

aBayesQR: A Bayesian Method for Reconstruction of Viral Populations Characterized by Low Diversity

SOYEON AHN and HARIS VIKALO

ABSTRACT

RNA viruses replicate with high mutation rates, creating closely related viral populations. The heterogeneous virus populations, referred to as viral quasispecies, rapidly adapt to environmental changes thus adversely affecting efficiency of antiviral drugs and vaccines. Therefore, studying the underlying genetic heterogeneity of viral populations plays a significant role in the development of effective therapeutic treatments. Recent high-throughput sequencing technologies have provided invaluable opportunity for uncovering the structure of quasispecies populations. However, accurate reconstruction of viral quasispecies remains difficult due to limited read lengths and presence of sequencing errors. The problem is particularly challenging when the strains in a population are highly similar, that is, the sequences are characterized by low mutual genetic distances, and further exacerbated if some of those strains are relatively rare; this is the setting where state-of-the-art methods struggle.

In this article, we present a novel viral quasispecies reconstruction algorithm, aBayesQR, that uses a maximum-likelihood framework to infer individual sequences in a mixture from high-throughput sequencing data. The search for the most likely quasispecies is conducted on long contigs that our method constructs from the set of short reads via agglomerative hierarchical clustering; operating on contigs rather than short reads enables identification of close strains in a population and provides computational tractability of the Bayesian method. Results on both simulated and real HIV-1 data demonstrate that the proposed algorithm generally outperforms state-of-the-art methods; aBayesQR particularly stands out when reconstructing a set of closely related viral strains (e.g., quasispecies characterized by low diversity).

Keywords: Bayesian inference, hierarchical clustering, low diversity, viral quasispecies.

1. INTRODUCTION

A NUMBER OF POTENTIALLY LIFE-THREATENING INFECTIOUS DISEASES are caused by RNA viruses, including human immunodeficiency virus (HIV), hepatitis C virus, influenza, and Ebola. RNA viruses have a relatively high mutation rate due to both their error-prone replication process and the lack of

sophisticated repair mechanisms (Duarte et al., 1994). Consequently, they rapidly evolve and exist as a set of nonidentical but closely related genetic variants, known as a viral quasispecies. Viral populations can readily adapt to dynamic environments and develop resistance to antiviral drugs and vaccines, which makes the design of effective and long-lasting treatments for RNA viral diseases exceedingly difficult (Lauring and Andino, 2010). Determining the structure of viral populations helps the understanding of viral diseases and provides guidance in the development of effective medical therapeutics. Quasispecies spectrum reconstruction (QSR) aims to assemble individual haplotype sequences in a population and estimate their prevalence using sequencing reads generated from a sample containing a set of viral variants. High-throughput next-generation sequencing (NGS) technologies have enabled affordable acquisition of data needed to assemble quasispecies. However, relatively short length of the NGS reads and the presence of errors in sequencing data render the QSR problem difficult. The QSR problem is particularly challenging when the strains in a viral population are highly similar, that is, the sequences are characterized by low mutual genetic distances, and further exacerbated if some of those strains are relatively rare (Posada-Cespedes et al., 2016).

Several software tools for solving the QSR problem by analyzing NGS data have been developed in recent years. ShoRAH (Zagordi et al., 2011), the earliest publicly available such software, was developed by combining a path cover-based approach and probabilistic clustering in Eriksson et al. (2008) and Zagordi et al. (2010a), respectively, and applied to analysis of HIV data (Zagordi et al., 2010b). Read-graph approach was the basis for ViSpA (Astrovskaya et al., 2011), developed as a variant of the network flow method proposed in Westbrook et al. (2008). Prosperi et al. (2011) proposed a combinatorial method for QSR and the resulting software, QuRe, was provided by Prosperi and Salemi (2012).

An approach that resulted in the software package PredictHaplo (Prabhakaran et al., 2014) relied on a Dirichlet process mixture model and was developed specifically targeting HIV population reconstruction; QuasiRecomb (Töpfer et al., 2013) is based on a hidden Markov model that explicitly models recombination events. In Schirmer et al. (2012), a benchmarking study that compares the performance of several publicly available quasispecies reconstruction softwares was presented. The study demonstrated that none of the tested methods could reconstruct populations characterized by low pairwise distance between the haplotype sequences. Following this study, other softwares, including HaploClique (Töpfer et al., 2014), based on max-clique enumeration of a read alignment graph, and viral genome assembler (VGA) (Mangul et al., 2014), a graph-coloring based heuristic method, were developed. Most recently, a reference-assisted de novo assembly pipeline, ViQuaS, was proposed in Jayasundara et al. (2014). ViQuaS extends an existing algorithm, QuRe (Prosperi et al., 2011), and outperforms various other techniques on a wide range of data set. However, performance of these more recent methods deteriorates dramatically in the scenarios where the genetic diversity of a population is low (Posada-Cespedes et al., 2016).

Both Posada-Cespedes et al. (2016) and Schirmer et al. (2012) have pointed out that the existing methods for viral quasispecies reconstruction struggle in the scenarios where the populations are characterized by low diversity. This, in part, is due to the presence of relatively long genetic regions that are common to pairs of closely related viral sequences; clearly, this makes distinguishing different strains challenging. The problem becomes even more difficult when the frequency of one (or more) of the close strains is low; in such settings, small genetic distances may be confused for sequencing errors and hence remain undetected. Such failures to detect may have serious consequences in antiviral treatment studies since undetected strains cannot be properly targeted for drug and vaccine design. It has been shown that even the viral strains existing at low frequencies can cause a drug treatment failure due to their resistance to the drug (Le et al., 2009; Simen et al., 2009). Therefore, complete recovery of the composition of viral populations is of critical importance for effective antiviral therapies.

In this article, we propose a novel QSR algorithm, aBayesQR (combining agglomerative hierarchical clustering and Bayesian inference), that overcomes limitations of the existing methods and reliably reconstructs quasispecies characterized by low diversity. The algorithm performs reconstruction of a quasispecies from NGS data in two stages. In the first stage, conflict-free short reads are hierarchically merged and assembled into longer sequences (contigs), which we refer to as super-reads. In the second stage, likelihoods of the probable quasispecies are computed using the assembled super-reads (rather than using the original set of short reads), and the most likely set of viral strains is selected. Note that the super-reads synthesized in the first stage of aBayesQR allow us to distinguish between closely related strains, which share long genetic regions, as well as reduce the search space and enable computational tractability of the Bayesian inference conducted in the second stage.

The second stage of aBayesQR involves sequential pruning of the solution space; in particular, the likely set of partial viral strains comprising n single-nucleotide variants (SNVs) is generated by extending

previously inferred partial viral strains having $n-1$ SNVs. The number of sequences in a set (i.e., the size of a viral population) is dynamically updated at each step by evaluating quality of the set of partially reconstructed viral strains and ultimately precisely inferred at the end of the search process. The relative frequencies of each strain are determined by counting the numbers of reads unambiguously associated with each of the reconstructed strains. Our tests on both simulated and experimental data demonstrate superior performance compared to state-of-the-art methods for quasispecies reconstruction. In particular, it is shown that unlike the competing methods, aBayesQR is capable of detecting and reliably reconstructing viral haplotypes having very small mutual genetic distances.

