## Supplementary Document

## Supplementary Document A: Proof of Convergence of the Alternating Minimization with Majority Voting

Recall the objective function of the optimization problem,

$$
f(\mathbf{M}, \mathbf{H})=\min _{\mathbf{M}, \mathbf{H}} \frac{1}{2}\left\|\mathcal{P}_{\Omega}\left(\mathbf{F}-\mathbf{M} \mathbf{H}^{\top}\right)\right\|_{F}^{2}
$$

Given the current estimate $\mathbf{M}_{t}$ and $\mathbf{H}_{t}$, we alternately update

$$
\mathbf{M}_{t+1}=\arg \min _{\mathbf{M} \in\{0,1\}^{m \times k}} \frac{1}{2}\left\|\mathcal{P}_{\Omega}\left(\mathbf{F}-\mathbf{M} \mathbf{H}_{t}^{\top}\right)\right\|_{F}^{2}
$$

and

$$
\mathbf{H}_{t+1}=\arg \min _{\mathbf{H} \in\{0,1\}^{4 n \times k}} \frac{1}{2}\left\|\mathcal{P}_{\Omega}\left(\mathbf{F}-\mathbf{M}_{t+1} \mathbf{H}^{\top}\right)\right\|_{F}^{2}
$$

Having fixed $\mathbf{H}_{t}, \mathbf{M}_{t+1}$ is updated by examining for each read all the possible viral haplotype associations and selecting the one that minimizes the number of base changes needed to make reads consistent with the observed information in $\mathbf{F}$, which implies that

$$
\begin{equation*}
f\left(\mathbf{M}_{t+1}, \mathbf{H}_{t}\right)-f\left(\mathbf{M}_{t}, \mathbf{H}_{t}\right) \leq 0 \tag{1}
\end{equation*}
$$

Now let us divide $\mathbf{F}$ into $k$ sub-matrices $\mathbf{F}^{i}, 1 \leq i \leq k$, each representing the collection of reads assigned to the $i^{t h}$ haplotype in $\mathbf{M}_{t+1}$, and rewrite $f\left(\mathbf{M}_{t+1}, \mathbf{H}_{t}\right)$ as

$$
f\left(\mathbf{M}_{t+1}, \mathbf{H}_{t}\right)=\frac{1}{2} \sum_{i=1}^{k} \sum_{j=1}^{n}\left\|\mathcal{P}_{\Omega_{\mathbf{F}^{i}(\cdot, j)}}\left(\mathbf{F}^{i}(\cdot, j)-\mathbf{M}_{t+1}^{i} \mathbf{H}_{t}(j, \cdot)^{\top}\right)\right\|_{2}^{2}
$$

where $\mathbf{F}^{(i)}(\cdot, j)$ denotes the $j^{\text {th }}$ column of $\mathbf{F}^{(i)}, \mathbf{M}_{t+1}^{i}$ has for rows the standard unit vectors $\mathbf{e}_{i}$, and $\mathbf{H}_{t}(j, \cdot)$ is the $j^{t h}$ row of $\mathbf{H}_{t}$ at the $t^{t h}$ iteration. Since $\mathbf{H}_{t+1}$ is updated by forming consensus sequences using reads in each sub-matrix
clustered by $\mathbf{M}_{t+1}$, we obtain

$$
\begin{align*}
& f\left(\mathbf{M}_{t+1}, \mathbf{H}_{t+1}\right)-f\left(\mathbf{M}_{t+1}, \mathbf{H}_{t}\right) \\
= & \frac{1}{2} \sum_{i=1}^{k} \sum_{j=1}^{n}\left\|\mathcal{P}_{\Omega_{\mathbf{F}^{i}(\cdot, j)}}\left(\mathbf{F}^{i}(\cdot, j)-\mathbf{M}_{t+1}^{i} \mathbf{H}_{t+1}(j, \cdot)^{\top}\right)\right\|_{2}^{2} \\
& -\frac{1}{2} \sum_{i=1}^{k} \sum_{j=1}^{n}\left\|\mathcal{P}_{\Omega_{\mathbf{F}^{i}(\cdot, j)}}\left(\mathbf{F}^{i}(\cdot, j)-\mathbf{M}_{t+1}^{i} \mathbf{H}_{t}(j, \cdot)^{\top}\right)\right\|_{2}^{2} \\
= & \frac{1}{2} \sum_{i=1}^{k} \sum_{j=1}^{n}\left(\left(N-N_{\mathbf{H}_{t+1}(j, i)}\right)-\left(N-N_{\mathbf{H}_{t}(j, i)}\right)\right) \\
= & \frac{1}{2} \sum_{i=1}^{k} \sum_{j=1}^{n}\left(N_{\mathbf{H}_{t}(j, i)}-N_{\mathbf{H}_{t+1}(j, i)}\right) \leq 0 \tag{2}
\end{align*}
$$

