative error in r_1 could lead to a larger error in r_j for $j > 1$. To get around the ill-conditioning problem, we focused on *constraining* possible values of $\mathbf{r}$. We observed empirically that ratio of consecutive spectral values r_{j+1}/r_j was tightly constrained. Fig. 2b traces r_2/r_1 as a function of $\frac{r_1}{\sum_{i \geq 1} r_i}$ on the training data and notes the tight correlation across all taxonomic groups. A similar, albeit less tight, constraint was observed for r_3/r_2 (Fig. 2c) and other values as well (Supplementary Figures S5 to S7).

These ideas provided the basis of a constrained linear-program for estimating $\mathbf{r}$. As a first step, we added the constraint that $\mathcal{L}_j \leq \frac{r_j}{r_{j+1}} \leq \mathcal{U}_j$ for each j , where $\mathcal{L}_j$ and $\mathcal{U}_j$ are the smallest and the largest $\frac{r_j}{r_{j+1}}$ ratios over the training genomes, and solved the following LP to find $\mathbf{r}$ (see Methods: ‘Linear programming for constrained optimization based estimates’)

$$\mathbf{r} = \arg \min_{\mathbf{r}} \mathcal{E} = \arg \min_{\mathbf{r}} \sum_{h=2}^n \left| o_h - \sum_{j=1}^n P_{hj} r_j \right| \quad (3)$$

This approach significantly improved the average error in estimating the spectra at multiplicity $j = 3$ and higher (Fig. 2a; LP method), and resulted in small improvement at $j = 1, 2$ as well.

Using the repeat spectra from 556 training genomes, we observed a strong correlation between r_2/r_1 and $r_1/\sum_{i \geq 1} r_i$ (Fig. 2b). Therefore, we estimated r_2/r_1 by using the LP estimate of $r_1/\sum_{i \geq 1} r_i$ and a spline fitted on the training data based on a generalized additive model^{21,22} (see Methods: ‘Spline Linear programming’). The estimated r_2/r_1 value and the LP estimated r_1 value provided a new estimate (named SLP) of r_2 . In a similar fashion, we computed SLP estimates of r_{j+1} from LP estimate of r_j and $r_j/\sum_{i \geq j} r_i$ for $j = 2, 3, 4, 5$ (Fig. 2c, Supplementary Figures S5 to S7, and Methods: ‘Spline Linear programming’). Using the additional information learned from the training genomes captured by the fitted splines, we obtained significant reduction in the average error of repeat spectra estimation (SLP vs. LP in Fig. 2a). To solve the full optimization problem in Eqn. 2, we used a simulated annealing procedure. Specifically, starting with initial estimates of parameters obtained under no-repeat assumption, at each iteration a new values for λ is suggested, and SLP method is used to estimate $\mathbf{r}$. If a candidate λ results in a reduction in error, the algorithm accepts the move. Moreover, to avoid getting stuck at local minima, occasionally moves to states with higher error are also accepted. Lastly, the initial estimate of ϵ is corrected for the repetitiveness of genome using a regression learned over a subset of training genomes (Supplementary Fig. S8). The algorithm is outlined below (also see Methods: ‘RESPECT algorithm’ for a detailed description).

1. Generate initial estimates of λ , ϵ , and $\mathbf{r}$.

2. Compute the initial values of $\mathbf{P}$ and error function $\mathcal{E}$.
3. For $t = 1, \dots, N$ repeat:
 - 3.1. Choose λ_{next} randomly within a neighborhood of current λ , and compute $\mathbf{P}_{\text{next}}$.
 - 3.2. Solve for $\mathbf{r}_{\text{next}}$ using SLP method.
 - 3.3. Use $\mathbf{P}_{\text{next}}$ and $\mathbf{r}_{\text{next}}$ to compute $\mathcal{E}_{\text{next}}$.
 - 3.4. Set $\lambda \leftarrow \lambda_{\text{next}}$, $\mathcal{E} \leftarrow \mathcal{E}_{\text{next}}$, and $\mathbf{r} \leftarrow \mathbf{r}_{\text{next}}$ with probability $\min\{1, \exp(-(\mathcal{E}_{\text{next}} - \mathcal{E})t/N)\}$.
4. Correct the initial estimate of ϵ , and update λ
5. Output $c = \lambda\ell/(\ell - k + 1)$, $\mathbf{r}$, $L = B/c$, and ϵ at the end of iterations (B is the total amount of nucleotides sequenced).

Estimating genome lengths

We applied RESPECT and CovEst to simulated genome-skims–Illumina reads sampled from the 66 test genomes skimmed at 1X coverage with 1% sequencing-error rate—and compared their relative error in the estimation of r_1 through r_5 and genome length (Fig. 3), after their convergence (see Supplementary Figures S9 to S14 for the convergence of RESPECT’s estimates). The median RESPECT error in estimating r_1 was less than 1.5% (average: 2.9%), while the median error of CovEst was 15% (average: 34%). The error profile extended to higher multiplicities, where, as noted earlier, CovEst used a parametric model. The tight relation between r_1 and r_2 and the large absolute differences between the two values implied that a small error in r_1 would translate into a large relative error for r_2 , and we observed that for r_2 . Similarly, the RESPECT estimates of genome length were highly accurate with median error 2.2% (average: 4.1%), in contrast to 27% (average: 40%) for CovEst (Fig. 3b). RESPECT estimates were better than CovEst in 62 out of 66 species, often by considerable margins (Fig. 3c). For example, in 54/66 species, RESPECT error was less than 5%, while CovEst error exceeded 50% in one third of test genomes. In fact, CovEst severely underestimates the length for these genomes (Supplementary Fig. S15). For 18/66 test genomes, the CovEST estimate was less than the true length by a factor of 4 or higher (Supplementary Fig. S16).

We repeated the same experiment at sequence level coverage of 0.5X, 2X, and 4X (Supplementary Fig. S17). At 0.5X coverage, the median error of RESPECT was 16% (average: 18%), while CovEst had 88% median error (average: 75%) and underestimated the length by a factor of 8 or more in half of the species (Supplementary Fig. S18). CovEst performance improved at higher coverage but RESPECT continued to have lower error (Supplementary Fig. S19). At 4X, CovEst had median error 3.3% (average: 7.6%), while RESPECT median error was $< 1\%$ (average: 1.9%). Moreover, CovEst error exceeded 10% error in a third of species, while RESPECT had $< 10\%$ error in 64/66 species (Supplementary Fig. S20)

Figure 3: **Iterative estimation of genome length** (a) Comparing the error of RESPECT and CovEst in estimating the repeat spectrum. The first 5 spectra are shown. (b) The distribution of error in CovEst and RESPECT. The absolute value of relative error in genome length estimation is used (in logarithmic scale). (c) Per-genome error of RESPECT and CovEst in estimating the genome length of 66 species with genomes skimmed at 1X coverage.

Estimating genome length using sequenced short reads

A key difference between sequenced reads versus simulated reads is the presence of ‘contaminants’ or reads from non-target species. Differences may also include presence of adapter sequences, duplications of reads from the sequencing platform, lower or higher sequencing error rates due to DNA quality, and length variation of reads. Therefore, we tested RESPECT in genome-skims obtained from NCBI’s Sequence Read Archive (SRA) database²³. We downloaded high-coverage raw reads from 29 test species (from all four major taxonomic groups of eukaryotes in RefSeq) including highly repetitive plant genomes, and compared the results with the corresponding genome assemblies of the same data. After preprocessing the raw reads using BBTools²⁴ to remove adapter sequences and duplicate reads, we used Kraken²⁵ to remove contaminant reads with microbial or human origin (see Methods: ‘SRA preprocessing and contamination filtering’). We note that this is an imperfect process as these tools work only when the contaminating organisms have a

highly related member in the reference databases²⁶. We discarded 10 samples because > 40% of reads (after removing adapters) were either duplicates of other reads, or came from external DNA sources (Supplementary Table S1). For the remaining 19 samples, duplicates and reads classified as contaminant were removed, and unclassified reads were sub-sampled to 1X coverage. In 16 out of 19 samples, RESPECT error was less than 11% (median: 4%), including highly repetitive genomes such as *A. tauschii* ($r_1/L = 0.29$), *Z. mays* (maize) ($r_1/L = 0.32$), *S. salar* (salmon) ($r_1/L = 0.48$), and *N. tabacum* ($r_1/L = 0.57$), where the abundance of repeats made the length estimation challenging (Fig. 4, Table 1). In contrast, CovEst had less than 30% error in only 4 samples (median error 80%) (Fig. 4). For the highly repetitive genomes, CovEst length estimates ranged from 1/11 to 1/7 of the assembled sequence lengths or 10 to 30 times larger error compared to RESPECT (see Table 1). In 3 samples, RESPECT had relatively high errors. For SRR085103 (domestic ferret), 99.9% of the reads did not in fact map to the available reference assembly of the domestic ferret *M. putorius*. Together with the relatively low percentage of duplication (9%) the data suggest a mislabeling of the sample species. For Coquerel’s sifaka (*P. coquereli*), we observed a large gap between the total sequence length (2.8 Gbp) and the total ungapped length (2.1 Gbp) of the assembly, suggesting some challenges with the assembly. Cape elephant shrew (*E. edwardii*) was the last sample where RESPECT length estimate of 4.5Gbp exceeded the RefSeq (GCF_000299155.1) assembly length (3.8Gbp) by over 10%. Interestingly, the uniqueness ratio of the assembly was $r_1/L = 0.72$, which contrasted with the RESPECT estimated uniqueness ratio of $r_1/L = 0.65$ from the short-read data. Upon investigation, we found that a more recent assembly for *E. edwardii* (GCA_004027355.1), not yet in RefSeq, had an assembled length equal to 4.3 Gbp, with $r_1/L = 0.66$, matching the RESPECT estimates (4.5Gb, 0.65, respectively). The difference between total sequence length and ungapped length in GCA_004027355.1 was only 1 Mbp, in contrast to > 500 Mbp for GCF_000299155.1. Together, these data suggest that GCA_004027355.1 better assembles repetitive regions, and the RESPECT length estimation error was < 5%, despite using only 1X coverage.