2. PROPOSED METHOD

Our algorithm for inferring spectrum of a viral population consists of the following two steps: (1) constructing super-reads by hierarchically clustering aligned paired-end reads and (2) inferring the most likely quasispecies from the set of super-reads and estimating the frequencies of the strains in the quasispecies.

2.1. Super-read construction via agglomerative clustering

In the first stage of aBayesQR, paired-end reads uniquely mapped to a reference genome are grouped into super-reads via agglomerative hierarchical clustering. This is facilitated by a weighted graph $G=(\mathcal{V}, \mathcal{E})$, which is constructed and recursively updated as the clustering proceeds. In particular, each vertex of G is associated with a cluster collecting read that originated from a single strain in a quasispecies; we denote the set of reads in the i^{th} cluster (i.e., the cluster associated with the i^{th} vertex) as $V_i = \{v_i^j, j=1, \dots, |V_i|\}$. Let sr_i denote a consensus sequence (i.e., a super-read) constructed from the reads in V_i . The i^{th} and j^{th} vertex of G are connected by an edge $e_{ij} \in \mathcal{E}$ if all the reads in V_i and V_j (or, equivalently, sr_i and sr_j) are conflict free and an overlap criterion, specified later in this subsection, is satisfied. The weight w_{ij} of the edge e_{ij} is a measure of similarity between V_i and V_j at each step, the algorithm merges a pair of vertices connected by the edge having the largest weight to form a new vertex and agglomerates the corresponding clusters.

The alleles at homozygous sites, common to all the components of a quasispecies, are not utilized in the reconstruction procedure. Instead, we separate reads having originated from different strains by clustering them using heterogeneous sites with reliable SNV information. An SNV information is considered reliable if the relative abundance of the allele is above a predetermined threshold, as in Di Giallonardo et al. (2014); alleles whose abundance is below the threshold are treated as sequencing errors and disregarded in the process of clustering. For convenience, let us denote the set of preprocessed paired-end reads by $R = \{r_i, i=1, \dots, |R|\}$. The agglomerative clustering is initialized with $|R|$ clusters, one for each read; in

other words, we start with $V_1 = r_1, \dots, V_{|R|} = r_{|R|}$, implying that $|\mathcal{V}| = \sum_{i=1}^{|\mathcal{V}|} |V_i| = |R|$, and proceed by sequentially merging judiciously chosen pairs of vertices (i.e., agglomerating the corresponding clusters). Intuitively, it is meaningful to reduce the number of vertices in the graph by merging those associated with conflict-free consensus sequences that have a large overlap. To formalize this, let $L_i = \{l_1, \dots, l_{|L_i|}\}$ denote an index set of the SNV positions covered by sr_i , let $L_{i \cap j} = \{l_1, \dots, l_{|L_{i \cap j}|}\}$ be the index set of SNV positions covered by both sr_i and sr_j , and let $L_{i \cup j} = \{l_1, \dots, l_{|L_{i \cup j}|}\}$ be the index set of SNV positions covered by either sr_i or sr_j . Then, the pairs of vertices (i, j) that we consider as candidates for merging and thus connect by an edge are those satisfying either

$$|L_{i \cap j}| \geq \theta \cdot |L_{i \cup j}| \quad \text{or} \quad |L_{i \cap j}| = \min(|L_i|, |L_j|),$$

where the second condition promotes merger of short super-reads, and the choice of θ is discussed below. To quantify uncertainty inherent to a clustering solution due to existence of nonoverlapping positions among the reads in each cluster, we define a position-specific confidence score as follows:

$$score_i[l] = \frac{cr_i[l] - cr[l]}{1 - cr[l]}$$

where l denotes the position, $cr[\cdot]$ is the overall coverage rate, and $cr_i[\cdot]$ denotes cluster-specific coverage rate for V_i (i.e., $cr_i[l]$ is the fraction of reads in $V_i = \{v_i^j, j=1, \dots, |V_i|\}$ covering position l). On the one hand, this score is penalized at a site where the fraction of cluster members (short reads) covering the site is

low; the score is negative if the cluster-specific coverage rate is below the global coverage rate, which implies uncertainty of the clustering decision. On the other hand, positive scores indicate high confidence in the decision to group the reads into the same cluster. Note that the highest possible score of 1 at position l is achieved when all the reads in a cluster cover the l^{th} position. Using the confidence scores, we define the weight w_{ij} assigned to an edge e_{ij} to quantify similarity between V_i and V_j as

$$w_{ij} = \frac{1}{|L_{i \cup j}|} \sum_{l \in L_{i \cup j}} \text{score}_{e_{ij}}[l].$$

Given the weights w_{ij} , we can now specify the clustering procedure. In each step, the pair of vertices connected by the edge with maximum weight is merged; the newly constructed vertex inherits edges from the merged vertices and the weights on those edges are re-evaluated. A new (longer) consensus sequence is constructed by combining the two super-reads associated with the merged vertices; recall that there are no conflicts between the super-reads being merged. If after such an update step no edges connect the new vertex with the rest of the graph (because no inherited edges satisfy the connectivity condition), θ is decreased and the above process is repeated. We initially set θ to 0.9 and gradually decrease it by 0.1 while $\theta > 0$. The above procedure is repeated until no pairs of vertices satisfy the connectivity condition. By that point, a set of long consensus sequences (the final super-reads) has been formed from the clusters of reads associated with the nodes of the final graph. While the complexity of agglomerative clustering is, in general, $O(N^3)$ where N denotes the input data size (Sasirekha and Baby, 2013), it has been shown that its time complexity can be reduced to $O(N^2)$ with accuracy equal to that of the brute-force method by using the partial maximum array technique (Jung and Kim, 2001). We exploit this to efficiently construct super-reads. The algorithm for super-read construction is formalized as Algorithm 1.

Algorithm 1: Agglomerative clustering for super-read construction

Input: Set of reads aligned to the reference genome

Output: Set of super-reads and the corresponding confidence scores

for $\theta > 0$ **do**

 Build a weighted graph $G=(\mathcal{V}, \mathcal{E})$

while $E \neq \emptyset$ **do**

 Merge two clusters connected with the largest weight

 Update $G=(\mathcal{V}, \mathcal{E})$ and weights using partial maximum array

end while

$\theta = \theta - 0.1$

end for

2.2. Maximum likely (MA) reconstruction of quasispecies from super-reads

Here we describe how to reconstruct the most likely set of strains in a viral quasispecies using super-reads from Section 2.1 and their confidence scores. While in principle the method outlined in this section could be applied directly to the short reads provided by a sequencing platform, such an approach would in general not only be computationally prohibitive due to a very large number of short reads but also limit the ability of the algorithm to distinguish strains with small mutual genetic distances due to having long conserved regions. Relying on a relatively small number of long super-reads constructed from short reads circumvents both of these problems and makes the reconstruction more accurate and practically feasible. Note that sequencing errors may undesirably prevent clusters of reads from being merged with other clusters due to a violation of conflict-free requirement; consequently, a set of short reads in a small cluster is likely to have a disproportionate amount of sequencing errors. For this reason, we ignore clusters with very small memberships (in particular, those containing fewer than $0.001 \cdot |R|$ reads), which limits the detection of strains to those constituting more than 0.1% of the quasispecies.

