

# Haplotype assignment of longitudinal viral deep-sequencing data using co-variation of variant frequencies

Cristina Venturini<sup>1†</sup>, Juanita Pang<sup>2†</sup>, Asif U. Tamuri<sup>3</sup>, Sunando Roy<sup>1</sup>, Claire Atkinson<sup>4</sup>, Paul Griffiths<sup>4</sup>, Judith Breuer<sup>1,5</sup>, Richard A. Goldstein<sup>2\*</sup>

<sup>1</sup> Infection, Immunity, Inflammation, Institute of Child Health, University College London, London WC1E 6BT, United Kingdom

<sup>2</sup> Division of Infection and Immunity, University College London, London WC1E 6BT, United Kingdom

<sup>3</sup> Research IT Services, University College London, London WC1E 6BT, United Kingdom

<sup>4</sup> Institute for Immunity and Transplantation, University College London, London, United Kingdom

<sup>5</sup> Great Ormond Street Hospital for Children, London WC1N 3JH, United Kingdom

† These authors contributed equally to this work.

\* To whom correspondence should be addressed. Email: [r.goldstein@ucl.ac.uk](mailto:r.goldstein@ucl.ac.uk)

## ABSTRACT

Longitudinal deep sequencing of viruses can provide detailed information about intra-host evolutionary dynamics including how viruses interact with and transmit between hosts. Many analyses require haplotype reconstruction, identifying which variants are co-located on the same genomic element. Most current methods to perform this reconstruction are based on a high density of variants and cannot perform this reconstruction for slowly evolving viruses. We present a new approach, HaROLD (Haplotype Reconstruction Of Longitudinal Deep sequencing data), which performs this reconstruction based on identifying co-varying variant frequencies using a probabilistic framework. We illustrate HaROLD on both RNA and DNA viruses with synthetic Illumina paired read data created from mixed human cytomegalovirus and norovirus genomes, and clinical datasets of human cytomegalovirus and norovirus samples, demonstrating high accuracy, especially when longitudinal samples are available.

## INTRODUCTION

Next generation sequencing (NGS) of virus populations derived from medical and biological samples can deepen our understanding of virus biology, pathogen evolution, host-pathogen interactions,

transmission dynamics and the development of drug resistance (1–3). While smaller than bacterial and eukaryotic genomes, virus genomes are still much larger than the individual reads that are obtained through Illumina next generation sequencing. Detailed analyses often require determining which variants are found together in the same genome or genomic segment, a process known as haplotype reconstruction. This is commonly performed by identifying variants at sites that are close enough to be found on the same reads. If these variants are sufficiently dense, co-occurring variants across the genome can be ‘stitched together’, resulting in the determination of whole genome haplotypes (4). Several computer programs have been developed over the last decade using this approach to reconstruct haplotypes from NGS data, including PredictHaplo (5) and CliqueSNV(6). A recent work by Eliseev and colleagues (7) benchmarked several of these tools and found that PredictHaplo and CliqueSNV outperformed the others. However, all these tools have been created and tested with small and fast-evolving viruses such HIV-1 and HCV. Unfortunately, viruses such as human cytomegalovirus (HCMV) can have long regions with few segregating sites, making it impossible to connect variants that span these regions.

There is increased focus on monitoring intra-host evolutionary dynamics using longitudinal sequencing, where samples are obtained from a single patient at multiple time points. Selection and drift result in changes in the relative frequencies of the haplotypes and thus in the frequencies of the variants that they contain. In such cases, we can use co-variation of variant frequencies to provide an additional source of information for haplotype reconstruction, even when these variants are far apart in the genome. To take advantage of data from longitudinal sampling and include bigger recombining viruses, such as herpes viruses, we created a new method for reconstructing whole-genome haplotypes from longitudinal sequence data (HAplotype Reconstruction Of Longitudinal Deep sequencing data, HaROLD). Few other tools have been developed that use frequency data to reconstruct haplotypes, such as EVORhA (8), which was specifically developed for bacteria. A new method has been described recently that has similarities to the first stage of the method described here; analysis of simulated data demonstrated good accuracy for haplotype frequencies but lower haplotype reconstruction accuracy than CliqueSNV (9).

Here, we describe HaROLD and compare its performance with CliqueSNV, PredictHaplo and EVORha. These comparisons were performed using synthetic NGS data obtained by simulating longitudinal sampling for two different types of viruses studied in our lab: norovirus, a highly diverse RNA virus and HCMV, a large (235k bp), slowly evolving DNA virus. We also illustrate how HaROLD works compared to other methods with real data from two immunocompromised patients, one infected

with HCMV and one with Norovirus. Application of this approach to real data from mixed-infected HCMV patients has been presented previously (10, 11).

## RESULTS

We consider that we have sets of reads from a number of samples analysed using NGS, where all of the samples share a common set of related haplotypes. These may, for instance, represent a series of virus samples that have been extracted from a single patient at various time points (longitudinal samples). Note that the number of samples can be as small as one, and each sample does not necessarily contain every haplotype (the frequency of a haplotype in some samples may be zero). We are interested in determining the sequences of the haplotypes and their frequencies in each of the samples based on the observed reads.

HaROLD performs the following steps:

- (1) Initial estimation of haplotype sequences and frequencies taking advantage of co-variation of variant frequencies.
- (2) Refinement of haplotypes through analysis of observed reads, incorporating information from co-occurring variants.

These steps are described briefly here, and in more detail in the Methods and Supplementary Information.

### Initial estimation

In this initial step we assume that the samples contain a common set of identical haplotypes but in differing proportions. In order to make an initial estimation of the haplotypes, we employ a statistical model that describes the observed sequence data consisting of 1) a set of haplotype frequencies, representing the frequency of each of the haplotypes in each sample, and 2) a distribution of sequencing error rates, represented by a Dirichlet distribution. The statistical model does not include the sequences of the haplotypes. Instead, we consider these sequences, as well as the rates for specific sequencing errors at specific sites, to represent unknown 'nuisance parameters'. In this initial stage, we consider each site in the genome separately. Although we lose information about the co-occurrence of variants at different sites along the reads, this simplification allows us to avoid a complicated and costly *exploration* of the space of possible haplotype sequences in favour of a simple *sum* over the  $4^H$  ways of assigning bases to each site in the  $H$  haplotypes.

We optimise the haplotype frequencies and error rate parameters to maximise the likelihood of the data, where the calculation of the likelihood involves an explicit summation over the sequences of the haplotypes and an integration over the error rates. (This is similar to other hybrid maximum-likelihood / Bayesian approaches, such as in phylogenetics where the phylogenetic tree is optimised based on a likelihood calculation that sums over all possible combinations of substitutions.)

Once we have derived the optimal haplotype frequencies and the error rate parameters, we can calculate the posterior probability that each possible base occurs at each site in each of the haplotypes. This provides a probabilistic reconstruction of the haplotypes, indicating the appropriate degree of confidence one should have about the haplotype reconstruction of each site. Whenever this probability is sufficiently high, we can assign a specific base to that haplotype.