Species	<i>A. tauschii</i> (goat grass)	<i>Z. mays</i> (maize)	<i>S. salar</i> (salmon)	<i>N. tabacum</i> (tobacco)
r_1/L	0.29	0.32	0.48	0.57
Assembly length (Gbp)	4.3	2.1	3.0	3.6
RESPECT	3.9 (-10.7%)	2.0 (-8.2%)	2.8 (-4.9%)	3.7 (2.6%)
CovEst	0.4 (-90%)	0.2 (-90%)	0.3 (-90%)	0.5 (-86%)

Table 1: Comparing RESPECT and CovEst accuracy on SRA’s of highly repetitive genomes. The numbers in parentheses are the percentage errors.

The role of WGD versus high copy repeat elements in shaping genome repeat structure

Predicting polyploidy and recent WGD is challenging because mutation and gene loss after a WGD event can reduce the polyploidy signal. Specifically, a WGD event results in the uniqueness ratio (r_1/L) becoming 0. Subsequently, as mutations accumulate, r_1/L ratio moves gradually towards

Figure 4: **Estimating genome length using SRA data.** Comparing the error of CovEst and RESPECT. High coverage SRA were preprocessed and later downsampled to 1X coverage. Both methods are applied to genome skims (after preprocessing) and the absolute values of the relative error in estimating the genome lengths are compared.

1 in a process that may be specific to the species, and hard to predict. Nevertheless, it should be skewed toward smaller values for recent WGD events. Independently, the presence of high copy repeats due to DNA transposons and retrotransposons can lead to very high copy numbers of a small set of oligomers. To capture the contribution of high copy repeat elements, we defined the ‘High Copy Repeats per Million (HCRM)’ value as the average count (per million base-pairs) of the 10 most highly repetitive k-mers. HCRM values varied across the species, ranging from 2 to 3738 among our set of 622 RefSeq genomes (Supplementary Fig. S21). We observed some correlation between HCRM values of species of the same genus, especially among vertebrates (Supplementary Fig. S22). However, similar to the case of uniqueness ratios, the phylogenetic signal was not pronounced enough to predict HCRM based on the taxonomy.

Analytical calculations showed that the probability of high HCRM values ≥ 200 in a genome with random set of k -mers was negligibly small ($P \leq 10^{-100}$) (See Methods: ‘Statistical analysis of the repeat structure’), suggesting that high HCRM values could not be explained solely by WGD events, and were likely due to high copy (transposon) repeats. Fig. 5 shows the $(r_1/L, \text{HCRM})$ value of 622 genome-skims, which tightly matched the true values computed from assembled genomes (Supplementary Fig. S23). To analyze the r_1/L and HCRM values of genomes with recent WGD, we compiled a partial list of species with known WGD events within the last 150M years based on the available literature^{27–29} (See Methods: ‘Selecting species with known recent WGD events’ and Supplementary Table S2).

Species with known recent WGD events had expectedly low r_1/L . For example, only 14% of species with recent WGD had r_1/L values ≥ 0.8 , in contrast with 64% of all species that had r_1/L values higher than 0.8. Surprisingly, 93% of species with recent WGD also had low HCRM values (≤ 200) (Fig. 5), and there was a strong association between the occurrence of recent WGD events and the $(r_1/L, \text{HCRM})$ values (p-value: 1.8×10^{-23} ; See Methods: ‘Statistical analysis of the repeat structure’). Our results suggest that genomes with low HCRM and r_1/L are strong candidates for WGD events.

Discussion

In this paper, we revisited the problem of estimating genomic parameters (length, sequence coverage, k -mer spectra) based on low coverage shotgun sequencing data. The problem has been studied previously and was considered challenging due to the need for simultaneous inference of coverage and sequencing errors along with the k -mer spectra. However, our results suggest that the problem remains challenging even when there is no error and the coverage is known. This is due to two factors. (a) The linear system is ill-conditioned, so that a small change in the k -mer counts due to random sampling can lead to large changes in the estimated k -mer spectra (b) Values in the k -mer spectra show a skewed and non-sparse distribution, where r_1 dominates; r_1 is important for length estimation, but controlling for small errors in r_1 leads to larger errors in the other r_h values. We provide evidence of both, but future work will clarify the importance of each facet of the identification.

Proposed solutions for ill-conditioning use regularization but those methods generally enforce sparse solutions. However, the true k -mer distribution is not sparse. Our work resolved this issue through an empirical estimation of k -mer ratios based on finished genomes. This approach is viable given the many finished genomes with different repeat characteristics. Our study, with 662 genomes of which around 10% were isolated for testing, is the largest empirical study of its kind.

As expected, accurately estimated k -mer spectra led to better estimation of genomic parameters such as length, with Skmer-genome performing significantly better than the previous best method, sometimes by orders of magnitude. Our results also have lower variance than those of other methods.

As coverage increases, all methods perform well. However, at coverage 8X and higher, partial assemblies are possible and small contigs can start to be assembled. In those cases, alternative

Figure 5: **High copy repeats per million versus uniqueness ratio among genomes with and without known recent WGD events.** Most of genomes with known recent WGD events had $r_1/L < 0.8$ and $HCRM < 200$. The y-axis is in a logarithmic scale. HCRM values are computed from genome-skims simulated at 1X coverage with no sequencing error. Some of the species with a recent WGD are labeled by their common names.

methods to estimate genome lengths may be possible, but our methods work well even for 0.5X coverage.

The presence of contaminants is a significant barrier to accurate estimations, and in fact is challenging even for assembling the data. As data sampling and DNA extraction methods improve, this problem will likely be less problematic. In parallel, we are also working to improve computational approaches to removing contamination.

While most k -mer based statistics were developed as an initial first step prior to deep sequencing and assembly, they may have an important role to play in independent analysis of genomes. Many genomes are $\leq 1\text{Gb}$ or lower. Therefore acquiring genome-skims for a majority of organisms and even multiple individuals in a population is a feasible goal. Methods that work on these reduced representations can be transformative for studying dramatic and short-term changes in bio-ecology. We can envision technologies where a sampled individual's genome-skim can be used to quickly estimate its genome-length, repeat structure, remove contaminating reads, identify the organism

or place it confidently in the tree of life, and finally, identify the robustness of population through analysis of heterozygosity. Our paper contributes to the first step of this vision.

Methods

Comparing $r1/L$ distribution over different sets. To compare two sets of values and see if the values in one set are greater than the other set, we used the Mann–Whitney U test. Formally, if X and Y are random samples from populations $\mathcal{X}$ and $\mathcal{Y}$, the test statistic, U , is given by the number of times x is greater than y for all $(x, y) \in \mathcal{X} \times \mathcal{Y}$. The Mann–Whitney U test is non-parametric and does not restrict the samples to be from a certain family of distributions. The test also allows the user to specify a location shift μ and examine the alternative hypothesis that $X - Y > \mu$. By gradually increasing μ and computing the p-value, we can understand the extent of difference between X and Y .

To test if two sets of numbers are drawn from the same distribution, we used the two-sample Kolmogorov–Smirnov (KS) test. The test statistic is a distance between the empirical distributions functions of the samples from the two sets. We used R ‘stats’ package³⁰ to compute the p-values for both tests.

Modeling genomic parameters. We consider k -mers in a genome of length L and assume that $k \gg \log_4 L$ so that any k -mer is unlikely to appear more than once, unless it is part of a repeated sequence. Denote the (unknown) k -mer spectrum of a genome that contains repeats using $\mathbf{r}$, where r_j describes the number of distinct k -mers that appear exactly j times in the genome.