Let $\mathcal{C} = \{C_m, m = 1, \dots, M\}$ denote the collection of clusters that remain after deleting clusters having only few reads; moreover, for convenience let us relabel the reads in C_m as c_m^j , that is, $C_m = \{c_m^j, j = 1, \dots, |C_m|\}$ where $c_m^j \in R$. We organize the super-reads obtained by Algorithm 1 in Section 2.1 into rows of an $M \times N$ matrix $\mathbf{S} = \{s_{mn}, m = 1, \dots, M, n = 1, \dots, N\}$ with entries $s_{mn} \in \{A, C, G, T, -\}$ where $-$ denotes a site

not covered by a super-read and N denotes the total number of SNV sites in the strains of a quasispecies. A nucleotide in the (m, n) position of \mathbf{S} is assigned confidence $score_m[n]$ defined in Section 2.1; the scores for the entire matrix are normalized so that they fall between 0 and 1 to use them in our Bayesian approach to assembly. Let ε_{mn} be the probability that s_{mn} was estimated erroneously due to either a sequencing error in reads on the n^{th} SNV position or the uncertainty induced by reads not covering the n^{th} SNV position. Note that negative scores indicate low confidence resulting from insufficient cluster-specific coverage rate, while positive scores imply relatively confident information. To map $score_m[n] \in (-\infty, 1]$ to the set $[0, 1]$, we set $\varepsilon_{mn} = 1 - e^{score_m[n]}$ for $score_m[n] < \ln(1 - \epsilon)$, where ϵ denotes the error rate of a sequencing platform. Otherwise, we set $\varepsilon_{mn} = \epsilon$.

Let $Q = \{q_k, k = 1, \dots, K\}$ denote the set of K strains of a viral quasispecies. The goal in the second stage of our method is to determine Q from the super-read matrix \mathbf{S} using a probabilistic framework. An exhaustive search over the entire solution space is computationally intractable even for small \mathbf{S} ; instead, we reconstruct the set of K viral strains sequentially, extending partially estimated strains on SNV position at each step. Since maintaining and extending all possible partial strains inevitably increase their number exponentially, unlikely sets of candidate strains are pruned in each step. Each step consists of three basic parts: (a) extension of the partially reconstructed strains, (b) selection of probable sets comprising K strains chosen among those generated in step (a), and (c) evaluation of the quality of the selected sets of strains and an update of K . The sequential Bayesian inference procedure in step t is illustrated in Supplementary Fig. S1 in Appendix A.

2.2.1. Extending partially reconstructed strains. Let $F_{1:t-1} = \{f_{1:t-1}^i, i = 1, \dots, |F_{1:t-1}|\}$ be the collection of partially reconstructed strains covering the first $t-1$ SNV sites and let $B_t = \{b_t^j, j = 1, \dots, |B_t|\}$ be the lists of distinct bases in the t^{th} column of \mathbf{S} , where $b_t^j \in \{A, C, G, T\}$ and $2 \leq |B_t| \leq 4$. Then, all the possible extensions of $f_{1:t-1}^i$ to the SNV site t can be enumerated as $\{[f_{1:t-1}^i, b_t^1], \dots, [f_{1:t-1}^i, b_t^{|B_t|}]\}$. Let $S_{1:t-1}^i = \{s_{1:t-1}^{i,c'}, c' = 1, \dots, |S_{1:t-1}^i|\}$ be the collection of super-reads covering some of the first t SNV sites, which are consistent with $f_{1:t-1}^i$ (ignoring “-” in $s_{1:t-1}^{i,c'}$) where $\{i_{c'}\}$ denotes indices of rows of \mathbf{S} that are placed in $S_{1:t-1}^i$, and let $S_t^i = \{s_t^{i,c}, c = 1, \dots, |S_t^i|\}$ denote the collection of nucleotides ($s_t^{i,c} \in \{A, C, G, T\}$, not “-”) observed at the t^{th} SNV site of the super-reads in $S_{1:t-1}^i$ where $\{i_c\}$ denotes the indices of rows in \mathbf{S} that contribute to S_t^i . Given $S_{1:t-1}^i, S_t^i$ and $f_{1:t-1}^i$, the probability of b_t^j being the true extension of $f_{1:t-1}^i$ is given by the following:

$$P(S_t^i | b_t^j, S_{1:t-1}^i, f_{1:t-1}^i) = \prod_{c=1}^{|S_t^i|} P(s_t^{i,c} | b_t^j),$$

$$P(s_t^{i,c} | b_t^j) = \begin{cases} 1 - \varepsilon_{i_c t}, & \text{if } b_t^j = s_t^{i,c}, \\ \frac{\varepsilon_{i_c t}}{|B_t|}, & \text{otherwise.} \end{cases}$$

We extend $f_{1:t-1}^i$ to $[f_{1:t-1}^i, b_t^j] \in F_{1:t-1,t}$ by appending the $b_t^j \in B_t$ which satisfies $\frac{P(S_t^i | b_t^j, S_{1:t-1}^i, f_{1:t-1}^i)^{\frac{1}{|S_t^i|}}}{\sum_{B_t} P(S_t^i | b_t^j, S_{1:t-1}^i, f_{1:t-1}^i)^{\frac{1}{|S_t^i|}}} \geq \delta_0$,

where the exponent ensures proper normalization and is needed since the number of super-reads, $|S_{1:t-1}^i|$, varies for each $\{f_{1:t-1}^i, i = 1, \dots, |F_{1:t-1}|\}$. For $f_{1:t-1}^i$, which has no matched super-reads, that is, $|S_{1:t-1}^i| = 0$, we keep all of $|B_t|$ possible extensions of $f_{1:t-1}^i$. By collecting probable extensions for each $f_{1:t-1}^i \in F_{1:t-1}$, we obtain the set of partial strains stretching over the first t SNV sites, $F_{1:t-1,t}$. This procedure is formalized as function *ExtendFrag* in Appendix A.