Although we use standard quality controls to distinguish spurious reads, we do not attempt to distinguish reliable and erroneous bases based on, for instance, number or frequency of observation. Rather, we explicitly model the probability that a specific base is observed, either correctly or erroneously. The estimation of this probability depends on a characterisation of the error rate, which may depend on the true base, the observed base, the location in the alignment, the direction of the read and the sample in which this read was present. Rather than assuming a fixed error rate, we model true-base- observed-base- location- direction- sample-specific error rates as independent draws from a Dirichlet distribution with parameters  $\alpha_0$  and  $\alpha_\epsilon$ , which are optimised during this initial step. The representation of the distribution of error rates as a Dirichlet distribution allows a closed-form integration over error rates.

This procedure is repeated for a range of different numbers of haplotypes. Increasing the number of haplotypes increases the number of ways of assigning bases to haplotypes, decreasing the prior probability of any given assignment. As is common in Bayesian methods, this results in the log likelihood decreasing as the number of haplotypes increases beyond that necessary to represent the data. We choose the number of haplotypes that maximises the log likelihood of the read data.

### **Refinement process**

The initial estimation step above assumes that the set of haplotype sequences are identical for the various samples, neglecting mutations that might occur between samples. It also ignores the information that forms the basis of most haplotype reconstruction methods, the presence of

multiple variants on the same read. The next step is to relax these assumptions and use variant co-localisation to refine the haplotypes.

In this refinement step (Figure 1), each sample is analysed individually. We start with the estimated frequencies of each haplotype in this sample and the *a posteriori* probability of each base at each site in each haplotype, as estimated in the initial step described above. Based on these parameters, each read is assigned probabilistically to each of the various haplotypes. The number of reads assigned to each haplotype is used to adjust the frequencies of that haplotype. The reads are then re-assigned until the haplotype frequencies have converged. The resulting assigned reads are then used to update the probability of the bases found in each site in all the reads assigned to each haplotype. The reads are then re-assigned based on these adjusted probabilities, and this procedure repeated until convergence. These two steps – estimation of haplotype frequencies and estimation of base probabilities – are then alternated until convergence. We also calculate the log likelihood of the read data given these parameter values, as well as a penalised log likelihood from which the number of adjustable parameters has been subtracted, equivalent to -0.5 times the Akaike information criterion (12).

If requested by the users, a number of structural modifications of the haplotypes are considered. These include a) recombination of two haplotypes, where corresponding regions of the haplotype sequences are swapped, b) gene conversion, where a region of one haplotype sequence is overwritten by the corresponding region of a different haplotype sequence, c) merging of two haplotypes into a single new haplotype, reducing the total number of haplotypes by one, and d) dividing a single haplotype into two new haplotypes, increasing the total number of haplotypes by one. After each of these modifications the haplotype frequencies and base probabilities are re-adjusted as described above, and the modification rejected or accepted based on whether it results in a decrease or increase in the penalised log likelihood. These modifications result in a final set of haplotypes whose size represents the number of haplotypes that can be justified, based on information theory, by the sequence data.

The output of the program includes, for each sample, the frequencies of the haplotypes as well as the probabilities of each of the bases at each of the sites in the haplotypes. When this probability is over a user-defined value, the site can be assigned to a specific base.

## Simulation results

To evaluate the ability of HaROLD to reconstruct haplotypes and estimate the relative haplotype frequencies, we created eight synthetic sequence datasets (four for norovirus and four for HCMV), each consisting of a set of sequential longitudinal samples drawn from differing mixtures of whole genome sequences from GenBank, as summarized in Tables 1 and 2. The norovirus datasets (Table 1) each consisted of five longitudinal samples of between two to four haplotypes with varying degrees of similarity (total number of norovirus samples = 20). The HCMV datasets (Table 2) were constructed in a similar manner, each with three longitudinal samples (total number of HCMV samples = 12) constructed from two or three haplotypes.

*Reconstruction of haplotype sequences and frequency.* The performance of HaROLD on the synthetic data is represented in Figures 2 and 3 (in sky blue). Performance was calculated as 1) the proportion of sites that are identical between Genbank sequences and the reconstructed haplotypes, and 2) the difference between the real haplotypes frequencies and frequencies calculated by HaROLD. With the norovirus data, the reconstructed haplotypes were identical to the Genbank sequences (accuracy 100% in all four datasets) (Figure 2A). The haplotype frequencies estimated by HaROLD were also highly accurate, with differences between the actual and estimated frequencies less than 0.002 (Figure 2B). Excellent results were also obtained with the synthetic data derived from HCMV; the reconstructed haplotypes were highly similar to the original sequences (similarity > 0.997) (Figure 3A) with differences between the actual and computed haplotype frequencies less than 0.06 (Figure 3B). HaROLD computational time (in an HPC node with a maximum of 50GB memory) was dependent on the number of haplotypes, average read depth and genome length and varied from 40 to 226 seconds for norovirus and 35 to 39 minutes for HCMV (Table 3). This can be longer in some cases, generally dominated by the estimation of the error rate parameters. The calculations can, correspondingly, be greatly sped up if these parameters are estimated and fixed.

*Utility of longitudinal sampling.* In contrast to most methods for haplotype reconstruction, HaROLD is formulated to take advantage of the availability of multiple longitudinal samples. To evaluate the importance of these longitudinal samples, we used HaROLD to reconstruct the haplotypes in our synthetic datasets without using this additional information. We compared performance of HaROLD using 1) single samples run independently (HaROLD-Single) and 2) longitudinal samples pooled together in one sample (HaROLD-Pooled). In both norovirus and HCMV the performance of HaROLD-Single and HaROLD-Pooled was high, varying from 0.99 to 1 similarity to the Genbank sequences, decreasing with the increase of number of haplotypes (Figure 2A and 3A in blue and in turquoise,

respectively). For HaROLD-Single, estimated frequencies were 93%-100% similar to actual frequencies in norovirus and 78-99% for HCMV, not as accurate as when HaROLD was run as designed with longitudinal samples (Figure 2B and 3B). In general, the performance of HaROLD on HaROLD-Single and HaROLD-Pooled was not as accurate as when longitudinal data were used, highlighting the advantage of using serial sampling. Even so, the accuracies of the haplotype reconstructions were generally quite high, especially for the shorter norovirus sequences and when there were relatively few haplotypes.

*Refinement step.* As described above and in more detail in Methods, HaROLD uses two different steps to reconstruct the final haplotypes and frequency. We evaluated the performance and utility of the additional refinement step by comparing results from HaROLD and HaROLD without this additional refinement step (HaROLD-Step1). In both Norovirus and HCMV the reconstructed haplotypes for HaROLD-Step1 were accurate, varying from 0.996 to 1 (Norovirus) and 0.99-0.999 (HCMV), decreasing as number of haplotypes increase (Figure 2A and 3A in green). Estimated frequencies were accurate (Norovirus 93-100%; HCMV 79-99%) (Figure 2B and 3B in green). However, both datasets achieved a lower accuracy in HaROLD-Step1 compared to results obtained with HaROLD (both steps). The difference is larger in the HCMV dataset, especially in the dataset with three haplotypes with low similarity.