The genome is shotgun sequenced using reads of length ℓ with average sequencing depth c . The total number of nucleotides sequenced is given by $B = cL$. As there are $l - k + 1$ k -mers in each read, the k -mer coverage is given by

$$\lambda = (1 - (k - 1)/\ell)c = \frac{(1 - (k - 1)/\ell)B}{L}. \quad (4)$$

Let $\mathbf{o}$ denote the histogram of observed k -mer counts. The observed number of k -mers of abundance h , o_h , can be thought of as a sample allocation to random variable O_h , whose expected value, $m_h = \mathbb{E}[O_h]$, depends upon $\mathbf{r}$, λ , L , and also on sequencing error. We assume that any base-pair is sequenced erroneously with probability ϵ , and sequencing errors only result in novel k -mers. We further assume that the number of times a unique k -mer repeated j times is sampled follows a Poisson distribution with rate $\lambda j(1 - \epsilon)^k$. Therefore

$$\mathbf{m} = \mathbf{rP}^T + \mathbf{1}_{h=1}E, \quad (5)$$

where $P_{hj} = e^{-j\lambda(1-\epsilon)^k} \frac{(j\lambda(1-\epsilon)^k)^h}{h!}$ denotes the probability that a k -mer repeated j times in the genome is observed with count h in the genome skim, $\mathbf{1}_{h=1} = [1, 0, 0, \dots]$, and $E = L\lambda(1 - (1 - \epsilon)^k)$ is the expected number of erroneous k -mers. As λ and L are connected through Eqn. 4, we choose λ as the independent variable and consider L as a function of λ . Under this model, we would like to

estimate $(\mathbf{r}, \lambda, \epsilon) = \arg \min_{\mathbf{r}, \epsilon} \mathcal{E}(\mathbf{P}, \mathbf{r}, \epsilon, \mathbf{o})$, where $\mathcal{E}$ is a weighted p-norm of the difference between expected and observed counts

$$\mathcal{E}_{\mathbf{w}, p}(\mathbf{P}, \mathbf{r}, \epsilon, \mathbf{o}) = \left(\sum_h w_h |m_h - o_h|^p \right)^{1/p} = \left(\sum_h w_h |(\mathbf{r}\mathbf{P}^T + \mathbf{1}_{h=1}E)_h - o_h|^p \right)^{1/p}. \quad (6)$$

Note that the optimization is non-trivial because $\mathbf{P}$ and E are functions of $(\mathbf{r}, \lambda, \epsilon)$, and must be simultaneously estimated.

A generic iterative optimization for parameter estimation. The dimensions of $\mathbf{o}$ and $\mathbf{r}$ in Eqn. 5 are determined entirely by data and are not necessarily identical. However, we truncated both to a common dimension $n = 50$ for computational expediency. A generic optimization method could be described as below.

1. Generate initial estimates of λ, ϵ, L .
2. Solve for $\mathbf{r}$ using Eqn. 6.
3. Use estimated $\mathbf{r}$ and grid-search to re-estimate λ, ϵ .
4. Repeat step 2 onwards until the error has converged.

Step 2 is the key step in this procedure, and we devised a number of approaches to solve it.

Least-squares estimate of repeat spectrum. Choosing $p = 2$ (Euclidean norm) and $w_h = 1, \forall h$ in Eqn. 6, the problem is turned into a Least-Squares (LS) optimization. To test an LS method for estimating $\mathbf{r}$, we considered the simplest sequencing-error-free case ($\epsilon = 0$), where coverage λ was known. Therefore, $\mathbb{E}[\mathbf{O}] = \mathbf{m} = \mathbf{r}\mathbf{P}^T$, where $\mathbf{P}$ is an $n \times n$ matrix with

$$P_{hj} = e^{-j\lambda} \frac{(j\lambda)^h}{h!}.$$

We showed (Supplementary Methods) that $\mathbf{P}$ is non-singular and in the error-free case, it should be possible to use the estimate $\mathbf{r}^{(\text{est})} = \mathbf{o}\mathbf{P}^{-T}$. However, we observed that its effective rank was very small as Λ, E each have rapidly diminishing eigenvalues. Therefore, instead of decomposing $\mathbf{P}$ and explicitly computing $\mathbf{P}^{-1}$, we used the non-negative least squares (NNLS) method³¹ to solve

$$\mathbf{r}^{(\text{est})} = \arg \min_{\mathbf{r}} \|\mathbf{o} - \mathbf{r}\mathbf{P}^T\|_2.$$

We used nnls method from SciPy's³² Optimize library. Unfortunately, the LS estimates were very unreliable and showed high error. In fact, we proved, for $\lambda = 1$ (see Supplementary Methods), that

$$\text{cond}(\mathbf{P}) \geq \frac{2^n}{n}.$$

The condition number grows exponentially with n suggesting a highly ill-conditioned matrix $\mathbf{P}$ where small changes in $\mathbf{o}$ from the expected values $\mathbf{m}$ would lead to large errors in estimate of $\mathbf{r}$. For these reasons, we adopted constrained optimization methods to solve for $\mathbf{r}$.

Linear programming for constrained optimization based estimates. We used Eqn. 6 with $\mathbf{w} = [0, 1, 1, \dots, 1]$ and $p = 1$ to design a Linear programming estimate of $\mathbf{r}$ as:

$$\min_{\mathbf{r}} \sum_{h=2}^n \left| o_h - \sum_{j=1}^n P_{hj} r_j \right|, \quad (7)$$

such that

$$\mathcal{L}_h \leq \frac{r_h}{r_{h+1}} \leq \mathcal{U}_h, \quad h = 1, 2, \dots, n-1$$

The rationale behind setting $w_1 = 0$ was that o_1 contains a large number of erroneous k-mers, so we exclude it from the objective function and use the rest of the bins to estimate $\mathbf{r}$. As ϵ is not known in general, o_1 was used to estimate the (average) sequencing error rate, and subsequently the k -mer coverage λ .

The lower and upper bounds on $\frac{r_j}{r_{j+1}}$ were determined based on the distribution R_j of spectral ratios in 556 training genomes, and therefore we only search for candidate solutions $\mathbf{r}$ that satisfy the constraints. Specifically, we profiled the repeat spectra of the training genomes and set $[\mathcal{L}_j, \mathcal{U}_j]$ equal to the empirical support of R_j distribution, i.e., $\mathcal{L}_j$ and $\mathcal{U}_j$ are the smallest and the largest samples observed from R_j over the training genomes. We use Gurobi Optimizer³³ to solve the constrained optimization problem formulated in Eqn. 7.

Spline Linear programming. The final method of estimating $\mathbf{r}$ is based on the LP estimate of $\mathbf{r}$ and the splines fitted on spectral ratios r_j/r_{j+1} as functions of $\frac{r_j}{\sum_{i \geq j} r_i}$. Formally, let r_j^{LP} denote the LP estimate of r_j by constraining the spectral ratios to be within the support of R_j among the training genomes, as discussed above. For each $j \in \{1, 2, 3, 4, 5\}$, we used a generalized additive model (GAM), learned from 556 training genomes, to predict r_j/r_{j+1} based on $\frac{r_j^{\text{LP}}}{\sum_{i \geq j} r_i^{\text{LP}}}$. Specifically, we model $y_j = r_j/r_{j+1}$ for different genomes as samples drawn from dependent random variable Y_j , which follows gamma distribution and its mean is determined by

$$g_j(\mathbb{E}[Y_j]) = s_j\left(\frac{r_j}{\sum_{i \geq j} r_i}\right), \quad (8)$$

where g_j is called the link function, and s_j is the smoothing spline. These functions allow us to capture nonlinear dependencies between the variables in our model. For $j = 1, 2$, we use a logarithmic link function to account for the large dynamic range of r_j/r_{j+1} over the training set, and use identity link for $j = 3, 4, 5$. For each fitted GAM, we empirically set the smoothing parameter to balance the over-fitting against the goodness of fit. We used R ‘mgcv’ package³⁴ for GAM fitting.