2.2.2. Inferring likely sets of K partial strains. Having generated the probable partial strain, $F_{1:t-1,t}$, we denote the set of all its possible subsets of K strains (i.e., the quasispecies population candidates) as $\mathcal{Q}_{1:t-1,t} = \left\{ Q_{1:t-1,t}^i, i = 1, \dots, \binom{|F_{1:t-1,t}|}{K} \right\}$ where $Q_{1:t-1,t}^i = \{q_{kn}^i, k = 1, \dots, K, n = 1, \dots, t\}$ and $q_{kn}^i \in F_{1:t-1,t}$. The log-likelihoods of $Q_{1:t-1,t}^i$ can be expressed as follows:

$$\ln P(\mathbf{S} | Q_{1:t-1,t}^i) = \sum_{m=1}^M \ln P(s_m | Q_{1:t-1,t}^i),$$

$$P(s_m | Q_{1:t-1,t}^i) = \frac{1}{K} \left(\sum_{k=1}^K \left(\prod_{n=1}^t P(s_{mn} | q_{kn}^i) \right) \right),$$

where s_m denotes the m^{th} row vector of the matrix of super-reads \mathbf{S} and

$$P(s_{mn} | q_{kn}^i) = \begin{cases} 1 - \epsilon_{mn}, & \text{if } q_{kn}^i = s_{mn}, \\ \frac{\epsilon_{mn}}{|B_n|}, & \text{if } q_{kn}^i \neq s_{mn} \text{ for } s_{mn} \neq -. \end{cases}$$

Let $Q_{1:t}^{\max} = \max_{Q_{1:t-1,t}^i \in \mathcal{Q}_{1:t-1,t}} P(\mathbf{S} | Q_{1:t-1,t}^i)$. Among the $\binom{|F_{1:t-1,t}|}{K}$ sets in $\mathcal{Q}_{1:t-1,t}$, we keep only those that satisfy $P(\mathbf{S} | Q_{1:t-1,t}^i) > \delta_1 \cdot Q_{1:t}^{\max}$, while the others are discarded; let us denote the collection of candidate sets that pass this test as $\mathcal{Q}_{1:t}$. For practical feasibility of the scheme, the collection of partially reconstructed strains $F_{1:t-1,t}$ is trimmed by excluding from it all the strains that are not part of at least one of the sets in $\mathcal{Q}_{1:t}$; we denote the resulting collection of partial strains by $F_{1:t} \in F_{1:t-1,t}$ and use it when extending the strains onto the $t+1$ SNV site. The described procedure is formalized as function *InferQuasi* in Appendix A.

2.2.3. Determining the number of strains K in a quasispecies. In this step, we assess appropriateness of K used in the inference of $\mathcal{Q}_{1:t}$ and update it if necessary. To this end, we rely on the minimum error correction (MEC) score, which has previously been broadly used as a criterion in the design of methods for haplotype assembly (Lancia et al., 2001; Lippert et al., 2002). In the context of polyploid haplotype assembly, the MEC score is defined as the smallest number of nucleotides that needs to be changed in data (i.e., in observed reads) so that the corrected reads are consistent with having originated from K haplotypes. Let $HD_t(\cdot, \cdot)$ denote the Hamming distance between two sequences counted over the observed nucleotides in the first t SNV positions.* Then, the MEC score of the most likely set $Q_{1:t}^{\max}$ of K viral strains evaluated on the first t SNVs is

$$\text{MEC}_t(K) = \sum_{m=1}^M \min_{k \in \{1, \dots, K\}} \sum_{j=1}^{|C_m|} HD_t(c_m^j, q_k^{\max}),$$

where q_k^{\max} is the k^{th} row vector of $Q_{1:t}^{\max}$. Let N_t be the total number of nucleotides observed in the first t SNV positions of all the reads of the data set. Note that the smaller the MEC scores, the higher the accuracy of a clustering. If $\text{MEC}_t(K)/N_t < 2\epsilon$, we use the same value K in the next step where the likely set of viral strains stretching over the first $t+1$ SNV positions is inferred. Otherwise, we increase K by 1, repeat the estimation of $\mathcal{Q}_{1:t}$, and evaluate the improvement rate of MEC score as follows:

$$\text{MECimpr}(K) = \frac{\text{MEC}_t(K) - \text{MEC}_t(K+1)}{\text{MEC}_t(K)}.$$

The reason for selecting K based on the MEC improvement rate (*MECimpr*) is that the MEC score drops significantly once K matches the actual number of clusters; our scheme attempts to detect that change to infer population size. If $\text{MECimpr}(K) > \eta$, where η denotes a prespecified threshold, the number of species is updated as $K \leftarrow \min\{K+n, |F_{1:t-1,t}|\}$ where n is the smallest integer number such that $\text{MECimpr}(K+n) < \eta$. If $\text{MECimpr}(K) < \eta$, we update the number of species as $K \leftarrow \max\{K-n, 2\}$ where n is the smallest integer such that $\text{MECimpr}(K-n) \geq \eta$. The choice of threshold η is discussed in Appendix B. The updated value of K is used for the inference of $\mathcal{Q}_{1:t+1}$. Note that the probable set of viral strains, $\mathcal{Q}_{1:t}$, is stored for each K to avoid performing redundant *MECimpr*(\cdot) calculations.

Once we obtain the most likely set of K viral sequences covering N SNVs, $Q_{1:N}^{\max}$, the full-length K quasispecies strains are reconstructed by inserting the consensus nucleotides observed in R into the non-SNV sites. We estimate relative frequencies p_k , $1 \leq k \leq K$, of quasispecies strains based on the Hamming distance between super-reads and the reconstructed sequences. In particular, for each super-read sr_i , we determine the nearest assembled strain q_j where $j = \arg \min_{k \in \{1, \dots, K\}} HD(sr_i, q_k)$, and the number of reads

*If either of the two sequences has a gap “-” in a position, that position is ignored in the computation of the aforementioned Hamming distance.

involved in constructing the super-read sr_i is counted toward p_j . The entire scheme proposed in this subsection is summarized as Algorithm 2.

Algorithm 2: Sequential Bayesian inference for quasispecies reconstruction

Input: Set of super-reads and the corresponding confidence scores

Output: Set of K strains of a viral quasispecies

Initial $K \leftarrow 2$, $F_{1:t} \leftarrow B_1$

for $t \in \{2, \dots, N\}$ **do**

$F_{1:t-1,t} = \text{ExtendFrag}(F_{1:t-1}, t, \delta_0)$

$\mathcal{Q}_{1:t} = \text{InferQuasi}(F_{1:t-1,t}, K, \delta_1)$

$K^* \leftarrow K$, $\mathcal{Q}_{1:t}^* \leftarrow \mathcal{Q}_{1:t}$

if $\text{MEC}_t(K)/N_t \geq 2\epsilon$ and $K < |F_{1:t-1,t}|$ **do**

$\mathcal{Q}_{1:t} = \text{InferQuasi}(F_{1:t-1,t}, K+1, \delta_1)$

if $\text{MECimpr}(K) < \eta$ **do**

while $\text{MECimpr}(K) < \eta$ and $K > 2$

$\mathcal{Q}_{1:t}^* \leftarrow \mathcal{Q}_{1:t}$, $K^* \leftarrow K$, $K \leftarrow K-1$