*Comparison with other methods.* We compared the performance of HaROLD with two haplotype reconstruction techniques reviewed in (7), namely CliqueSNV and PredictHaplo; these two methods performed well in term of accuracy in their HIV validation. CliqueSNV is a reference-based method to reconstruct haplotypes from NGS short reads data, which constructs an allele graph based on linkage between variants and identifies true viral variants by merging cliques of that graph via combinatorial optimization techniques (6). PredictHaplo implements a fully probabilistic approach to quasispecies reconstruction. Given a set of aligned reads, it uses a Bayesian mixture model with a Dirichlet process prior to estimate the unknown number of underlying haplotypes (5). We also added a comparison with a third method, EVORhA (8), that was developed for bacterial haplotype reconstruction and combines phasing information in regions of overlapping reads with the estimated frequencies of inferred local haplotypes. This method was chosen because it is one of the few other haplotype reconstruction methods which also considers variant frequencies. We ran these three methods using default parameters unless otherwise stated. All analyses were run on an HPC node with a maximum of 48h and 50GB of memory. In both the HCMV and norovirus datasets, EVORhA generally estimated a larger number of haplotypes than present in the sample (ranging from 1 to 5

additional haplotypes), and consistently yielded haplotypes that most resembled the input reference sequence used for mapping. The sequence accuracy ranged from 0.972 to 0.999 for HCMV (Figure 3A in yellow) and from 0.983 to 0.999 for Norovirus (Figure 2A in yellow), consistently lower than HaROLD. The performance of EVORhA in estimating the relative haplotype frequencies was uneven and overall worse compared to HaROLD (Figure 2B and 3B). On the norovirus datasets, CliqueSNV yielded more accurate haplotype sequences than EVORhA; frequency accuracy was, however, uneven (Figure 2A and B in red). PredictHaplo performed similarly to CliqueSNV (sequence accuracy from 0.988 to 1, Figure 2A in purple) and again frequency accuracy was uneven, especially with four haplotypes (Figure 2B). HaROLD consistently outperformed these other techniques in both sequence and frequency accuracy, even when single samples were run independently (as explained in the previous section *Utility of longitudinal sampling*) (Figure 2). We were not able to analyse the HCMV datasets using CliqueSNV and PredictHaplo due to memory constraints; both programs were developed for smaller RNA viruses such as HIV and were not able to analyse a genome as large as ~250k bp with available computational resources. Harold was generally faster than the other methods for the norovirus datasets, although EVORhA was faster for the HCMV datasets where the average read depth was low (Table 3).

*Diversity calculation.* As an example of the consequences of the different reconstruction accuracies on downstream analyses, we estimated the average heterozygosity of the various samples based on the reconstructed haplotypes, as shown in Figure 4. The haplotypes generated by HaROLD produced accurate estimates of the average heterozygosity, especially in the longitudinal dataset. PredictHaplo generally produced accurate heterozygosity in norovirus, albeit the accuracy decreased when four haplotypes were present. CliqueSNV underestimated heterozygosity in almost all conditions, except for when we had two norovirus haplotypes that were very similar. Finally, EVORhA underestimated heterozygosity in both HCMV and norovirus in almost all datasets, except in norovirus with two haplotypes.

### **Application to real data**

We applied our approach to two datasets: an unpublished dataset of HCMV samples from kidney/liver recipients where there was contamination from a lab strain, and a dataset including longitudinal samples of norovirus from an immunocompromised patient described by Ruis et al. (13).

*HCMV contamination dataset.* During analysis of a set of 5 longitudinal samples taken from a 42-year-old patient following a liver transplant (PatientA, T1-T5), we observed a high degree of within-



host diversity in two of the samples. HaROLD was used to look for the presence of distinct haplotypes, yielding two haplotypes for samples T1 and T3 and only a single haplotype for the other three samples. The second haplotype in T1 and T3 was nearly identical to the Merlin laboratory strain (NC.006273.2), a strain present in the sequencing lab. This was subsequently identified as a sample contaminant. Following discovery of the contamination, these two samples were re-sequenced without the contaminant, providing a real-world scenario for validation of HaROLD with typical Illumina sequencing errors and uneven read depth and coverage. We built a maximum-likelihood phylogenetic tree including the haplotypes reconstructed by HaROLD (indicated as H0 and H1) and the consensus sequences obtained from re-sequenced samples T1 and T3 (Figure 5). The two haplotypes (H0 and H1) for each of the contaminated samples clearly clustered separately in the phylogenetic tree: H0 clustered together with the consensus sequences of T1 and T3 (blue cluster) whereas H1 clustered with Merlin GenBank sequence (pink cluster); the two clusters were approximately 98.5% similar, with the 4000+ differences largely in the ‘hypervariable genes’ (14). We directly compared the HaROLD haplotypes and the consensus sequence for both samples. Sample T1 resulted in only 7 differences between patientA\_T1\_H0 and the consensus sequence patientA\_T1; these SNP differences clustered in IRS1 gene which is a region full of repeats in HCMV. Sample T3 showed 312 SNP differences between patientA\_T3\_H0 and patientA\_T3 and, again, these were in IRS1 and TRS1 genes, which contained repeats and are difficult to assemble and align.

*Norovirus dataset.* We used HaROLD to analyse norovirus deep sequencing samples from an immunocompromised 48-year-old patient with chronic norovirus infection previously published (13) (Pang J. et al, manuscript in preparation). We collected 12 longitudinal samples over almost a year during which time the patient was treated with anti-viral drug Favipiravir. The patient showed symptomatic response to Favipiravir treatment and phylogenetic analysis showed evidence for selective pressure in the infecting norovirus population. To better understand whether and how different viral populations evolved overtime and in response to treatment, we reconstructed haplotypes from all samples using HaROLD. Each sample yielded 2 to 5 haplotypes which we used to build a multiple sequences alignment together with the closest GenBank reference sequence (FJ537136). Analysis of pair-wise genetic distances showed a clear bimodal distribution (Supplementary figure 1) with two main clusters observed with multi-dimensional scaling (Figure 6). The two clades were also present in the maximum-likelihood phylogenetic tree (Figure 7). The first viral cluster (orange) was present since the first time point and was dominant in almost all samples (Figure 8). However, following treatment a second viral population (grey) appeared which become the dominant viral strain in time point 6 and 7, which correspond to the time when the patient received extensive treatment with Favipiravir.

We then compared these results with those obtained using other haplotype reconstruction methods. PredictHaplo generally gave similar results compared to HaROLD: it generated 3-8 haplotypes for each sample which generally clustered in two main viral populations which, however, were not as clearly distinct as for HaROLD, with each cluster divided into two sub-clusters (Supplementary Figure 2-4). Even though PredictHaplo did perform similarly to HaROLD, we encountered computational issues due to time and memory limits; it did not finish on five out of twelve samples (HPC node with 50GB and 14-day time limit). Both CliqueSNV and EVORhA yielded many low-frequency haplotypes (EVORhA 2-10, CliqueSNV 4-8) that tended to form diffuse clusters or were similar to the reference sequence (Supplementary Figures 2-3, Supplementary figures 5-6) which did not give information about the evolution of viral populations over time. In addition, when sequences from all methods were compared together, EVORhA haplotypes were genetically dissimilar from haplotypes obtained with other methods (Supplementary figure 3).