Using the LP estimates of r_j ’s and plugging them into Eqn. 8, we predict the spectral ratios. Let y_j^{SLP} denote the estimate of y_j using Eqn. 8 on previous estimates of $\mathbf{r}$. We *recursively* re-estimate

r_j for $j \in \{2, 3, 4, 5, 6\}$ and call them r_j^{SLP} :

$$r_j^{\text{SLP}} = \begin{cases} r_j^{\text{LP}} & j = 1 \text{ and } j > 6 \\ r_{j-1}^{\text{SLP}}/y_{j-1}^{\text{SLP}} & 2 \leq j \leq 6 \end{cases} \quad (9)$$

RESPECT algorithm. For the RESPECT algorithm, we replaced the basic iterative method described above with a simulated annealing procedure outlined in Algorithm 1 to speed up the computations. To initialize the algorithm, we started with the assumptions that genome has no repeats $\mathbf{r} = [L, 0, 0, \dots]$, and the error-free k -mer counts follow a Poisson distribution (Eqn. 5). Defining $\lambda_{\text{ef}} = \lambda(1 - \epsilon)^k$ as the *error-free k -mer coverage*, we estimate its initial value from the ratio of observed counts

$$\lambda_{\text{ef}} = \frac{(h^* + 1)o_{h^*+1}}{o_{h^*}}, \text{ where } h^* = \arg \max_{h>1} o_h,$$

and set

$$\lambda = e^{-\lambda_{\text{ef}}} \frac{\lambda_{\text{ef}}^{h^*}}{h^*!} \frac{o_1}{o_{h^*}} + \lambda_{\text{ef}}(1 - e^{-\lambda_{\text{ef}}}), \quad \epsilon = 1 - (\lambda_{\text{ef}}/\lambda)^{1/k}$$

(see Supplementary Methods). The above estimate of ϵ is used throughout the algorithm, but is corrected at the end based on the estimated uniqueness ratio (described below). Using the estimate of λ_{ef} , we compute $\mathbf{P}$, and thus the error function $\mathcal{E}$ at the start of the algorithm. For $\mathcal{E}$, we chose $\mathbf{w} = [0, 1, 1, \dots, 1]$ and $p = 1$ in Eqn. 6, so

$$\mathcal{E} = \sum_{h=2}^n \left| o_h - \sum_{j=1}^n P_{hj} r_j \right|$$

With the initial values of the parameters known, RESPECT runs a simulated annealing optimization until the error converges. At each iteration, a candidate λ_{next} in $[\frac{1}{2}\lambda, 3\lambda]$ is selected uniformly at random, and $\mathbf{P}_{\text{next}}$ is computed from $\lambda_{\text{next}}(1 - \epsilon)^k$. Next, we run SLP method on $(\mathbf{o}, \mathbf{P}_{\text{next}})$ to get $\mathbf{r}_{\text{next}}$. Throughout the algorithm, we used truncated $\mathbf{o}_{1 \times m}$, $\mathbf{r}_{1 \times n}$, and $\mathbf{P}_{m \times n}$ where the number of spectra is fixed at $n = 50$ (a reasonable compromise between accuracy and speed), and the number of observed counts $m = n \cdot \max(1, \lambda_{\text{ef}})$ scales proportionally with the initial estimate of error-free k -mer coverage. Using $(\mathbf{o}, \mathbf{P}_{\text{next}}, \mathbf{r}_{\text{next}})$, error function for the candidate state $\mathcal{E}_{\text{next}}$ is calculated. If moving to the candidate state results in a reduction in the error ($\mathcal{E}_{\text{next}} < \mathcal{E}$), the algorithm accepts the move and updates the current estimate of parameters. In addition, to help the algorithm deal with local minima and find better solutions, a simulated annealing scheme is implemented such that the algorithm probabilistically decides to move to states with higher error. Specifically, at iteration t , even if $\mathcal{E}_{\text{next}} > \mathcal{E}$, the algorithm accepts the move with probability $\exp(-(\mathcal{E}_{\text{next}} - \mathcal{E})t/N)$.

At the end of iterations, the initial estimate of ϵ (obtained under no-repeats assumption) is corrected based on the estimated value of r_1/L . The correction was learned over 120 genomes

Algorithm 1: The RESPECT method

Start with $\lambda_{\text{ef}} = \lambda^{(0)}(1 - \epsilon)^k = \frac{(h^*+1)o_{h^*+1}}{o_{h^*}}$, where $h^* = \arg \max_{h>1} o_h$;
 Compute $\mathbf{P}^{(0)}$, $\mathcal{E}^{(0)} = \min_{\mathbf{r}} \mathcal{E}(\mathbf{P}^{(0)}, \mathbf{r}^{(0)}, \mathbf{o})$, and $\mathbf{r}^{(0)} = \arg \min_{\mathbf{r}} \mathcal{E}(\mathbf{P}^{(0)}, \mathbf{r}^{(0)}, \mathbf{o})$;
 Find $E = o_1 - \sum_j P_{1j}^{(0)} r_j^{(0)}$;
 Set $\lambda^{(0)} = e^{-\lambda_{\text{ef}} \frac{\lambda_{\text{ef}}^{h^*}}{h^*!} \frac{o_1}{o_{h^*}}} + \lambda_{\text{ef}}(1 - e^{-\lambda_{\text{ef}}})$, and compute ϵ from λ_{ef} and $\lambda^{(0)}$;
for $1 \leq t \leq N$ **do**
 $\lambda^{(t)} \leftarrow \mathcal{U}[\frac{1}{2} \cdot \lambda^{(t-1)}, 3 \cdot \lambda^{(t-1)}]$;
 Use $\lambda^{(t)}$ and ϵ to compute $\mathbf{P}^{(t)}$, $\mathbf{r}^{(t)} = \arg \min_{\mathbf{r}} \mathcal{E}(\mathbf{P}^{(t)}, \mathbf{r}^{(t)}, \mathbf{o})$, and $\mathcal{E}^{(t)} = \min_{\mathbf{r}} \mathcal{E}(\mathbf{P}^{(t)}, \mathbf{r}^{(t)}, \mathbf{o})$;
 Move to $\lambda^{(t)}$ with probability $\min \left\{ 1, \exp \left(\frac{\mathcal{E}^{(t-1)} - \mathcal{E}^{(t)}}{N-t+1} \right) \right\}$;
end
 Correct ϵ and set $\lambda = \lambda^{(N)}(1 - \epsilon)^k / (1 - \epsilon_{\text{corrected}})^k$;
 Output $c = \frac{\ell}{\ell-k+1} \lambda$, $L = B/c$, $\epsilon_{\text{corrected}}$, and $\mathbf{r}^{(N)}$

randomly selected from the training set, and applied if the estimated coverage is smaller than 1.5X. Then, λ is re-computed based on the corrected ϵ , and is used to compute the final estimates of coverage and genome length. The estimated sequencing error rate and repeat spectrum are also provided by the algorithm.

SRA preprocessing and contamination filtering. After downloading SRA accessions and converting them to FASTQ using SRA Toolkit³⁵, we used BBduk and Dedupe from BBTools package to trim adapter sequences and remove duplicate reads. We then ran Kraken2 to remove contamination with prokaryotic or human origin. For plant and invertebrate samples, we filtered out any read that was classified to the Kraken database at 0 confidence level (very sensitive, a single matched k -mer is enough for the classification). For vertebrates, due to their smaller evolutionary distance to homo sapiens, we required 0.5 confidence level (more specific, half of the read's k -mers should match) for human classification, and 0 confidence level for everything else in the database.

Implementation details and running time. We use ‘count’ and ‘histo’ commands from Jellyfish³⁶ command line tool to compute the k -mer histogram of input genome-skims. In each iteration of RESPECT algorithm, we solve a constrained optimization problem using the tools provided by Gurobi Python interface in ‘gurobipy’ package. The running time of RESPECT slowly increases with the coverage as the size of $\mathbf{P}$ (and hence the size of optimization problem at each iteration) scales with the (initial) estimate of coverage. On average, for a typical 0.5X-4X coverage of genome-skims, it takes about 2 hours for RESPECT algorithm to converge and produce the final estimate of the parameters.

Selecting species with known recent WGD events. From the total of 83 RefSeq genomes in our database, we obtained the WGD annotation (with estimated age) for 44 plant species²⁸. WGD annotations for the remaining 32 plant species in our database were based on the data provided by the 1000 plants project²⁹, where either the exact same species or a species from the same genus is identified to have undergone a WGD event using transcriptomic data. We also have 7 Salmonid genomes where their common ancestor is thought to have had a WGD event about 80My ago²⁷.

Statistical analysis of the repeat structure. In a random genome with length L , there are $L - k + 1 \simeq L$ k -mers, and assuming the random selection of k -mers is uniform over the space of all 4^k possible k -mers, the probability distribution for the copy number (CN) of each k -mer is

$$\text{Prob}[\text{CN} = x] = \binom{L}{x} \left(\frac{1}{4^k}\right)^x \left(1 - \frac{1}{4^k}\right)^{L-x}.$$

For typical values of $L \sim 100 - 1000$ Mbp and $k = 31$, the conditions to use a Poisson distribution to approximate a Binomial (see e.g., Section 5.4 of³⁷) are met, i.e., $L \gg 1$ and $4^{-k} \ll 1$, hence we have

$$\text{Prob}[\text{CN} = x] = e^{-L/4^k} \frac{(L/4^k)^x}{x!}.$$

If the genome subsequently undergoes n_w whole genome duplication events, the genome length is multiplied by 2^{n_w} . However, the multiplicity of each k -mer increases by at most 2^{n_w} , as mutations reduce the copy number of k -mers. Therefore, to have an HCRM value of H , there should exist at least a k -mer with copy number $x \geq HL$ in the original random genome. Now, considering that under random-genome model the selection of *any* k -mer is equally likely, we can use the union bound (see e.g., Section 1.5 of³⁷) and have

$$\begin{aligned} \text{Prob}[\text{HCRM} \geq H] &< \sum_{\substack{\text{all} \\ k\text{-mers}}} \sum_{x=HL} e^{-L/4^k} \frac{(L/4^k)^x}{x!} \\ &< 4^k \sum_{x=HL} e^{-L/4^k} \frac{(L/4^k)^x}{x!}. \end{aligned} \quad (10)$$

We used WolframAlpha³⁸ to compute the bound in (10) for several values of H . For $H = 200$ and $L \in [100 - 1000]$ Mbp, the resulting p -values were less than 10^{-100} .