$\mathcal{Q}_{1:t} = \text{InferQuasi}(F_{1:t-1,t}, K, \delta_1)$

end while

else do

while $\text{MECimpr}(K) \geq \eta$ and $K < |F_{1:t-1,t}|$

$\mathcal{Q}_{1:t}^* \leftarrow \mathcal{Q}_{1:t}$, $K^* \leftarrow K$

$\mathcal{Q}_{1:t} = \text{InferQuasi}(F_{1:t-1,t}, K+1, \delta_1)$

end while

end if

end if

$K \leftarrow K^*$, $\mathcal{Q}_{1:t} \leftarrow \mathcal{Q}_{1:t}^*$

Get $F_{1:t}$ by pruning $F_{1:t-1,t}$ based on $\mathcal{Q}_{1:t}$

end for

Reconstruct full-length quasispecies Q from $Q_{1:N}^{max} \in \mathcal{Q}_{1:t}$ and R

Estimate frequencies of each strain $q_k \in Q$ based on $HD(sr_i, q_k)$ and $|C_i|$

3. RESULTS AND DISCUSSION

3.1. Performance comparison on simulated data

To evaluate performance of the proposed method for quasispecies reconstruction, we use metrics *Recall*, *Precision*, *Predicted Proportion*, and *Reconstruction Rate*. *Recall* is defined as the ratio of the number of correctly reconstructed strains to the total number of true strains in the quasispecies, that is, $\text{Recall} = \frac{TP}{TP + FN}$, while *Precision* is defined as the fraction of correctly reconstructed strains among all the assembled sequences, that is, $\text{Precision} = \frac{TP}{TP + FP}$. Noting that *Precision* usually reports high scores when the number of strains is underestimated while penalizing overestimation of the population size, we also report the ratio of the number of reconstructed sequences to the true population size, *Predicted Proportion*. The closer the *Predicted Proportion* to 1, the more accurate the number of reconstructed strains. Moreover, to assess the degree of reconstruction accuracy, we define *Reconstruction Rate* = $\frac{1}{K} \sum_{k=1}^K \left(1 - \frac{HD(q_k, \hat{q}_k)}{G}\right)$, where G is the length of a genome, K is the number of strains in a quasispecies, and q_k and \hat{q}_k denote the k^{th} true strain and its nearest sequence among the K estimated ones, respectively. To assess the accuracy of estimated frequencies, we use the Jensen–Shannon divergence (JSD), which quantifies similarity between two distributions. Given a true distribution P and its approximation Q , the Kullback–Leibler (KL) divergence $D(P||Q) = \sum_{i=1}^n P(i) \log \frac{P(i)}{Q(i)}$ is undefined when $Q(i) = 0$. JSD, a symmetrized and smoothed version of the KL divergence, circumvents this problem by defining similarity of P and Q as $\text{JSD}(P||Q) = \frac{1}{2}D(P||M) + \frac{1}{2}D(Q||M)$, where M is defined as $M = \frac{1}{2}(P + Q)$.

We compare our algorithm with publicly available ShoRAH (Zagordi et al., 2011), PredictHaplo (Prabhakaran et al., 2014), and ViQuaS (Jayasundara et al., 2014). Since ViQuaS is an extension of the

algorithm in Prosperi et al. (2011) and Prosperi and Salemi (2012), and was shown to have superior performance compared to its predecessor, we omit the comparison with the software QuRe in Prosperi et al. (2011) and Prosperi and Salemi (2012). It is worth pointing out that for the synthetic data sets we study, ShoRAH could not reconstruct strains in the regions where the simulated sequencing coverage is relatively low compared to the average, resulting in reconstruction of strains that are shorter than the true length G . To facilitate a fair comparison with ShoRAH, we aligned its reconstructed strains to the reference genome and completed missing sites with bases from the reference. ViQuaS, on the contrary, tends to reconstruct many more strains than actually present; thus, we followed ViQuaS's authors' recommendation and retained only those having frequencies greater than f_{min} when calculating *Precision*. Finally, not all of the synthetic data sets could be processed with PredictHaplo, preventing us from reporting its performance in some of the scenarios.

We generated synthetic data sets by emulating high-throughput sequencing of a viral population consisting of a number of closely related viral genomes having length of 1300 bp; this particular length was chosen to coincide with the longest region of the HIV *pol* gene. Quasispecies sequences are generated by introducing independent mutations at uniformly random locations along the length of a randomly generated reference genome so as to obtain a predefined level of diversity (*div%*), that is, a predefined average Hamming distance between quasispecies strains. Simulating Illumina's MiSeq data, 2×250 bp-long paired-end reads are sampled uniformly from each viral strain with a mean coverage of $cov \times$ per strain. Inserts of the paired-end reads are on average 150 bp long with standard deviation of 30. In our benchmarking tests, we focus on exploring the effects of diversity (*div%*) on the accuracy of the quasispecies reconstruction.

Two sets of viral populations are considered: (1) a mix of 5 viral strains with abundance levels 50%, 30%, 15%, 4%, and 1%; and (2) a mix of 10 strains with abundance levels 36%, 24%, 16%, 8%, 5.5%, 4%, 3%, 2%, 1%, and 0.5%. Note that the abundances are chosen to approximately follow geometric distribution and that the populations include low-abundant strains. For each combination of the parameters, 100 data sets were generated and the reported results were obtained by averaging over those data instances. For PredictHaplo, which did not produce results in each instance, the averaged results are reported if more than 50 instances were successfully processed.

In all of the following experiments, potential SNVs are called if their abundance is higher than 1%, which is set relatively high to avoid false positives (FPs); FPs prevent reads to be merged with existing clusters in Section 2.1. We execute the function *ExtendFrag* with parameter $\delta_0 = 0.1$. Parameter δ_1 in function *InferQuasi* is initially set to 0.001, but adaptively increases if the number of combinations of partially reconstructed strains exceeds 10,000; this is done to limit the number of likelihood calculations performed in each run of *InferQuasi*. The recommended value of η , a threshold used to determine population size K based on *MECimpr*(\cdot), is discussed in Appendix B.

We compare performances of aBayesQR, ShoRAH, ViQuaS, and PredictHaplo when applied to the reconstruction of a quasispecies spectrum with diversity levels varying between 1% and 5% (i.e., $div \in \{1\%, 2\%, 3\%, 4\%, 5\%\}$). To test the ability of different methods to reconstruct quasispecies with low diversity, we assume low-sequencing error rate of $err = 0.1\%$ [median mismatch error rates for 454 Life Sciences and Illumina platforms are 0.1% and 0.12%, respectively (Archer et al., 2012)]. Coverage per strain $cov \times$ is set to $500 \times$, implying total coverage of $2500 \times$ and $5000 \times$ for the 5-strain and 10-strain population, respectively; strains having frequencies 0.23% or higher in the 5-strain case and those with frequencies 0.46% or higher in the 10-strain case are covered with probability 0.99 (Eriksson et al., 2008).