where $N$ denotes the total number of the observed nucleotides in $\mathbf{F}^{i}(\cdot, j)$ and $N_{\mathbf{H}_{t}(j, i)}$ is the number of nucleotides corresponding to $\mathbf{H}_{t}^{(j, i)}$. The combination of expressions (1) and (2) shows that the proposed alternating minimization procedure is guaranteed to converge.

## Supplementary Document B: Additional results

(1) Comparing speed of tensor factorization that relies on majority voting with the one that relies on gradient descent

Table S1: Runtime comparison of majority voting and gradient descent.

|  |  | div\% |  |  |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $k$ | method | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 |
| 5 | MV | 1.00 | 1.50 | 1.83 | 2.64 | 3.63 | 3.51 | 4.17 | 4.49 | 4.59 | 5.95 |
|  | GD | 31.84 | 42.40 | 39.60 | 78.60 | 66.04 | 105.55 | 93.87 | 112.49 | 116.83 | 202.12 |
| 10 | MV | 3.93 | 8.75 | 13.13 | 19.62 | 24.96 | 25.03 | 28.73 | 33.51 | 38.32 | 36.20 |
|  | GD | 233.88 | 321.14 | 320.79 | 374.16 | 426.62 | 535.04 | 513.29 | 752.01 | 604.85 | 678.27 |

Running time comparisons (sec) of majority voting (MV) and gradient descent (GD) for tensor
factorization on the simulated data with $\varepsilon=2 \times 10^{-3}$ and $\operatorname{cov}=500 \times$ for a mixture of 5 and 10 strains vs diversity (div), measured on a Linux OS servers with Intel Xeon Phi 7250 (1.4GHz) and 96GB RAM.

In Table S2, we compare runtimes of majority voting and gradient descent applied to solving equation (3) in Section 2.2 of the main paper. The dataset is the same as the one used in Section 3.1: mixtures of 5 and 10 strains are sequenced with error rate $\varepsilon=2 \times 10^{-3}$ and coverage cov $=500 \times$. The speed is measured on a Linux OS server with Intel Xeon Phi 7250 ( 1.4 GHz ) and 96GB RAM. As evident from the table, tensor factorization implemented with majority voting is much more time-efficient than the gradient-descent based approach for all strain diversity levels.
(2) Comparing accuracy of viral quasispecies reconstruction based on single-pass tensor factorization (AltHap) with the one that employs multiple tensor factorizations and read removal (TenSQR)

We compare the performance of TenSQR and AltHap (?) on the simulated dataset consisting of 5 viral strains whose synthesis is described in Section 3.1. We measure the performances of two schemes in terms of Recall, Precision, Predicted Proportion (PredProp), Reconstruction Rate (ReconRate) and JSD. Since AltHap requires the number of haplotypes a priori, we do not report Predicted Proportion. As shown in Figure S1, TenSQR significantly outperforms AltHap. This, of course, is expected: AltHap is designed for reconstruction of sequence communities characterized by uniform abundances - an assumption violated in viral quasispecies.


Figure S1: Performance comparison of TenSQR and AltHap in terms of Recall, Precision, Reconstruction Rate (ReconRate) and $J S D$ on the simulated data with $\varepsilon=2 \times 10^{-3}$ for a mixture of 5 viral strains.