To test the association between WGD events and the values of r_1/L and HCRM, we used the assembled genomes of 622 RefSeq species and constructed a two by two contingency table where columns represent the species with or without an identified recent WGD, and the rows specify whether or not the genome has r_1/L and HCRM values less than 0.8 and 200, respectively. We filled the table by the count of genomes that satisfied each of these four conditions, and performed a Fisher's exact test (using R 'stats' package³⁰) and got the p -value = 1.8×10^{-23} for the correlation between the rows and columns of the table.

Availability of data and material. RESPECT software and the trained models are publicly available on <https://github.com/shahab-sarmashghi/RESPECT> under a BSD 3-Clause license. Run accessions of SRA files used in the tests are provided in Supplementary Table S1.

Competing interests. The authors declare that they have no competing interests.

References

- [1] E. Brondizio, J. Settele, S. Diaz, and H. Ngo, “Global assessment report on biodiversity and ecosystem services of the intergovernmental science-policy platform on biodiversity and ecosystem services,” *IPBES Secretariat, Bonn*, 2019.
- [2] K. V. Rosenberg, A. M. Dokter, P. J. Blancher, J. R. Sauer, A. C. Smith, P. A. Smith, J. C. Stanton, A. Panjabi, L. Helft, M. Parr, and P. P. Marra, “Decline of the North American avifauna,” *Science*, p. eaaw1313, sep 2019.
- [3] H. A. Lewin, G. E. Robinson, W. J. Kress, W. J. Baker, J. Coddington, K. A. Crandall, R. Durbin, S. V. Edwards, F. Forest, M. T. P. Gilbert, M. M. Goldstein, I. V. Grigoriev, K. J. Hackett, D. Haussler, E. D. Jarvis, W. E. Johnson, A. Patinos, S. Richards, J. C. Castilla-Rubio, M.-A. van Sluys, P. S. Soltis, X. Xu, H. Yang, and G. Zhang, “Earth BioGenome Project: Sequencing life for the future of life,” *Proceedings of the National Academy of Sciences*, vol. 115, pp. 4325–4333, apr 2018.
- [4] P. D. N. Hebert, A. Cywinska, S. L. Ball, and J. R. deWaard, “Biological identifications through DNA barcodes,” *Proceedings of the Royal Society B: Biological Sciences*, vol. 270, no. 1512, pp. 313–321, 2003.
- [5] V. Savolainen, R. S. Cowan, A. P. Vogler, G. K. Roderick, and R. Lane, “Towards writing the encyclopaedia of life: an introduction to DNA barcoding,” *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 360, no. 1462, pp. 1805–1811, 2005.
- [6] P. TABERLET, E. COISSAC, F. POMPANON, C. BROCHMANN, and E. WILLERSLEV, “Towards next-generation biodiversity assessment using DNA metabarcoding,” *Molecular Ecology*, vol. 21, pp. 2045–2050, 4 2012.
- [7] M. J. Hickerson, C. P. Meyer, C. Moritz, and M. Hedin, “DNA Barcoding Will Often Fail to Discover New Animal Species over Broad Parameter Space,” *Systematic Biology*, vol. 55, pp. 729–739, 10 2006.
- [8] D. L. J. Quicke, M. Alex Smith, D. H. Janzen, W. Hallwachs, J. Fernandez-Triana, N. M. Laurence, A. Zaldívar-Riverón, M. R. Shaw, G. R. Broad, S. Klopstein, S. R. Shaw, J. Hrecek, P. D. N. Hebert, S. E. Miller, J. J. Rodriguez, J. B. Whitfield, M. J. Sharkey, B. J. Sharanowski, R. Jussila, I. D. Gauld, D. Chesters, and A. P. Vogler, “Utility of the DNA barcoding gene fragment for parasitic wasp phylogeny (Hymenoptera: Ichneumonidae): Data release and new measure of taxonomic congruence,” *Molecular Ecology Resources*, vol. 12, pp. 676–685, 7 2012.
- [9] S. Liu, Y. Li, J. Lu, X. Su, M. Tang, R. Zhang, L. Zhou, C. Zhou, Q. Yang, Y. Ji, D. W. Yu, and X. Zhou, “SOAP Barcode: revealing arthropod biodiversity through assembly of

- Illumina shotgun sequences of PCR amplicons,” *Methods in Ecology and Evolution*, vol. 4, pp. 1142–1150, 12 2013.
- [10] “DNAMark.” <http://dnamark.ku.dk/english/>.
- [11] “France Génomique - Mutualisation des compétences et des équipements français pour l’analyse génomique et la bio-informatique.” <https://www.france-genomique.org/>.
- [12] E. Coissac, P. M. Hollingsworth, S. Laverigne, and P. Taberlet, “From barcodes to genomes: Extending the concept of DNA barcoding,” 2016.
- [13] K. Bohmann, S. Mirarab, V. Bafna, and M. T. P. Gilbert, “Beyond DNA barcoding: The unrealized potential of genome skim data in sample identification,” *Molecular Ecology*, vol. 29, pp. 2521–2534, jul 2020.
- [14] S. Sarmashghi, K. Bohmann, M. T. P. Gilbert, V. Bafna, and S. Mirarab, “Skmer: assembly-free and alignment-free sample identification using genome skims,” *Genome Biol.*, vol. 20, p. 34, 02 2019.
- [15] M. Balaban, S. Sarmashghi, and S. Mirarab, “APPLES: Scalable Distance-based Phylogenetic Placement with or without Alignments,” *Systematic Biology*, 2019.
- [16] X. Li and M. S. Waterman, “Estimating the repeat structure and length of DNA sequences using L-tuples,” *Genome research*, vol. 13, pp. 1916–22, aug 2003.
- [17] D. Williams, W. L. Trimble, M. Shilts, F. Meyer, and H. Ochman, “Rapid quantification of sequence repeats to resolve the size, structure and contents of bacterial genomes,” *BMC Genomics*, 2013.
- [18] M. Hozza, T. Vinař, and B. Brejová, “How Big is that Genome? Estimating Genome Size and Coverage from k-mer Abundance Spectra,” in *String Processing and Information Retrieval*, (Cham), pp. 199–209, Springer International Publishing, 2015.
- [19] P. Melsted and J. K. Pritchard, “Efficient counting of k-mers in DNA sequences using a bloom filter,” *BMC Bioinformatics*, 2011.
- [20] P. Melsted and B. V. Halldórsson, “KmerStream: Streaming algorithms for k-mer abundance estimation,” *Bioinformatics*, 2014.
- [21] G. Wahba, *Spline models for observational data*. SIAM, 1990.
- [22] T. J. Hastie and R. J. Tibshirani, *Generalized additive models*, vol. 43. CRC press, 1990.
- [23] R. Leinonen, H. Sugawara, M. Shumway, and I. N. S. D. Collaboration, “The sequence read archive,” *Nucleic acids research*, vol. 39, no. suppl_1, pp. D19–D21, 2010.
- [24] B. Bushnell, “BBMap.” <https://sourceforge.net/projects/bbmap/>.

- [25] D. E. Wood, J. Lu, and B. Langmead, “Improved metagenomic analysis with kraken 2,” *Genome biology*, vol. 20, no. 1, p. 257, 2019.
- [26] E. Rachtman, M. Balaban, V. Bafna, and S. Mirarab, “The impact of contaminants on the accuracy of genome skimming and the effectiveness of exclusion read filters,” *Molecular Ecology Resources*, vol. 20, pp. 1755–0998.13135, may 2020.
- [27] S. Lien, B. F. Koop, S. R. Sandve, J. R. Miller, M. P. Kent, T. Nome, T. R. Hvidsten, J. S. Leong, D. R. Minkley, A. Zimin, *et al.*, “The atlantic salmon genome provides insights into rediploidization,” *Nature*, vol. 533, no. 7602, pp. 200–205, 2016.
- [28] Y. Van de Peer, E. Mizrachi, and K. Marchal, “The evolutionary significance of polyploidy,” *Nature Reviews Genetics*, vol. 18, no. 7, p. 411, 2017.
- [29] O. T. P. T. Initiative *et al.*, “One thousand plant transcriptomes and the phylogenomics of green plants,” *Nature*, vol. 574, no. 7780, p. 679, 2019.
- [30] R Core Team, *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2019.
- [31] C. L. Lawson and R. J. Hanson, *Solving least squares problems*. SIAM, 1995.
- [32] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, Í. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors, “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python,” *Nature Methods*, vol. 17, pp. 261–272, 2020.
- [33] Gurobi Optimization, LLC, “Gurobi optimizer reference manual.” <http://www.gurobi.com>, 2020.
- [34] S. N. Wood, “Fast stable restricted maximum likelihood and marginal likelihood estimation of semiparametric generalized linear models,” *Journal of the Royal Statistical Society (B)*, vol. 73, no. 1, pp. 3–36, 2011.
- [35] SRA Toolkit Development Team, “SRA-Tools.” <http://ncbi.github.io/sra-tools/>.
- [36] G. Marçais and C. Kingsford, “A fast, lock-free approach for efficient parallel counting of occurrences of k-mers,” *Bioinformatics*, vol. 27, pp. 764–770, 3 2011.
- [37] M. H. DeGroot and M. J. Schervish, *Probability and statistics*. Pearson Education, 2012.
- [38] Wolfram Alpha LLC, “Wolfram|Alpha.” <https://www.wolframalpha.com/widgets/view.jsp?id=74e8bb60ad4e38d6a1b0dc865d7197ff>. Accessed: Dec. 09, 2020.