Table 1 demonstrates that the proposed aBayesQR algorithm outperforms existing schemes. In terms of *Recall* and *Precision*, aBayesQR exhibits exceptionally good performance compared to competing methods when reconstructing quasispecies strains with diversity $div < 4\%$. The performance of ViQuaS deteriorates at low diversities in terms of most of the criteria (i.e., *Recall*, *Precision*, *Predicted Proportion*, and *JSD*). PredictHaplo could not perform reconstruction in most of the low diversity instances, yet it overall achieves the highest *Precision* because it typically underestimates the number of strains as shown by *Predicted Proportion* (e.g., estimating only 2–3 out of 10 strains), which is in agreement with the results reported by a previous study (Schirmer et al., 2012).

Among all methods, ShoRAH has the lowest performance in terms of *Recall* and *Precision*. As indicated by *Predicted Proportion*, aBayesQR is the most accurate method in terms of estimating the population size although it often misses a strain with the lowest frequency when applied to reconstruction of a quasispecies consisting of 10 strains. ViQuaS and ShoRAH typically overestimate the number of strains especially at low diversity levels. aBayesQR is the best method in terms of *Reconstruction Rate* at all levels of diversity.

TABLE 1. PERFORMANCE COMPARISON OF DIFFERENT METHODS FOR VARIED DIVERSITIES (*div*) ON SIMULATED DATA

div(%)	5 strains					10 strains				
	1	2	3	4	5	1	2	3	4	5
Recall										
aBayesQR	0.7080	0.7120	0.6840	0.6560	0.6320	0.5810	0.6380	0.6080	0.5860	0.5550
ShoRAH	0.1920	0.1600	0.1300	0.1060	0.0780	0.0150	0.0380	0.0740	0.0640	0.0930
ViQuaS	0.3700	0.5240	0.6040	0.6360	0.5960	0.0980	0.1700	0.3730	0.4720	0.5050
PredictHaplo	—	—	—	0.6918	0.6808	—	—	0.1021	0.1550	0.2010
Precision										
aBayesQR	0.7113	0.7130	0.6826	0.6447	0.6319	0.6210	0.6881	0.6610	0.6373	0.6140
ShoRAH	0.1062	0.1418	0.1240	0.1078	0.0790	0.0050	0.0170	0.0498	0.0506	0.0824
ViQuaS	0.1960	0.3206	0.4559	0.4982	0.5298	0.0485	0.1079	0.2973	0.4690	0.5596
PredictHaplo	—	—	—	0.9373	0.8822	—	—	0.4509	0.6000	0.6833
PredProp										
aBayesQR	1.0180	1.0120	1.0120	1.0360	1.0140	0.9680	0.9440	0.9240	0.9240	0.9100
ShoRAH	1.9660	1.2200	1.0780	1.0000	1.0180	3.2000	2.9100	1.6710	1.3520	1.1860
ViQuaS	2.1100	1.7220	1.4080	1.3340	1.2180	2.0860	1.8580	1.5450	1.2320	1.0730
PredictHaplo	—	—	—	0.7388	0.7737	—	—	0.1947	0.2430	0.2890
ReconRate										
aBayesQR	0.9990	0.9982	0.9971	0.9961	0.9953	0.9975	0.9967	0.9952	0.9942	0.9924
ShoRAH	0.9948	0.9903	0.9891	0.9851	0.9827	0.9941	0.9900	0.9899	0.9897	0.9911
ViQuaS	0.9963	0.9949	0.9917	0.9936	0.9897	0.9944	0.9910	0.9899	0.9881	0.9858
PredictHaplo	—	—	—	0.9906	0.9896	—	—	0.9850	0.9797	0.9747
JSD										
aBayesQR	0.0022	0.0008	0.0008	0.0014	0.0008	0.0043	0.0026	0.0023	0.0023	0.0025
ShoRAH	0.0762	0.0174	0.0047	0.0009	0.0012	0.1390	0.1110	0.0422	0.0238	0.0109
ViQuaS	0.0651	0.0255	0.0222	0.0097	0.0180	0.0993	0.0747	0.0495	0.0469	0.0454
PredictHaplo	—	—	—	0.1020	0.1036	—	—	0.1971	0.1636	0.1312

Performance comparison of aBayesQR, ShoRAH, ViQuaS, and PredictHaplo in terms of *Recall*, *Precision*, *Predicted Proportion* (*PredProp*), *Reconstruction Rate* (*ReconRate*) and *JSD* on the simulated data with $err=0.1\%$ and $cov=500\times$ versus *div* for a mixture of 5 and 10 viral strains. Averaged PredictHaplo results are reported if they provide answers for more than 50% of data sets. Boldface values indicate the best performance for each *div*(%).

JSD, Jensen–Shannon divergence.

In terms of frequency estimation, aBayesQR overall outperforms all the other methods, whereas PredictHaplo shows the highest *JSD* due to its drawback of underestimating the number of strains. Note that both ViQuaS and ShoRAH exhibit significantly increased (i.e., deteriorated) *JSD* at low diversity levels. This fact, along with the low *Recall* and *Precision* scores they have in low diversity settings, indicates that state-of-the-art methods experience major difficulties when attempting to reconstruct viral quasispecies in those settings, as also observed in Eriksson et al. (2008), Schirmer et al. (2012), and Jayasundara et al. (2014).

We further study the effects that sequencing error rate ($err\%$) and coverage per strain ($cov\times$) have on the performance of the algorithms. Those results are reported in Supplementary Tables S2 and S3 in Appendix C, demonstrating superiority of aBayesQR compared to the competing methods. The runtimes of the tested algorithms are shown in Supplementary Table S4 in Appendix C.

3.2. Performance comparison on real HIV data

To further test the performance of our proposed method, we use it for the analysis of the HIV 5-virus-mix data set published in Di Giallonardo et al. (2014). Specifically, we apply our algorithm to reconstruct an in vitro generated quasispecies population consisting of five known HIV-1 strains: HIV-1_{HXB2}, HIV-1_{89.6}, HIV-1_{JR-CSF}, HIV-1_{NL4-3}, and HIV-1_{YU2}. Compared to the simulated data set, relative frequencies of the five HIV-1 strains are more evenly distributed (about 10%–30%) and the pairwise distances between strains are higher (2.61%–8.45%) (Di Giallonardo et al., 2014). We use the 2×250 bp-long paired-end

reads provided by Illumina's MiSeq Benchtop Sequencer. The reads are aligned to the HIV-1_{HXB2} reference genome; the reads shorter than 150nt and those having bases with quality scores less than a Phred threshold of 60 are discarded. We compare the performance of our method applied to gene-wise quasispecies reconstruction of the above described HIV data with that of the competing techniques. Since the current version of ViQuaS software does not support specifying genomic regions, we could not use it in this experiment. When running aBayesQR, we set the parameter η to 0.09 (the setting recommended in Appendix B). Other parameters are set to the same values as the ones used in Section 3.1.