## DISCUSSION

The majority of methods for reconstructing haplotypes rely on reads that contain multiple polymorphic sites, and thus require a sufficiently density of polymorphic sites so that the distances between such sites are closer together than the read length (5, 6). This approach can be used, for example, with HIV-1 and HCV, two small and fast-evolving viruses, frequently used in testing haplotype reconstruction's methods. However, this is not always the case, especially for viruses such as HCMV where much of the observed sequence diversity is confined to short intervals. Even when there is copious variation, there may be closely related haplotypes where the haplotype-defining variants are separated by distances greater than the read length, making it difficult to assign these variants correctly to the otherwise similar haplotypes. HaROLD was motivated by the increasing availability of multiple samples that are likely to share closely related haplotypes, such as longitudinal studies of within host evolution or samples from an outbreak cluster. Under such conditions, variant frequencies can provide an important additional source of information for making accurate haplotype reconstructions (9). Notably, HaROLD generates haplotypes as accurate or more accurate than other tested methods even when multiple samples were not available. This greater accuracy was achieved with significantly less computing power and memory than the other methods we used for comparison, allowing rapid analysis of sequence data. It is important to note that both PredictHaplo and CliqueSNV were not able to produce results for HCMV at all, due to the size of the virus, showing that these programs would be impractical for bigger microbes such as herpes viruses and bacteria.

Even in a small virus such as norovirus, PredictHaplo encountered computational issues for 5/12 samples due to time and memory limits (HPC node with 50GB and 14-day time limit).

Even if HaROLD was created to deal with big dsDNA viruses, it performed well with RNA viruses, such as norovirus. In such context, it is difficult to determine how many haplotypes there are in the sample, even with perfect information. One could consider every unique sequence in the sample as a different haplotype, but in this case the number of haplotypes would generally be so large as to make any further analysis impractical. Alternatively, one could consider haplotypes as representing clusters of closely related sequences that do not need to be all identical. In this case, there is some flexibility in how one defines the term 'closely related'. HaROLD is generally conservative about the number of haplotypes. In particular, the refinement method does not add an additional haplotype unless the improvement in the log likelihood is sufficient to justify the resulting increase in the number of parameters. The resulting haplotypes then include some amount of variation, which is provided as output to the user. In particular, the output reports the probability that a sequence belonging to a haplotype would have any of the four bases found in each site. When these probabilities are sufficiently definitive, a base is assigned in the multiple sequence alignment. An ambiguous base is presented when a definitive assignment cannot be made.

We described the performance of HaROLD in the analysis of synthetic datasets, as well as its use for two real data examples, one for HCMV and one for Norovirus. In the HCMV example, HaROLD was able to reconstruct the patient's sequence and the lab strain contaminate sequence with high accuracy. HaROLD was successfully used to detect and confirm the contamination in the first place and then reconstruct the "real" sequence allowing us to use a sample which would have been discarded otherwise. This was a lab-created "mixed infection" and no different from a real situation where a liver recipient may have a superinfection from reactivation of a HCMV strain already present and a new infection from a HCMV-positive donor following transplant. In the second example, we showed longitudinal samples from an immune-compromised patient chronically infected with norovirus. In this patient we were able to distinguish two main viral populations with one selected during the drug treatment. Evohra (8) and CliqueSNV (6) generated larger numbers of haplotypes that either clustered by sample or very close to the reference used for mapping. PredictHaplo (5) performed similarly to HaROLD although it often did not converge within a reasonable timescale. In these examples, we illustrated how a precise haplotype reconstruction can be useful in determining the likelihood of mixed-infection and/or how the viral populations respond to treatment and evolve over time.

In a previous paper, we described another application on real clinical data for HCMV (10, 11), where HaROLD was able to reconstruct individual viral haplotypes within patients with mixed infections. By

reconstructing the full-length genome, we were able to pinpoint the likely timing, origins, and natural history of HCMV superinfections and uncover within-host viral recombination.

By providing a tool for viral haplotype reconstruction which is also suitable for DNA viruses with large genomes, we aim to simplify the investigation of mixed-infections and within-host evolution for all viruses, both when longitudinal sequences are and are not available.

## **MATERIAL AND METHODS**

HaROLD involves an initial estimation step followed by a refinement step. We describe the methods here. Further details are included in the Supplementary materials.

### **Initial estimation**

We start with a set of sequence reads from related samples that have been analysed using NGS. We initially assume that these samples contain a common set of haplotypes but in differing proportions, an assumption that will be relaxed at a later stage.

We start with an assumed total number of haplotypes for the set of samples. Following quality control and assembly of the reads, for each sample we count the number of each type of base observed at each position in the resulting alignment. The observed number of each base depends on a) the frequencies of the haplotypes in that sample, b) the base found at that position in each of the haplotypes, and c) the probability of making an erroneous measurement at that site. As the error rate may be different at different sites and on different strands, we consider that this rate is drawn from a Dirichlet distribution. We first find the maximum likelihood estimate of the haplotype frequencies in each sample and the parameters defining the error rate distribution. We account for our initial ignorance of the haplotype sequences by summing this likelihood over all possible ways the different bases observed at that position can occur in the different haplotypes. We also integrate over the distribution of error rates.

Following estimation of the haplotype frequencies and error rate distribution parameters, we determine how much each assignment of bases to haplotypes contributes to the overall likelihood. This allows us to calculate the posterior probability of each assignment of bases to haplotypes. By summing over these posterior probabilities, we can compute the marginal posterior probability that

a base is found at that site in each of the haplotypes. If these probabilities are sufficiently definitive, an assignment is made. The *a posteriori* marginal probability of each base is included in the output.

We perform this procedure for a range of different numbers of haplotypes. As increasing the number of haplotypes increases the number of ways of assigning bases to each of the haplotypes, decreasing the prior probability of any given assignment, the log likelihood typically decreases when the number of haplotypes increases beyond that necessary to represent the data. We select the number of haplotypes that maximise the log likelihood.

### Further refinement

The method described above takes advantage of the presence of the same haplotype in multiple samples at various frequencies. It assumes that, although the haplotype sequences are described probabilistically, these probabilities are identical for all of the various samples, neglecting processes such as mutations. It also ignores the information that forms the basis of most haplotype reconstruction methods, the presence of multiple variants on the same read. The next step is to relax these assumptions and use variant co-localisation to refine the haplotypes.

In this refinement step (Figure 1), each sample is analysed individually. We start with the estimated frequencies of each haplotype in this sample, and the *a posteriori* probability of each base at each site in each haplotype, as output from the previous program. The haplotypes are then optimised by assigning the reads, probabilistically, to the various haplotypes. The number of reads assigned to each haplotype is used to adjust the frequencies of each haplotype. The reads are then re-assigned until the haplotype frequencies have converged. The resulting assigned reads are then used to update the probability of the bases found in each site in all of the reads assigned to each haplotype. This process of is performed until convergence.

If requested by the users, a number of modifications of the haplotypes are considered. These include a) recombination of two haplotypes, where corresponding regions of the haplotype sequences are swapped, b) gene conversion, where a region of one haplotype sequence is overwritten by the corresponding region of a different haplotype sequence, c) merging of two haplotypes into a single new haplotype, reducing the total number of haplotypes by one, and d) dividing a single haplotype into two new haplotypes, increasing the total number of haplotypes by one.

First, the program considers an adjustable number of possible recombination of the haplotypes. These recombination events involve a) picking two haplotypes at random, b) picking a region of the alignment, of length chosen from a Normal distribution with standard deviation of 10 sites, and then c) either swapping the values of the probabilities of the different bases in this region between the two haplotypes (50% probability) or over-writing the values in one haplotype with the values from the other (25% probability for each direction). Following such a step, the haplotype frequencies and base probabilities are then re-optimised as described above, and the recombination event is either accepted or rejected based on whether the penalised log likelihood, that is, the log likelihood minus the number of adjustable parameters defining the haplotypes, is increased or decreased.