- [39] C. D. Meyer, *Matrix analysis and applied linear algebra*, vol. 71. Siam, 2000.
- [40] K. B. Petersen and M. S. Pedersen, “The matrix cookbook.” <http://www2.compute.dtu.dk/pubdb/pubs/3274-full.html>, Nov 2012. Version 20121115.
- [41] W. Gautschi, “On inverses of vandermonde and confluent vandermonde matrices,” *Numerische Mathematik*, vol. 4, pp. 117–123, Dec 1962.
- [42] H. Robbins, “A remark on stirling’s formula,” *The American mathematical monthly*, vol. 62, no. 1, pp. 26–29, 1955.

Supplementary Material

A Supplemental Methods

Initial estimate of parameters. With the assumption that a genome has no repeating k -mers $\mathbf{r} = [L, 0, 0, \dots]$, Eq. (5) reduces to

$$\mathbb{E}[o_h] = m_h = \begin{cases} L\lambda(1-\epsilon)^k e^{-\lambda(1-\epsilon)^k} + L\lambda(1-(1-\epsilon)^k) & h = 1 \\ L \frac{(\lambda(1-\epsilon)^k)^h}{h!} e^{-\lambda(1-\epsilon)^k} & h > 1 \end{cases}. \quad (\text{S1})$$

We use the method of moments (see e.g., Section 7.6 of³⁷) and set $m_h = o_h$ to estimate the underlying parameters λ and ϵ . Specifically, let $h^* = \arg \max_{h>1} o_h$ be the multiplicity with the largest number of observed k -mers (excluding the unique ones $h = 1$). We use o_{h^*+1}/o_{h^*} to estimate $\lambda_{\text{ef}} = \lambda(1-\epsilon)^k$

$$\lambda_{\text{ef}} = \frac{(h^* + 1)o_{h^*+1}}{o_{h^*}}. \quad (\text{S2})$$

Then, using o_{h^*}/o_1 , we estimate λ as

$$\lambda = \lambda_{\text{ef}}^h e^{-\lambda_{\text{ef}}} \frac{o_1}{h^*! o_{h^*}} - \lambda_{\text{ef}} e^{-\lambda_{\text{ef}}} + \lambda_{\text{ef}}, \quad (\text{S3})$$

and estimate ϵ from the ratio of λ_{ef} and λ

$$\epsilon = 1 - \left(\frac{\lambda_{\text{ef}}}{\lambda}\right)^{1/k}. \quad (\text{S4})$$

Least-squares estimate of repeat spectrum. Consider the cost function defined in Eq. (6) with $p = 2$ and $w_h = 1$ for all h

$$\mathcal{E}_{\mathbf{w},p}(\mathbf{P}, \mathbf{r}, \epsilon, \mathbf{o}) = \left(\sum_h |m_h - o_h|^2 \right)^{1/2} = \left(\sum_h |(\mathbf{rP}^T + \mathbf{1}_{h=1}E)_h - o_h|^2 \right)^{1/2}. \quad (\text{S5})$$

We considered the simplest sequencing-error-free case ($\epsilon = 0$), where coverage λ was known. Therefore, $\mathbb{E}[\mathbf{O}] = \mathbf{m} = \mathbf{rP}^T$, where $\mathbf{P}$ is an $n \times n$ matrix with

$$P_{hj} = e^{-j\lambda} \frac{(j\lambda)^h}{h!}. \quad (\text{S6})$$

$\mathbf{P}$ can be decomposed as $\mathbf{P} = \mathbf{\Lambda V E}$ where $\mathbf{\Lambda}$ is a diagonal matrix with $\Lambda_{hh} = \frac{\lambda^h}{h!}$, $\mathbf{E}$ is a diagonal matrix with $\mathbf{E}_{jj} = j e^{-j\lambda}$, and $\mathbf{V}$ is the transpose of a Vandermonde matrix with the second column given by the vector $(1, 2, 3, \dots, n)^T$; thus $V_{hj} = j^{h-1}$, and

$$\mathbf{V} = \begin{bmatrix} 1 & 1 & 1 & \cdots & 1 \\ 1 & 2 & 3 & \cdots & n \\ 1 & 2^2 & 3^2 & \cdots & n^2 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 2^{n-1} & 3^{n-1} & \cdots & n^{n-1} \end{bmatrix}. \quad (\text{S7})$$

Note that $\mathbf{\Lambda}$ and $\mathbf{E}$ are diagonal matrices with non-zero diagonal elements and hence, they are non-singular. Also, since V is a Vandermonde matrix, we have $\det(\mathbf{V}) = \prod_{1 \leq i < j \leq n} (j - i) > 0$ which renders $\mathbf{V}$ a non-singular matrix. Thus, it enables us to use the estimate $\mathbf{r}^{(\text{est})} = \mathbf{oP}^{-\mathbf{T}}$. However, for the case $\lambda = 1$ we prove that

$$\text{cond}(\mathbf{P}) \geq c \frac{2^n}{n},$$

suggesting a highly ill-conditioned matrix, and making the LS estimates unreliable.

Bound on $\text{cond}(\mathbf{P})$. To establish the aforementioned bound for $\text{cond}(\mathbf{P})$, as $\mathbf{P}$ is non-singular, we have (see e.g., Section 3.8 of³⁹)

$$\text{cond}(\mathbf{P}) = \|\mathbf{P}\| \|\mathbf{P}^{-1}\| = \|\mathbf{\Lambda V E}\| \cdot \|\mathbf{E}^{-1} \mathbf{V}^{-1} \mathbf{\Lambda}^{-1}\|, \quad (\text{S8})$$

where $\|\mathbf{A}\|$ is the induced 2-norm of the matrix $\mathbf{A}$. Since both $\mathbf{E}$ and $\mathbf{\Lambda}$ are diagonal, we have

$$\begin{aligned} \text{cond}(\mathbf{\Lambda}) = \text{cond}(\mathbf{\Lambda}^{-1}) &= \frac{\max_i \{\lambda^i / i!\}}{\min_i \{\lambda^i / i!\}} = \frac{\lambda^{\lfloor \lambda \rfloor} / \lfloor \lambda \rfloor!}{\min(1, \lambda^n / n!)} \\ \text{cond}(\mathbf{E}) = \text{cond}(\mathbf{E}^{-1}) &= \frac{\max_i \{i e^{-i\lambda}\}}{\min_i \{i e^{-i\lambda}\}} = \frac{f(\lambda)}{\min(e^{-\lambda}, n e^{-n\lambda})}, \end{aligned}$$

where

$$f(\lambda) = \begin{cases} \max(\lfloor \lambda \rfloor e^{-\lambda \lfloor \lambda \rfloor}, (\lfloor \lambda \rfloor + 1) e^{-\lambda(\lfloor \lambda \rfloor + 1)}) & \lambda < 1 \\ e^{-\lambda} & \lambda \geq 1 \end{cases}.$$

For the simplicity of exposition, we use $\lambda = 1$. Hence,

$$\begin{aligned} \text{cond}(\mathbf{\Lambda}) = \text{cond}(\mathbf{\Lambda}^{-1}) &= n!, \\ \text{cond}(\mathbf{E}) = \text{cond}(\mathbf{E}^{-1}) &= \frac{1}{n} e^{n-1}. \end{aligned} \quad (\text{S9})$$

Note that for the 2-norm $\|\cdot\|$, we have

$$\|\mathbf{AB}\| \leq \|\mathbf{A}\|\|\mathbf{B}\| \quad (\text{S10})$$

hence, for any matrix $\mathbf{B}$ and any invertible matrix $\mathbf{A}$, we have $\|\mathbf{B}\| = \|\mathbf{A}^{-1}\mathbf{AB}\| \leq \|\mathbf{A}^{-1}\|\|\mathbf{AB}\|$, implying

$$\|\mathbf{AB}\| \geq \frac{\|\mathbf{B}\|}{\|\mathbf{A}^{-1}\|}. \quad (\text{S11})$$

By repeated application of this inequality to Eq. (S8) and using (S9), we have

$$\text{cond}(\mathbf{P}) \geq \frac{\text{cond}(\mathbf{V})}{\text{cond}(\mathbf{E}) \cdot \text{cond}(\mathbf{A})} = \frac{n}{n!e^{n-1}} \text{cond}(\mathbf{V}). \quad (\text{S12})$$

To show that $\text{cond}(\mathbf{P})$ grows exponentially, we use the following lemma.