We evaluate and report the *Predicted Proportion* (i.e., the fraction of correctly estimated strains as defined in Section 3.1) and *Reconstruction Rate* in Table 2. On this real HIV-1 data set that (as pointed above) has different properties than the simulated data set in Section 3.1, aBayesQR is the most accurate among the considered methods in terms of *Predicted Proportion*. PredictHaplo underestimates the population size and reconstructs three or four strains in the 8 considered genes, and ShoRAH greatly overestimates the population size for all 13 genes of the HIV-1 data set (e.g., it reconstructs 119 strains in gp120), which is consistent with our simulation results as well as with the results in Schirmer et al. (2012). aBayesQR and PredictHaplo are tied for the number of genes where all the strains are perfectly reconstructed (five each); for the remaining genes, PredictHaplo provides a larger number of perfectly reconstructed strains. However, it is worth pointing out that PredictHaplo, designed for identification of HIV haplotypes, missed at least one strain in each of the remaining eight genes, while aBayesQR reconstructed most of the strains on all but two genes, gp120 and gp41. ShoRAH did not perfectly reconstruct any of the 13 genes, which is consistent with the simulation results. Moreover, overestimating the number of strains negatively affects the accuracy of ShoRAH's frequency estimation; for instance, the sum of frequencies corresponding to the most abundant five strains does not exceed 50% in 9 of 13 genes (71% is the largest such sum, on *vpu*) (see Supplementary Table S5 in Appendix C).

To complement the gene-wise quasispecies reconstruction study with that of a global reconstruction, we consider the HIV-1 *gap-pol* region spanning 4307 bp. To efficiently process 355,241 paired-end reads that remain after applying a quality filter, we organize the region into a sequence of windows of length 400 bp

TABLE 2. PERFORMANCE COMPARISONS ON A REAL HIV-1 5-VIRUS-MIX DATA SET

	<i>p17</i>	<i>p24</i>	<i>p2-p6</i>	<i>PR</i>	<i>RT</i>	<i>RNase</i>	<i>int</i>	<i>vif</i>	<i>vpr</i>	<i>vpu</i>	<i>gp120</i>	<i>gp41</i>	<i>nef</i>
aBayesQR													
PredProp	1	1	1	1	1	1	1	1	12	1	0.8	0.8	1.2
RR _{HXB2}	100	99.4	100	100	98.5	100	99.9	100	100	99.6	98	0	95.8
RR _{89.6}	100	98.7	100	100	98.6	100	100	100	100	92	96.5	98.9	95.5
RR _{JR-CSF}	100	99.6	100	100	99	100	100	100	100	98.8	97.7	99.1	98.2
RR _{NL4-3}	100	100	100	100	98.9	100	100	99.8	100	100	96.3	98.8	100
RR _{YU2}	100	99.7	100	100	99.2	100	99.5	99.7	100	100	0	98.6	99.2
ShoRAH													
PredProp	13	16.4	13.8	8.8	21.8	11.8	13.6	12.8	7.8	4	23.8	19.8	17.4
RR _{HXB2}	100	96.7	100	100	98.2	100	97.5	100	100	100	97.7	98.4	98.2
RR _{89.6}	100	99.7	100	100	98.6	100	98.9	99.8	100	93.6	96.1	98.6	98.9
RR _{JR-CSF}	100	100	100	100	99.8	96.4	99	100	100	98	96.9	96.3	94.7
RR _{NL4-3}	100	99.1	97.3	100	98.9	99.2	99.3	99.3	100	100	96.1	98.5	98.6
RR _{YU2}	94.2	99	100	98.3	98	94.5	98.6	95	93.2	90.8	97	95.4	97.9
PredictHaplo													
PredProp	1	0.6	1	1	1	0.8	0.8	0.8	1	0.8	0.8	0.8	0.8
RR _{HXB2}	100	0	100	100	100	98.9	100	100	100	93.17	0	0	0
RR _{89.6}	100	100	100	100	100	100	99.8	100	100	0	97.8	100	98.87
RR _{JR-CSF}	100	100	100	100	100	100	100	100	100	100	99.7	100	100
RR _{NL4-3}	100	99.1	100	100	100	100	100	100	100	100	100	100	100
RR _{YU2}	100	0	100	100	100	0	0	0	100	100	98.6	100	100

Predicted Proportion (PredProp) and reconstruction rate (RR) for aBayesQR, ShoRAH, and PredictHaplo applied to reconstruction of HIV-1_{HXB2}, HIV-1_{89.6}, HIV-1_{JR-CSF}, HIV-1_{NL4-3}, and HIV-1_{YU2} for all 13 genes of the HIV-1 data set. (Note: RR expressed in percentages.) Boldface values indicate the genes where all the strains are perfectly reconstructed. The inferred frequencies are shown in Supplementary Table S5 in Appendix C.

where the consecutive windows overlap by 150 bp and run aBayesQR on those windows. The entire region is assembled by connecting strains in the consecutive windows while testing consistency in the overlapping intervals. The number of strains retrieved in the global reconstruction is decided by majority voting of the number of strains obtained in each window. The frequencies are estimated by counting reads nearest (in terms of Hamming distance) to each of the reconstructed strains. Following this procedure, both aBayesQR and PredictHaplo could reconstruct all five HIV-1 strains in the *gap-pol* region correctly, that is, they both achieved *Reconstruction Rate* of 100 for all five strains and *Predicted Proportion* of 1. The frequencies estimated by aBayesQR are 15.21%, 19.34%, 25.56%, 27.61%, and 12.27%, while those estimated by PredictHaplo are 13.21%, 13.56%, 25.67%, 19.69%, and 27.86%. ShoRAH highly overestimated the number of strains and reported *Predicted Proportion* of 41.8; its five most abundant strains estimated are reported to have frequencies 8.51%, 5.04%, 3.41%, 3.24%, and 3.09%.

4. CONCLUSIONS

In this article, we presented a novel maximum-likelihood-based approximate algorithm for reconstructing viral quasispecies from high-throughput sequencing data. aBayesQR assembles paired-end short reads into longer fragments based on similarity of the read overlaps and the uncertainty level of nonoverlapping regions. The probable sets of partially reconstructed strains are inductively searched and a subset of those strains is extended to efficiently deduce the most likely set of strains in a quasispecies. Detection of the population size is embedded into the algorithm and is empirically shown to be very accurate; the number of strains is dynamically adjusted based on the reliability of the partially assembled quasispecies in each extension step. Performance of the developed method is tested on both synthetic data sets and a real HIV-1 data set. In both settings, the new algorithm outperforms existing techniques in terms of accuracy of the quasispecies size estimation, perfect reconstruction of strains, proportion of correct bases in each reconstructed strain, and the estimation of their abundance.