The program then implements an iterative process of refinement. At the start of each iteration, if requested, pairs of haplotypes are chosen and merged, with the frequencies of the resulting haplotype equal to the sum of that of the parents, and the base frequencies equal to the average of the two parents. This results in a reduction in the number of haplotypes by one. The haplotypes are then re-optimised. This process is repeated for every pair of haplotypes. The merge that most increases the penalised log likelihood is recorded.

If requested, a haplotype is chosen and split into two haplotypes, increasing the total number of haplotypes by one. The resulting set of haplotypes is then re-optimised. This is repeated for every original haplotype. The split that results in the largest increase in penalised log likelihood is recorded. Finally, the recombination process described above is performed. Again, the recombination event that results in the largest increase in penalised log likelihood is recorded. Following these attempted modifications of the haplotypes, the modification – merge, split, or recombination – that most increases the penalised log likelihood is compared with the penalised log likelihood at the beginning of the iteration. If this results in a net increase in the penalised log likelihood, this modification is accepted, and becomes the starting position for the next iteration. This iterative process is then repeated until convergence.

### **Preparation of synthetic test datasets**

The first four synthetic datasets consisted of mixtures of two to four norovirus sequences (approximately 7.5 kb in length) (Table 1); four additional datasets were assembled from two to three human cytomegalovirus (HCMV) sequences (approximately 230 kb) (Table 2). SimSeqNBProject (15) (<https://github.com/jstjohn/SimSeq>) was used to create 1,000,000 paired end

reads of length 250 for each GenBank norovirus sequence listed in Table 1, and 100,000 paired end reads for each GenBank CMV sequence listed in Table 2. SimSeq includes the `getErrorProfile` module which generates the error model for the sequence simulator. The output SAM files from SimSeq were then converted into Fastq files using Picard version 2.21.1 'SamToFastq' (16). In order to construct the datasets, Seqtk 1.3 (<https://github.com/shenwei356/seqtk>) (17) was used to mix the reads from each ensemble according to the relative fractions listed in Tables 1 and 2. Reads were then trimmed for adapters using Trim galore version 0.6.0 (18). Duplicate reads were removed using Picard version 2.21.1 'MarkDuplicates'. Reads were mapped to the GII.Pe-GII.4 Sydney 2012 reference strain JX459907 for norovirus, and the Merlin reference strain NC\_006273.2 for CMV using BWA version 0.7.17 (19). The `Makereadcount.jar` (<https://github.com/ucl-pathgenomics/HaROLD/tree/master/jar>) was used to obtain the strand specific nucleotide counts from BAM files. These strand count files were used as the input for HaROLD.

### Evaluation of performance

We evaluated performance of haplotype reconstruction based on accuracy of reconstructed sequences and accuracy in reporting haplotype frequency or abundance in the sample. Accuracy of reconstructed sequences was calculated as SNPs differences between the GenBank sequences and the reconstructed haplotypes using the "dist.dna" function in R library "ape" (5.4-1) (20) which produces a matrix of pairwise distances from DNA sequences. The model used was "raw", simply the proportion of sites that differ between each pair of sequences. Frequency accuracy was calculated as a difference between the real haplotypes frequencies and frequencies calculated by the software (as  $1 - \text{abs}(\text{real haplotype frequency} - \text{estimated haplotypes frequency})$ ).

We estimated average heterozygosity (with an in-house script) for each sample as a measure of genetic diversity based on the reconstructed haplotypes.

### Comparison with other haplotype reconstruction programs

We compared HaROLD performance with the latest version of EVORhA, CliqueSNV and PredictHaplo (v. 5). EVORhA was run with default parameters. We ran CliqueSNV with `-tf` option (minimum threshold for frequency relative to the read's coverage) set to 0.01 (default was 0.05, decreasing the parameters increase the sensitivity of the program and `-cm` option set as "fast"). PredictHaplo was run with default parameters as the HIV example included in the programs, except for entropy

threshold, which was set to 0.05, max gap fraction to 0.05, local window size factor to 0.9, MCMC interaction to 100 and deletions were not included.

### **Validation with HCMV dataset (real data)**

This patient was sequenced as part of the Wellcome Trust Collaborative Awards 204870. The project includes HCMV sequencing from liver and kidney transplant recipients and donors. Seven samples were available for five time points for PatientA. All samples were mapped to Merlin GenBank sequence (NC\_006273.2). The average read depth varied from 10x to 360x. Data were prepared for HaROLD in the same way as for the synthetic datasets. HaROLD analysis was run in the five samples from the initial run. Data were then aligned with Mafft (21) and trees were obtained with Iqtree 1.6.12 (22) with GTR model and 500 bootstrap and plotted with Figtree (<https://github.com/rambaut/figtree/releases/tag/v1.4.4>).

### **Validation on norovirus dataset (real data)**

We have described this patient previously (13). Fastq files were mapped to the closest Genbank reference (FJ537136) using the same pipeline used with the synthetic datasets. Haplotypes from HaROLD, EVORhA, CliqueSNV and PredictHaplo were obtained similarly to the synthetic datasets. Data were then aligned with Mafft (21) and trees were obtained with RAXML version 8 with GTRGAMMA model and 1000 bootstraps (23). Pairwise distances were retrieved with “ape” package in R (20). Trees were plotted using ggtree (version 2.4.1) (24).

### **R version and packages**

Unless otherwise stated, all statistical analysis and plots were done in R 4.0.3, ggplot2 (3.3.5) (25) and ggpubr (0.4.0) (<https://rpkgs.datanovia.com/ggpubr/index.html>).

### **DATA AND SOFTWARE AVAILABILITY**

The software HaROLD and other materials used is available in the GitHub repository <https://github.com/ucl-pathgenomics/HAROLD>. Current release (v2.0) can be found in <https://github.com/ucl-pathgenomics/HAROLD/releases/tag/v2.0.0>. HaROLD supports multiple threads.



## ACKNOWLEDGEMENTS

This study was supported by Wellcome Trust Collaborative Awards 203268 and 204870 awarded to JB, Rosetrees Trust PhD Studentship M876 awarded to JP and Medical Research Council Grant MC\_U117573805 awarded to RAG.

## CONFLICT OF INTEREST

We declare no conflict of interests.