Lemma 1 *For the matrix $\mathbf{V}$ given in (S7), we have*

$$c \frac{1}{n^{3/2}} (2n)^n \leq \text{cond}(\mathbf{V}) \leq C (2n)^n. \quad (\text{S13})$$

for some $c, C > 0$.

Proof. Since $\mathbf{V}$ is non-singular, $\text{cond}(\mathbf{V}) = \|\mathbf{V}\|\|\mathbf{V}^{-1}\|$ and hence, it remains to bound $\|\mathbf{V}\|$ and $\|\mathbf{V}^{-1}\|$. To do this, we use the following inequalities relating several norms of a matrix $\mathbf{A}$ (see e.g., Section 10.4.4 of⁴⁰)

$$\frac{1}{\sqrt{n}} \|\mathbf{A}\|_F \leq \|\mathbf{A}\| \leq \|\mathbf{A}\|_F, \quad (\text{S14})$$

$$\frac{1}{\sqrt{n}} \|\mathbf{A}\|_\infty \leq \|\mathbf{A}\| \leq \sqrt{n} \|\mathbf{A}\|_\infty, \quad (\text{S15})$$

where $\|\mathbf{A}\|_F$ and $\|\mathbf{A}\|_\infty$ are the Frobenius and the induced ∞ -norm of $\mathbf{A}$, respectively.

1. **Bounding $\|\mathbf{V}\|$:** Using (S14) and $\|\mathbf{V}\|_F^2 = \sum_{i=1}^n \sum_{j=1}^n j^{2i-2} \geq n^{2(n-1)}$, we get

$$\|\mathbf{V}\|_F \geq n^{n-1}.$$

For an upper bound on $\|\mathbf{V}\|_F$, we have

$$\begin{aligned}
\|\mathbf{V}\|_F^2 &= \sum_{i=1}^n \sum_{j=1}^n j^{2i-2} \\
&\leq \sum_{i=1}^n \left(n^{2i-2} + \int_1^n x^{2i-2} dx \right) \\
&= \sum_{i=1}^n \left(n^{2i-2} + \frac{n^{2i-1} - 1}{2i-1} \right) \\
&= \sum_{i=1}^n n^{2i-2} + \sum_{i=1}^n \frac{n^{2i-1}}{2i-1},
\end{aligned}$$

where the inequality follows from x^k being a monotonically increasing function for positive x and all $k \geq 1$. But

$$\sum_{i=1}^n n^{2i-2} = \sum_{i=1}^n (n^2)^{i-1} = \frac{n^{2n} - 1}{n^2 - 1} \leq \frac{n^{2n}}{n^2 - 1}. \quad (\text{S16})$$

Therefore, for $n \geq 3$, we have

$$\begin{aligned}
\|\mathbf{V}\|_F^2 &\leq \frac{n^{2n}}{n^2 - 1} + \sum_{i=1}^n \frac{n^{2i-1}}{2i-1} \\
&= \frac{n^{2n}}{n^2 - 1} + \frac{n^{2n-1}}{2n-1} + \sum_{i=1}^{n-1} \frac{n^{2i-1}}{2i-1} \\
&\leq \frac{n^{2n}}{n^2 - 1} + \frac{n^{2n-1}}{2n-1} + (n-1) \frac{n^{2n-3}}{2(n-1)-1} \\
&\leq \left(\frac{1}{1-n^{-2}} + \frac{1}{2-n^{-2}} + \frac{1}{2n-3} \right) n^{2n-2} \\
&\leq 2n^{2n-2},
\end{aligned}$$

where the second inequality follows from $\frac{n^{2i-1}}{2i-1}$ being an increasing function of i for $n \geq 3$. It can be verified that the above bound holds for $n = 1, 2$ too. Using this fact, and the lower bound on $\|\mathbf{V}\|_F$, we have

$$\frac{1}{\sqrt{n}} n^{n-1} \leq \|\mathbf{V}\| \leq \sqrt{2} n^{n-1}. \quad (\text{S17})$$

2. **Bounding $\|\mathbf{V}^{-1}\|$:** Using Theorem 1 by Gautschi (1962)⁴¹, we have:

$$\|\mathbf{V}^{-1}\|_\infty = \max_i \prod_{j \neq i} \frac{1+j}{|j-i|}. \quad (\text{S18})$$

Note that for a fixed i ,

$$\prod_{j \neq i} (1+j) = \frac{(n+1)!}{i+1},$$

and

$$\prod_{j \neq i} \frac{1}{|j-i|} = \frac{1}{(i-1)!(n-i)!}.$$

Therefore,

$$\prod_{j \neq i} \frac{1+j}{|j-i|} = \frac{(n+1)!}{(i+1)(i-1)!(n-i)!} = \binom{n}{i} \frac{i}{i+1} (n+1).$$

Replacing this in (S18), we get

$$\frac{1}{2}(n+1) \max_i \binom{n}{i} \leq \|\mathbf{V}^{-1}\|_{\infty} \leq (n+1) \max_i \binom{n}{i}. \quad (\text{S19})$$

But $\max_i \binom{n}{i} = \binom{n}{\lfloor n/2 \rfloor}$ and hence, by the Sterling's approximation⁴², we have

$$a \frac{1}{\sqrt{n}} 2^n \leq \max_i \binom{n}{i} \leq A \frac{1}{\sqrt{n}} 2^n$$

for some $a < A$. Using this in (S19), we arrive at

$$b\sqrt{n}2^n \leq \|\mathbf{V}^{-1}\|_{\infty} \leq B\sqrt{n}2^n,$$

for some $b, B > 0$. Therefore, using (S15), we get

$$b2^n \leq \|\mathbf{V}^{-1}\| \leq Bn2^n. \quad (\text{S20})$$

Finally, combining the bounds (S17) and (S20) on $\|\mathbf{V}\|$ and $\|\mathbf{V}^{-1}\|$, respectively, and using $\text{cond}(\mathbf{V}) = \|\mathbf{V}\| \|\mathbf{V}^{-1}\|$, we get the desired result

$$c \frac{1}{n^{3/2}} (2n)^n \leq \text{cond}(\mathbf{V}) \leq C(2n)^n,$$

for some constants $c, C > 0$.

■

Now, we are ready to show that $\mathbf{P}$ is a highly ill-conditioned matrix.

Lemma 2 *For the matrix $\mathbf{P}$ (given by (S6)), we have*

$$\text{cond}(\mathbf{P}) \geq c \frac{2^n}{n}, \quad (\text{S21})$$

for some $c > 0$.

Proof. By the application of the lower-bound of Lemma 1 to (S12), we have

$$\text{cond}(\mathbf{P}) \geq \frac{n}{n!e^{n-1}} \text{cond}(\mathbf{V}) \geq \frac{n}{n!e^{n-1}} \cdot c \frac{1}{n^{3/2}} (2n)^n = \frac{c(2n)^n}{\sqrt{n}n!e^{n-1}}.$$

Therefore, using the Sterling's approximation $n! \leq n^{n+1/2}e^{-(n-1)}$, we get

$$\text{cond}(\mathbf{P}) \geq c \frac{2^n}{n}$$

for some constant $c > 0$. ■

Since the bound (S21) can be written as

$$\text{cond}(\mathbf{P}) \geq c 2^{n-\log_2 n},$$

and for any $\epsilon > 0$, we have $n - \log_2 n \geq (1 - \epsilon)n$ for sufficiently large n , we have $\text{cond}(\mathbf{P}) \geq c 2^{(1-\epsilon)n}$.

Therefore, $\text{cond}(\mathbf{P})$ grows exponentially and in fact, $\text{cond}(\mathbf{P}) = \Omega(\alpha^n)$ for any $\alpha < 2$.

B Supplementary figures and tables

Figure S1: **Whole RefSeq taxonomy with r_1/L annotation.** (a) Plants. (b) Invertebrates. (c) Mammals. (d) Other vertebrates.

Figure S2: Distributions of intra-generic versus inter-generic differences in r_1/L for pairs of RefSeq species. (a) Plants. (b) Invertebrates. (c) Mammals. (d) Other vertebrates.

Figure S3: Correlation of r_1/L with spectral ratios. (a) r_1/r_3 versus r_1/L , (b) r_1/r_5 versus r_1/L .

Figure S4: **Comparing the distributions of r_1/L among test and all RefSeq genomes.** The p-value for the hypothesis that the distributions are different using two-sided Kolmogorov–Smirnov test is 0.93. Highly-repetitive genomes are slightly over-represented in the test set.

Figure S5: Correlation between true r_4/r_3 and estimated $r_3/\sum_{i=3} r_i$.

Figure S6: Correlation between true r_5/r_4 and estimated $r_4/\sum_{i=4} r_i$.