A particularly high accuracy is observed in estimating the population size (i.e., the number of strains) and the relative abundance. Tests on synthetic data sets demonstrate that aBayesQR is capable of reconstructing quasispecies at low diversity, showing superior performance in those settings compared to state-of-the-art algorithms. Furthermore, the study on a real HIV-1 data set demonstrates that our proposed algorithm outperforms or has performance comparable to that of the existing methods in the general setting of viral quasispecies reconstruction.

aBayesQR can be extended and applied to the problem of estimating the population size and the degree of variation among the constituent species in related fields such as immunogenetics. On a related note, bacterial populations are characterized by having relatively lower mutation rates than viral and thus typically have fewer segregating sites on the sequences in a population. The ability of our method to perform highly accurate reconstructions in such settings should be further investigated.

A software aBayesQR is available at <https://github.com/SoyeonA/aBayesQR>. An appendix can be found in a bioRxiv version of this article, which is available at <http://biorxiv.org/content/early/2017/02/06/103630>.

ACKNOWLEDGMENT

This work was funded by the National Science Foundation under grants CCF 1507998 and CCF 1618427.

AUTHOR DISCLOSURE STATEMENT

No competing financial interests exist.

REFERENCES

- Archer, J., Baillie, G., Watson, S.J., et al. 2012. Analysis of high-depth sequence data for studying viral diversity: A comparison of next generation sequencing platforms using segminator II. *BMC Bioinformatics* 13, 47.
- Astrovskaya, I., Tork, B., Mangul, S., et al. 2011. Inferring viral quasispecies spectra from 454 pyrosequencing reads. *BMC Bioinformatics* 12, 1.

- Di Giallonardo, F., Töpfer, A., Rey, M., et al. 2014. Full-length haplotype reconstruction to infer the structure of heterogeneous virus populations. *Nucleic Acids Res.* 42, e115–e115.
- Duarte, E., Novella, I., Weaver, S., et al. 1994. RNA virus quasispecies: Significance for viral disease and epidemiology. *Infect. Agents Dis.* 3, 201–214.
- Eriksson, N., Pachter, L., Mitsuya, Y., et al. 2008. Viral population estimation using pyrosequencing. *PLoS Comput. Biol.* 4, e1000074.
- Jayasundara, D., Saeed, I., Maheswararajah, S., et al. 2014. ViQuaS: An improved reconstruction pipeline for viral quasispecies spectra generated by next-generation sequencing. *Bioinformatics*, vol. 31, btu754.
- Jung, S.Y., and Kim, T.-S. 2001. An agglomerative hierarchical clustering using partial maximum array and incremental similarity computation method, 265–272. *In Data Mining, 2001. ICDM 2001, Proc. IEEE Int. Conf. on. IEEE*, San Jose, CA.
- Lancia, G., Bafna, V., Istrail, S., Lippert, R., and Schwartz, R. 2001. SNPs problems, complexity, and algorithms, 182–193. *In Algorithms—ESA 2001*. ESA (European Symposium on Algorithms). Springer.
- Lauring, A.S., and Andino, R. 2010. Quasispecies theory and the behavior of RNA viruses. *PLoS Pathog.* 6, e1001005.
- Le, T., Chiarella, J., Simen, B.B., et al. 2009. Low-abundance HIV drug-resistant viral variants in treatment-experienced persons correlate with historical antiretroviral use. *PLoS One.* 4, e6079.
- Lippert, R., Schwartz, R., Lancia, G., et al. 2002. Algorithmic strategies for the single nucleotide polymorphism haplotype assembly problem. *Brief. Bioinform.* 3, 23–31.
- Mangul, S., Wu, N.C., Mancuso, N., et al. 2014. Accurate viral population assembly from ultra-deep sequencing data. *Bioinformatics* 30, i329–i337.
- Posada-Céspedes, S., Seifert, D., and Beerenwinkel, N. 2016. Recent advances in inferring viral diversity from high-throughput sequencing data. *Virus Res.* 239, 14–32.
- Prabhakaran, S., Rey, M., Zagordi, O., et al. 2014. HIV haplotype inference using a propagating dirichlet process mixture model. *IEEE/ACM Trans. Comput. Biol. Bioinform.* 11, 182–191.
- Prosperi, M.C., Prospero, L., Bruselles, A., et al. 2011. Combinatorial analysis and algorithms for quasispecies reconstruction using next-generation sequencing. *BMC Bioinformatics* 12, 5.
- Prosperi, M.C., and Salemi, M. 2012. QuRe: Software for viral quasispecies reconstruction from next-generation sequencing data. *Bioinformatics* 28, 132–133.
- Sasirekha, K., and Baby, P. 2013. Agglomerative hierarchical clustering algorithm—A review. *Int. J. Sci. Res. Publ.* 3, 1–3.
- Schirmer, M., Sloan, W.T., and Quince, C. 2012. Benchmarking of viral haplotype reconstruction programmes: An overview of the capacities and limitations of currently available programmes. *Brief. Bioinform.* Vol. 15, bbs081.
- Simen, B.B., Simons, J.F., Hullsiek, K.H., et al. 2009. Low-abundance drug-resistant viral variants in chronically HIV-infected, antiretroviral treatment-naïve patients significantly impact treatment outcomes. *J. Infect. Dis.* 199, 693–701.
- Töpfer, A., Marschall, T., Bull, R.A., et al. 2014. Viral quasispecies assembly via maximal clique enumeration. *PLoS Comput. Biol.* 10, e1003515.
- Töpfer, A., Zagordi, O., Prabhakaran, S., et al. 2013. Probabilistic inference of viral quasispecies subject to recombination. *J. Comput. Biol.* 20, 113–123.
- Westbrooks, K., Astrovskaya, I., Campo, D., et al. 2008. HCV quasispecies assembly using network flows, 159–170. *In Bioinformatics Research and Applications*. Bioinformatics Research and Applications, Atlanta, GA.
- Zagordi, O., Bhattacharya, A., Eriksson, N., et al. 2011. ShoRAH: Estimating the genetic diversity of a mixed sample from next-generation sequencing data. *BMC Bioinformatics* 12, 119.
- Zagordi, O., Geyrhofer, L., Roth, V., et al. 2010a. Deep sequencing of a genetically heterogeneous sample: Local haplotype reconstruction and read error correction. *J. Comput. Biol.* 17, 417–428.
- Zagordi, O., Klein, R., Däumer, M., et al. 2010b. Error correction of next-generation sequencing data and reliable estimation of HIV quasispecies. *Nucleic Acids Res.* 38, 7400–7409.

Address correspondence to:

Dr. Soyeon Ahn
 Department of Electrical and Computer Engineering
 The University of Texas at Austin
 Austin, TX 78712

E-mail: soyeon.ahn@utexas.edu