## REFERENCES

1. Houldcroft, C.J., Beale, M.A. and Breuer, J. (2017) Clinical and biological insights from viral genome sequencing. *Nature Reviews Microbiology*, **15**, 183–192.
2. Leung, P., Eltahla, A.A., Lloyd, A.R., Bull, R.A. and Luciani, F. (2017) Understanding the complex evolution of rapidly mutating viruses with deep sequencing: Beyond the analysis of viral diversity. *Virus Research*, **239**, 43–54.
3. Moncla, L.H., Weiler, A.M., Barry, G., Weinfurter, J.T., Dinis, J.M., Charlier, O., Lauck, M., Bailey, A.L., Wahl-Jensen, V., Nelson, C.W., *et al.* (2017) Within-Host Evolution of Simian Arteriviruses in Crab-Eating Macaques. *J Virol*, **91**, e02231-16, e02231-16.
4. Posada-Céspedes, S., Seifert, D. and Beerenwinkel, N. (2017) Recent advances in inferring viral diversity from high-throughput sequencing data. *Virus Research*, **239**, 17–32.
5. Prabhakaran, S., Rey, M., Zagordi, O., Beerenwinkel, N. and Roth, V. (2014) HIV Haplotype Inference Using a Propagating Dirichlet Process Mixture Model. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, **11**, 182–191.
6. Knyazev, S., Tsyvina, V., Shankar, A., Melnyk, A., Artyomenko, A., Malygina, T., Porozov, Y.B., Campbell, E.M., Switzer, W.M., Skums, P., *et al.* (2020) CliqueSNV: An Efficient Noise Reduction Technique for Accurate Assembly of Viral Variants from NGS Data. *bioRxiv*, 10.1101/264242.
7. Eliseev, A., Gibson, K.M., Avdeyev, P., Novik, D., Bendall, M.L., Pérez-Losada, M., Alexeev, N. and Crandall, K.A. (2020) Evaluation of haplotype callers for next-generation sequencing of viruses. *Infection, Genetics and Evolution*, **82**, 104277.
8. Pulido-Tamayo, S., Sánchez-Rodríguez, A., Swings, T., Van den Bergh, B., Dubey, A., Steenackers, H., Michiels, J., Fostier, J. and Marchal, K. (2015) Frequency-based haplotype reconstruction from deep sequencing data of bacterial populations. *Nucleic Acids Res*, **43**, e105–e105.
9. Pelizzola, M., Behr, M., Li, H., Munk, A. and Futschik, A. (2021) Multiple haplotype reconstruction from allele frequency data. *Nat Comput Sci*, **1**, 262–271.
10. Cudini, J., Roy, S., Houldcroft, C.J., Bryant, J.M., Depledge, D.P., Tutill, H., Veys, P., Williams, R., Worth, A.J.J., Tamuri, A.U., *et al.* (2019) Human cytomegalovirus haplotype reconstruction reveals high diversity due to superinfection and evidence of within-host recombination. *PNAS*, **116**, 5693–5698.

11. Pang,J., Slyker,J.A., Roy,S., Bryant,J., Atkinson,C., Cudini,J., Farquhar,C., Griffiths,P., Kiarie,J., Morfopoulou,S., *et al.* (2020) Mixed cytomegalovirus genotypes in HIV-positive mothers show compartmentalization and distinct patterns of transmission to infants. *eLife*, **9**, e63199.
12. Akaike,H. (1998) Information Theory and an Extension of the Maximum Likelihood Principle. In Parzen,E., Tanabe,K., Kitagawa,G. (eds), *Selected Papers of Hirotugu Akaike*. Springer, New York, NY, pp. 199–213.
13. Ruis,C., Brown,L.-A.K., Roy,S., Atkinson,C., Williams,R., Burns,S.O., Yara-Romero,E., Jacobs,M., Goldstein,R., Breuer,J., *et al.* (2018) Mutagenesis in Norovirus in Response to Favipiravir Treatment. *New England Journal of Medicine*, **379**, 2173–2176.
14. Suárez,N.M., Musonda,K.G., Escriva,E., Njenga,M., Agbueze,A., Camiolo,S., Davison,A.J. and Gompels,U.A. (2019) Multiple-Strain Infections of Human Cytomegalovirus With High Genomic Diversity Are Common in Breast Milk From Human Immunodeficiency Virus–Infected Women in Zambia. *The Journal of Infectious Diseases*, **220**, 792–801.
15. Benidt,S. and Nettleton,D. (2015) SimSeq: a nonparametric approach to simulation of RNA-sequence datasets. *Bioinformatics*, **31**, 2131–2140.
16. Broad Institute (2019) Picard GitHub Repository.
17. Shen,W., Le,S., Li,Y. and Hu,F. (2016) SeqKit: A Cross-Platform and Ultrafast Toolkit for FASTA/Q File Manipulation. *PLOS ONE*, **11**, e0163962.
18. The Babraham Institute (2019) TrimGalore.
19. Li,H. and Durbin,R. (2009) Fast and accurate short read alignment with Burrows-Wheeler transform. *Bioinformatics*, **25**, 1754–1760.
20. Paradis,E. and Schliep,K. (2019) ape 5.0: an environment for modern phylogenetics and evolutionary analyses in R. *Bioinformatics*, **35**, 526–528.
21. Katoh,K., Misawa,K., Kuma,K. and Miyata,T. (2002) MAFFT: a novel method for rapid multiple sequence alignment based on fast Fourier transform. *Nucleic Acids Res*, **30**, 3059–3066.
22. Nguyen,L.-T., Schmidt,H.A., von Haeseler,A. and Minh,B.Q. (2015) IQ-TREE: A Fast and Effective Stochastic Algorithm for Estimating Maximum-Likelihood Phylogenies. *Molecular Biology and Evolution*, **32**, 268–274.
23. Stamatakis,A. (2014) RAxML version 8: a tool for phylogenetic analysis and post-analysis of large phylogenies. *Bioinformatics*, **30**, 1312–1313.
24. Yu,G., Smith,D.K., Zhu,H., Guan,Y. and Lam,T.T.-Y. (2017) ggtree: an r package for visualization and annotation of phylogenetic trees with their covariates and other associated data. *Methods in Ecology and Evolution*, **8**, 28–36.
25. Wickham,H. (2016) ggplot2: Elegant Graphics for Data Analysis Springer-Verlag New York.

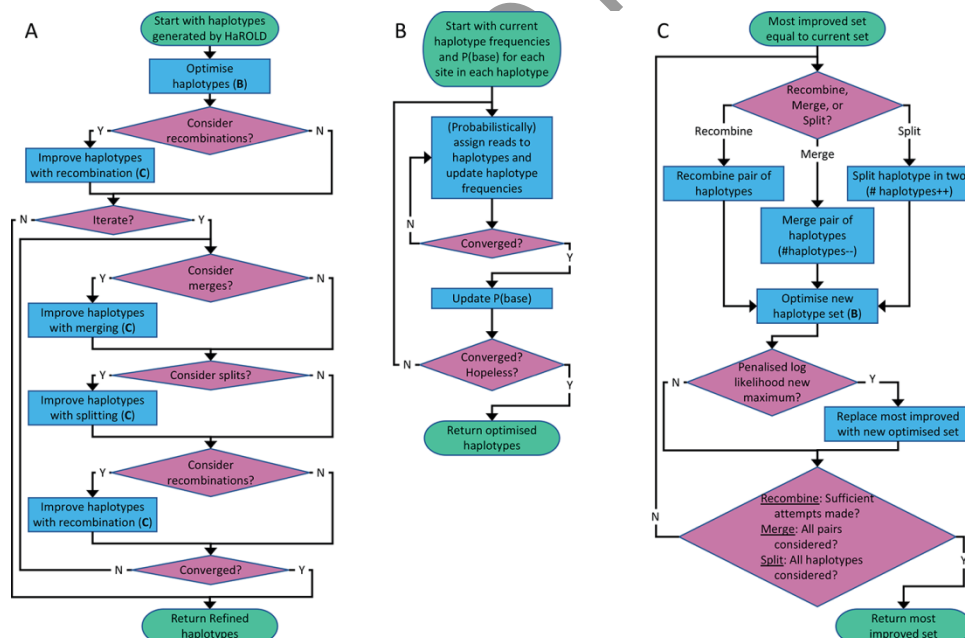
## TABLES AND FIGURES LEGENDS

**Table 1. Summary of the longitudinal norovirus synthetic datasets used to test the accuracy of the haplotype reconstruction methods.** Four synthetic datasets with five samples each were created for norovirus by mixing GenBank sequences for a total of 20 samples.