Figure S2: Performance comparison of TenSQR, aBayesQR, ShoRAH, ViQuaS and PredictHaplo in terms of Recall, Precision, Predicted Proportion (PredProp), Reconstruction Rate (ReconRate) and $J S D$ on the simulated data with $\varepsilon=2 \times 10^{-3}$ for a mixture of (a) 5 viral strains and (b) 10 viral strains.


Figure S3: Performance comparison of TenSQR, aBayesQR, ShoRAH, ViQuaS and PredictHaplo in terms of Recall, Precision, Predicted Proportion (PredProp), Reconstruction Rate (ReconRate) and JSD on the simulated data with $\varepsilon=7 \times 10^{-3}$ for a mixture of (a) 5 viral strains and (b) 10 viral strains.
Table S2: Performance comparisons of TenSQR, aBayesQR and PredictHap on a real HIV-1 5-virus-mix data.

Predicted Proportion (PredProp) and Reconstruction Rate (RR (\%)) for TenSQR, aBayesQR and PredictHaplo applied to reconstruction of HIV
HIV-189.6, HIV-1JR-CSF, HIV-1NL4-3 and HIV-1YU2 for all 13 genes of the HIV-1 dataset. Frequencies are reported in parenthesis.

## Supplementary Document C: Estimating insertions

In addition to deletions, strains of a quasispecies may also contain insertions as compared to a reference genome. Here we propose a pipeline to detect and reconstruct insertions by processing paired-end reads that were previously filtered out due to inability to match them confidently with the reference genome; therefore, those reads have not been used in tensor factorization. The idea is that some of the discarded reads that have high base-calling quality scores might have not been matched successfully to the reference because they originate from a region that was inserted into one of the strains of the quasispecies.

The three major steps of the proposed procedure are the following:

1. Infer the origin of the filtered-out paired-end reads having high basecalling quality scores.

Let us denote two sequences in a paired-end read as $\left(r_{h}, r_{l}\right)$ where the first one is a read having high mapping score and the other one a read which is not mapped or is only partially aligned to the reference (i.e., has a low mapping score) due to an insertion. In order to infer the origin of the paired-end read $\left(r_{h}, r_{l}\right)$, we align the mapped read $r_{h}$ to the position identified by the alignment software and count the number of mismatches between the read $r_{h}$ and each of $k$ reconstructed strains. The strain $q_{I}$ having the smallest Hamming distance to the paired-end read is inferred as the origin of the paired-end read $\left(r_{h}, r_{l}\right)$ and thus contains an insertion.

## 2. Alignment of $r_{l}$

Having identified strain $q_{I}$ as the origin of the paired-end read $\left(r_{h}, r_{l}\right)$, we next attempt to find the most likely position of $r_{l}$ relative to $q_{I}$. To this end, we examine subsequences that consist of the first and last $l$ nucleotides of $r_{l}$ and compare them to every $l$-nucleotides long subsequence of $q_{I}$. Note that we assume $l<L_{r_{l}}$, where $L_{r_{l}}$ denotes the length of read $r_{l}$. To identify the best placement of $r_{l}$ along $q_{I}$, we measure alignment scores for $2\left(L_{g}-l\right)$ windows, where $L_{g}$ is the length of $q_{I}$, and choose the one with the largest score. More specifically, the alignment score is updated for each window at each position of $1 \leq i \leq l ;$ score $_{i+1} \leftarrow$ score $_{i}+N_{\text {match }}$ if the $i^{t h}$ nucleotide of the $l$-nucleotides long subsequences is matched with one in $q_{I}$, otherwise the score is updated as score $_{i+1} \leftarrow$ score $_{i}-\left(N_{\text {mismatch }}+1\right)^{2}$ where $N_{\text {match }}$ and $N_{\text {mismatch }}$ are the number of consecutive matches and mismatches, respectively; the scoring function penalizes consecutive mismatches while favoring consecutive matches to form a reliable estimate of the placement for $r_{l}$.