Figure S7: **Correlation between true r_6/r_5 and estimated $r_5/\sum_{i=5} r_i$.**

Figure S8: Correlation between the relative error in the estimated sequencing error and the uniqueness ratio A subset of 120 training genomes were selected as the cross-validation set, and genome-skims were simulated at 1X coverage with 1% sequencing error rate. There is a strong correlation ($R = -0.995$) between the error in estimating ϵ and r_1/L ratio. We capped the correction at 20% (red dashed line).

Figure S9: r_1 estimation convergence with time.

Figure S10: r_2 estimation convergence with time.

Figure S11: r_3 estimation convergence with time.

Figure S12: r_4 estimation convergence with time.

Figure S13: r_5 estimation convergence with time.

Figure S14: Genome length convergence with time.

Figure S15: **Genome length estimation error of RESPECT and CovEst.** The coverage is 1X, and the y-axis is in square-root scale. The sign of error indicates overestimation or underestimation. The dashed lines mark the region that the absolute value of error is less than 5%.

Figure S16: **Estimated to true genome length ratio.** Comparing RESPECT and CovEst over 66 test species with genomes skimmed at 1X coverage. The y-axis is plotted in log scale, and the red dashed line at $y = 1$ is the grand truth (no error). Two genomes (*A. tauschii* (0.002) and *Z. mays* (0.003)) that CovEst had extremely low estimated to true ratios were removed to improve readability.

Figure S17: **Length estimation error on simulated data at different coverages.** The distribution of error made by RESPECT and CovEst in estimating the length of 66 test genomes skimmed at 0.5X, 1X, 2X, and 4X coverage. The y-axis is plotted in log scale.

Figure S22: Distributions of intra-generic versus inter-generic differences in HCRM for pairs of RefSeq species. (a) Plants. (b) Invertebrates. (c) Mammals. (d) Other vertebrates.

Figure S23: **High copy repeats per million versus uniqueness ratio among genomes with and without known recent WGD events.** HCRM values are computed directly from the genome assemblies.

Species	Common name	Run accession	Genome length (RefSeq)	Duplication	Contamination	Total	RESPECT error	CovEst error
<i>Pelecanus crispus</i>	Dalmatian pelican	SRR959397	1160924693	4%	3%	7%	-4%	-42%
<i>Zea mays</i>	Maize	SRR2960981	2135083061	1%	7%	8%	-8%	-90%
<i>Leptosomus discolor</i>	Cuckoo roller	SRR956935	1136244952	6%	3%	9%	4%	-17%
<i>Merops nubicus</i>	Northern carmine bee-eater	SRR958515	1062961556	6%	3%	9%	4%	-13%
<i>Nicotiana tabacum</i>	Cultivated tobacco	SRR955758	3643471356	7%	3%	10%	3%	-85%
<i>Aegilops tauschii</i>	Tausch's goatgrass	SRR5170323	4327321625	5%	7%	12%	-11%	-91%
<i>Frankliniella occidentalis</i>	Western flower thrips	SRR1300141	274989634	7%	5%	12%	0%	-2%
<i>Mustela putorius</i>	Domestic ferret	SRR085103	2410879678	9%	4%	13%	-60%	-61%
<i>Salmo salar</i>	Atlantic salmon	SRR1264544	2966890203	7%	8%	15%	-5%	-89%
<i>Ceratitis capitata</i>	Mediterranean fruit fly	SRR847379	436490799	10%	7%	17%	-3%	-40%
<i>Athene cunicularia</i>	Burrowing owl	SRR6670174	1157069330	15%	4%	19%	-5%	-6%
<i>Manis javanica</i>	Sunda pangolin	SRR3929782	2547395906	7%	16%	23%	-2%	-58%
<i>Elephantulus edwardii</i>	Cape sengi	SRR387354	3843982861	9%	15%	24%	18%	-90%
<i>Propithecus coquereli</i>	Coquerel's sifaka	SRR1657020	2798152141	7%	20%	27%	-41%	-54%
<i>Perca flavescens</i>	Yellow perch	SRR8482300	877456336	1%	28%	29%	3%	-69%
<i>Fopius arisanus</i>	A Braconid wasp	SRR1560668	153631861	28%	2%	30%	-2%	-43%
<i>Sturnus vulgaris</i>	Common starling	SRR2240710	1036755994	15%	20%	35%	7%	-44%
<i>Charadrius vociferus</i>	Killdeer	SRR944000	1219859583	37%	1%	38%	-3%	-84%
<i>Linepithema humile</i>	Argentine ant	SRR059226	219500750	10%	29%	39%	-8%	-34%
<i>*Cyprinodon variegatus</i>	Sheepshead minnow	SRR1261820	1035184475	24%	19%	43%		
<i>*Bombus impatiens</i>	Eastern bumble bee	SRR1575221	246856484	12%	32%	44%		
<i>*Priapulus caudatus</i>	Penis worm	SRR649590	511738253	18%	27%	45%		
<i>*Diaphorina citri</i>	Citrus psylla	SRR189236	485705082	39%	11%	50%		
<i>*Drosophila miranda</i>	Fruit fly	SRR789694	136728780	49%	12%	61%		
<i>*Musca domestica</i>	Housefly	SRR650115	750403944	12%	50%	62%		
<i>*Stegastes partitus</i>	Bicolor damselfish	SRR649426	800491834	34%	41%	75%		
<i>*Bos mutus</i>	Wild yak	SRR361227	2645161911	64%	12%	76%		
<i>*Populus euphratica</i>	Desert poplar	SRR616245	496032534	94%	4%	98%		
<i>*Carrito syricta</i>	Philippine tarsier	SRR3502922	3453864774	59%	82%	141%		

Table S1: **SRA preprocessing results.** The percentage of identified duplicate and contaminant reads are provided. Samples with more than 40% total duplication plus contamination (marked with an asterisk) were discarded.

Table S2: List of species with recent WGD events.

Species	UR	HCRM
Amborella trichopoda	0.85	48
Ananas comosus	0.71	113
Arabidopsis lyrata	0.71	26
Arabidopsis thaliana	0.91	14
Asparagus officinalis	0.51	60
Beta vulgaris	0.74	77
Brachypodium distachyon	0.83	5
Brassica napus	0.40	25
Brassica oleracea	0.71	21
Brassica rapa	0.80	112
Cajanus cajan	0.67	38
Cannabis sativa	0.51	67
Carica papaya	0.81	21
Chenopodium quinoa	0.50	165
Chlamydomonas reinhardtii	0.85	242
Chondrus crispus	0.51	6
Cicer arietinum	0.66	106
Citrus clementina	0.75	41
Citrus sinensis	0.68	154
Coccomyxa subellipsoidea	0.96	27
Cucumis melo	0.84	49
Cucumis sativus	0.82	26
Daucus carota	0.66	35
Dendrobium catenatum	0.79	30
Eucalyptus grandis	0.69	50
Fragaria vesca	0.82	57
Glycine max	0.65	107
Glycine soja	0.63	95
Gossypium raimondii	0.76	13
Ipomoea nil	0.55	547
Ipomoea triloba	0.67	193
Jatropha curcas	0.79	55
Juglans regia	0.60	45
Lactuca sativa	0.53	36
Lupinus angustifolius	0.73	60
Malus domestica	0.65	38

Continued on next page

Table S2 – *Continued from previous page*

Species	UR	HCRM
Manihot esculenta	0.67	17
Medicago truncatula	0.77	57
Morus notabilis	0.72	57
Musa acuminata	0.79	85
Nelumbo nucifera	0.82	27
Nicotiana attenuata	0.49	46
Nicotiana glauca	0.56	41
Nicotiana glauca	0.57	20
Nicotiana glauca	0.60	37
Nymphaea colorata	0.78	152
Olea europaea	0.65	40
Oncorhynchus kisutch	0.52	370
Oncorhynchus mykiss	0.59	221
Oncorhynchus nerka	0.68	182
Oncorhynchus tshawytscha	0.57	92
Oryza sativa	0.75	44
Panicum hallii	0.56	28
Papaver somniferum	0.44	18
Phalaenopsis equestris	0.72	25
Phoenix dactylifera	0.77	21
Physcomitrella patens	0.73	47
Populus trichocarpa	0.77	49
Prunus avium	0.68	155
Prunus mume	0.77	60
Prunus persica	0.75	69
Punica granatum	0.72	42
Pyrus x bretschneideri	0.54	133
Quercus lobata	0.63	62
Quercus suber	0.62	60
Ricinus communis	0.73	69
Rosa chinensis	0.61	70
Salmo salar	0.48	356
Salmo trutta	0.55	330
Salvelinus alpinus	0.69	185
Selaginella moellendorffii	0.42	39
Sesamum indicum	0.84	12

Continued on next page

Table S2 – *Continued from previous page*

Species	UR	HCRM
Setaria italica	0.76	14
Solanum lycopersicum	0.80	39
Solanum pennellii	0.76	39
Solanum tuberosum	0.74	54
Sorghum bicolor	0.52	46
Syzygium oleosum	0.75	48
Theobroma cacao	0.79	14
Vitis vinifera	0.73	44
Volvox carteri	0.80	151
Zea mays	0.32	32
Ziziphus jujuba	0.58	90