**Table 2. Summary of the longitudinal human cytomegalovirus synthetic datasets used to test the accuracy of the haplotype reconstruction methods.** Four synthetic datasets with three samples each were created for HCMV by mixing GenBank sequences for a total of 12 samples.

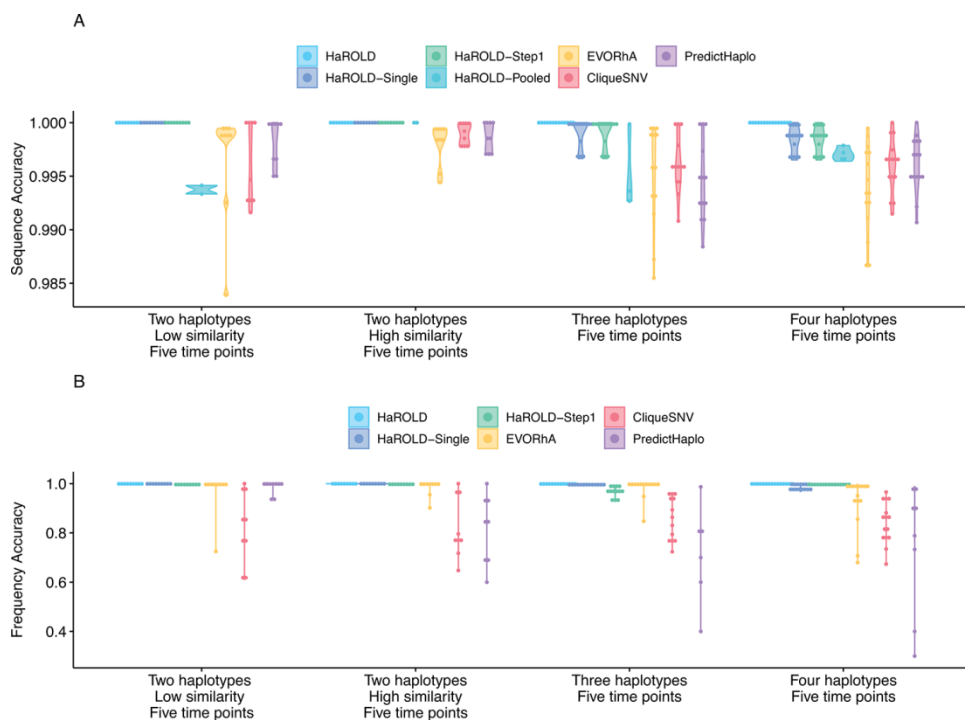
**Table 3. Computational time for HaROLD and other software for analysing norovirus and HCMV datasets.**

**Figure 1. Flowchart of refinement process.** (A) Overall process. (B) Subprocess for haplotype optimisation. (C) Subprocess for considering recombination, merging and splitting; merging decreases the number of haplotypes by one, whilst splitting increases this number by one. Support for these three operations is evaluated by considering a penalised log likelihood equal to -0.5 times the Akaike Information Criterion (AIC).

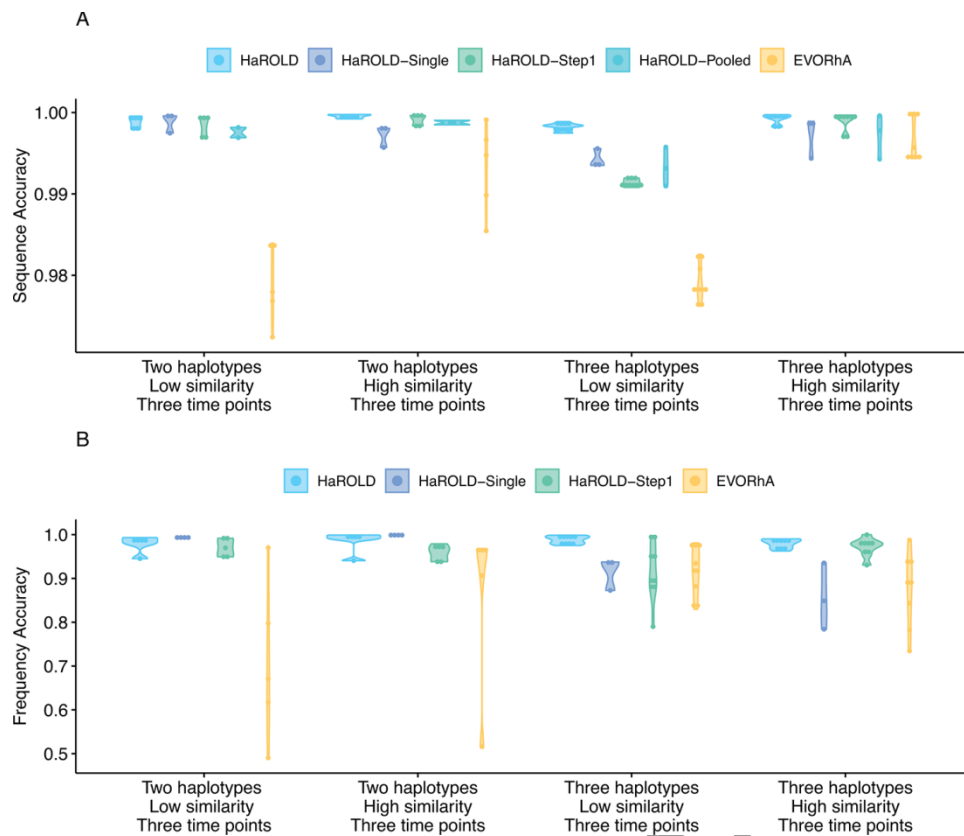


**Figure 2. Violin plots showing the accuracy of haplotype reconstruction in the norovirus test datasets.** (A) The accuracy of reconstructed sequences (pairwise identity between the actual sequences and reconstructed sequences). (B) The accuracy of estimated frequencies. Colours