Given a relative position of $r_{l}$ along $q_{I}$, we infer the inserted position (i.e., the position at which $r_{l}$ starts to differ from $q_{I}$ due to insertions) by relying on the same scoring function used to determine the best alignment position. In particular, the score is computed for each position from 1 to $L_{r_{I}}$ for reads whose $l$-nucleotides-long prefix maps to $q_{I}$ and from $L_{r_{I}}$ to $i$ for those whose $l$-nucleotides-long suffix maps to $q_{I}$; then the position at which insertion starts is estimated as $I+1$, where $I=\arg \max _{i}$ score $_{i}$. For computational efficiency, the updates of $s c o r e_{i}$ are terminated once score $_{i}<0$.

Table S3: Performance of recovering insertions.

|  |  | div\% |  |  |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $l_{\text {ins }}$ | $f_{\min } \%$ | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 |
| 100 | 5 | 0.28 | 0.12 | 0.08 | 0.06 | 0.12 | 0.10 | 0.02 | 0.10 | 0.04 | 0.18 |
|  |  | $(0.95)$ | $(0.98)$ | $(0.99)$ | $(0.99)$ | $(0.99)$ | $(0.99)$ | $(0.99)$ | $(0.99)$ | $(0.99)$ | $(0.99)$ |
|  | 10 | 0.14 | 0.08 | 0.10 | 0.04 | 0.04 | 0.04 | 0.04 | 0.02 | 0.02 | 0 |
|  |  | $(0.98)$ | $(0.99)$ | $(0.99)$ | $(0.97)$ | $(0.98)$ | $(0.98)$ | $(0.99)$ | $(0.99)$ | $(0.98)$ | $(0.99)$ |
| 200 | 5 | 0.18 | 0.20 | 0.14 | 0.06 | 0.12 | 0.06 | 0.12 | 0.04 | 0.14 | 0.08 |
|  |  | $(0.99)$ | $(0.99)$ | $(0.99)$ | $(0.99)$ | $(0.99)$ | $(0.99)$ | $(0.99)$ | $(0.99)$ | $(0.99)$ | $(0.99)$ |
|  | 10 | 0.12 | 0.10 | 0.02 | 0.06 | 0 | 0.06 | 0.04 | 0.02 | 0.02 | 0 |
|  |  | $(1)$ | $(1)$ | $(0.99)$ | $(1)$ | $(1)$ | $(0.99)$ | $(0.99)$ | $(0.99)$ | $(0.99)$ | $(0.99)$ |

Performance of insertions recovery in terms of the ratio of False Negative and Reconstruction Rate of insertion (in parentheses) averaged over 50 instances on the simulated dataset with
$\varepsilon=2 \times 10^{-3}$ and $\operatorname{cov}=500 \times$ for a mixture of 2 strains, as a function of the diversity (div) and the frequency of the low abundant strain $\left(f_{\min }\right)$ which contains an insertion of the length ( $l_{\text {ins }}$ ) 100 bp and 200bp.

## 3. Recover inserted sequences

Given the alignment and the position at which the insertion starts, recovery of the insertion is readily performed by constructing two consensus sequences one built using the collection of reads whose $l$-nucleotides-long prefix maps to $q_{I}$ and the other whose $l$-nucleotides-long suffix maps to $q_{I}$. Finally, the two consensus sequences are aligned and merged to recover the entire insertion.

To study insertions recovery, we generated data that emulate the same sequencing platform considered in Sections 3.1 and 3.2 of the paper. Sets of quasispecies consisting of two strains where the length of the abundant strain is 1300bp while the other strain (present at frequencies $f_{\min }$ of 5 and $10 \%$ ) includes insertions of length $l_{\text {ins }}=100$ and 200bp. 50 instances are generated for each of 40 benchmark sets of sequencing reads. Strains in the quasispecies have diversity that varies from $1 \%$ to $10 \%$; sequencing is performed at the coverage of $500 \times$ per strain and is affected by sequencing errors incurring with probability $\varepsilon=2 \times 10^{-3}$. Performance of insertion detection is evaluated by means of the false negative rate of the detected insertions and the reconstruction rate of the inserted sequences. Table S 3 reports the results of recovering insertions in viral strains. As indicated by the reconstructed rate shown in parentheses, once insertions in the strains are detected, our approach is able to accurately reconstruct the inserted sequences.