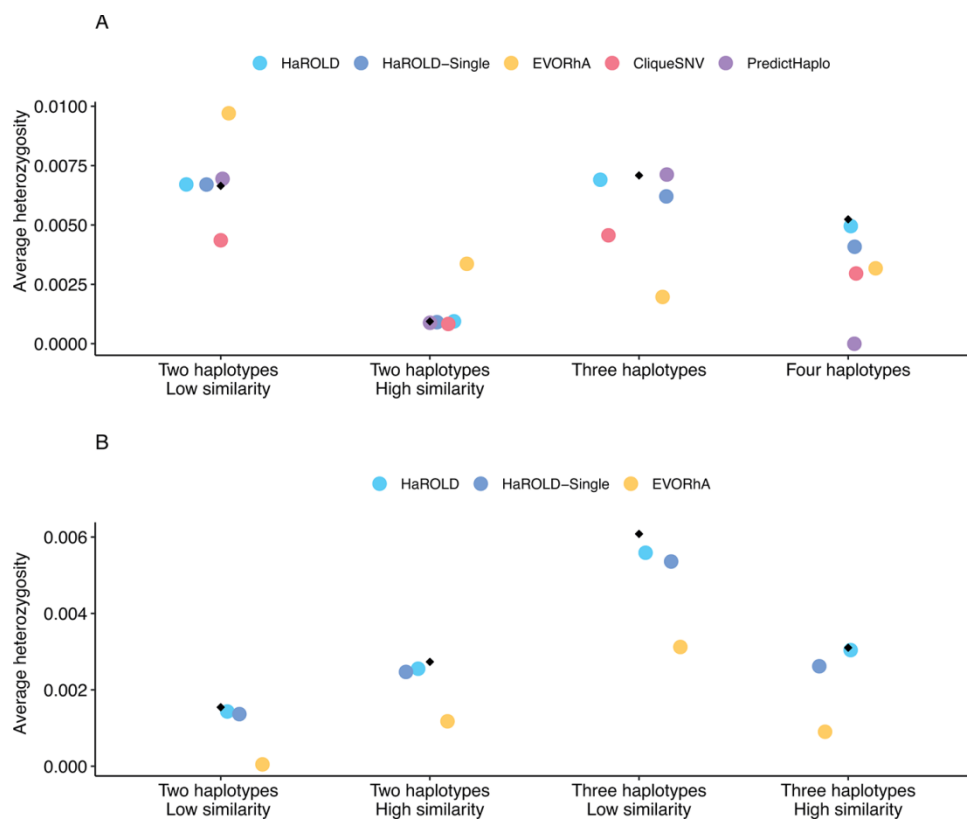
indicate different haplotype reconstruction methods. Each dot represents a sequence from each sample.



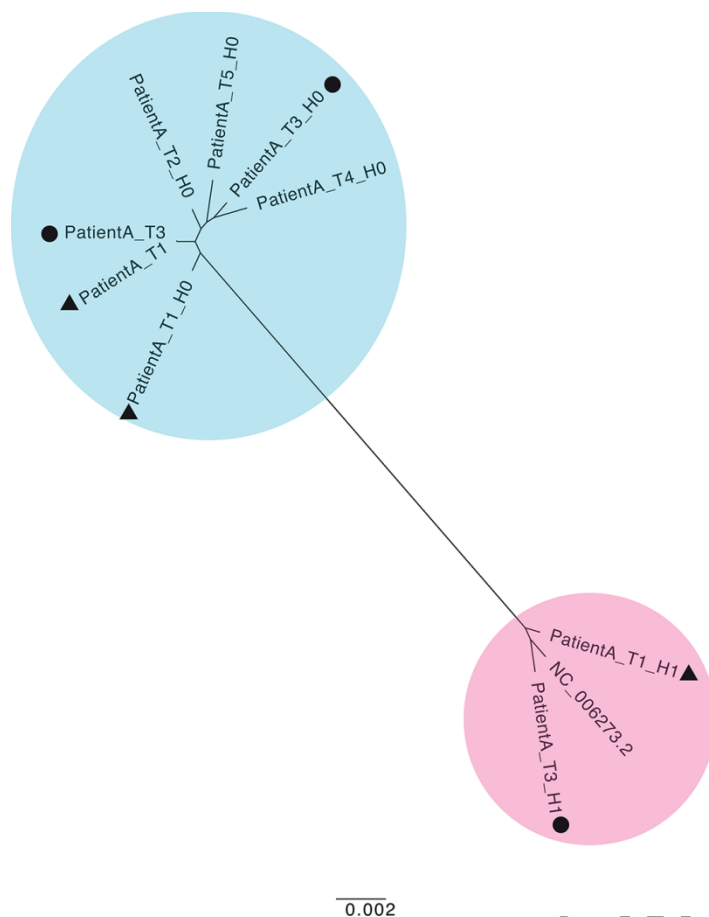
**Figure 3. Violin plots showing the accuracy of haplotype reconstruction in the HCMV test set. (A)** The accuracy of reconstructed sequence (pairwise identity between the actual sequences and reconstructed sequences). **(B)** The accuracy of estimated frequencies. Colours indicate different haplotype reconstruction methods. Each dot represents a sequence from each sample.



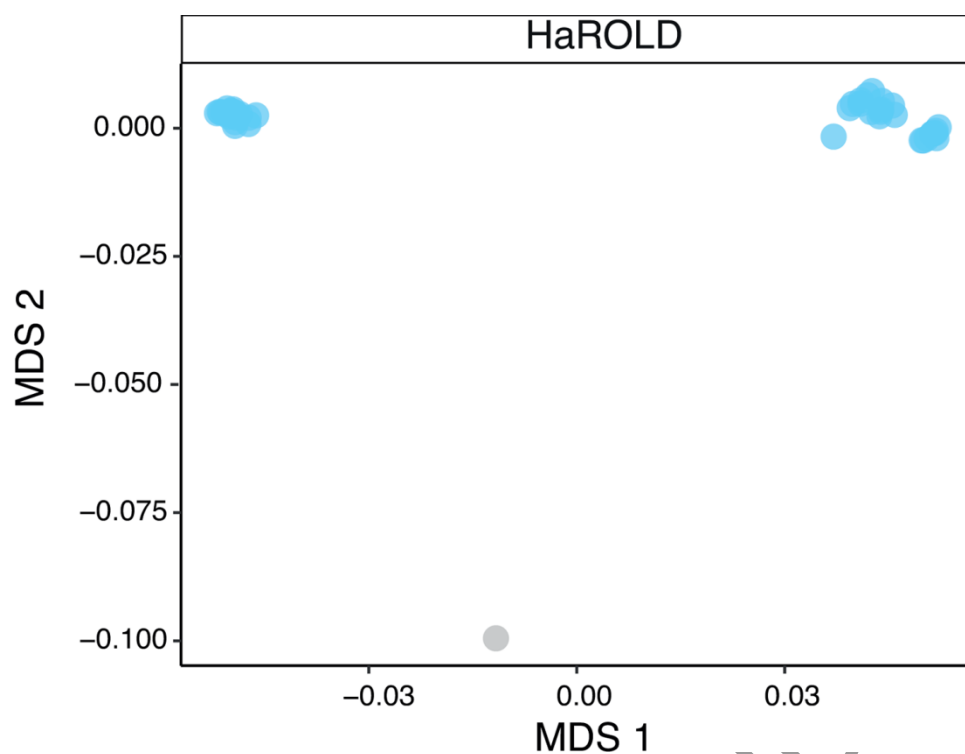
**Figure 4. Accuracy of sample diversity (measured as average heterozygosity) estimations based on reconstructed haplotypes for norovirus dataset (A) and HCMV dataset (B).** Average heterozygosity was estimated for one sample for each set for both norovirus (Set1-Sample4, Set2-Sample3, Set3-Sample1, Set4-Sample4) and HCMV (Set1-sample2, Set2-Sample3, Set3-Sample2, Set4-Sample2) True sequence heterozygosity shown with black diamonds.



**Figure 5. Maximum likelihood phylogenetic tree of reconstructed haplotypes and consensus sequences for the patient infected with HCMV.** Samples of HCMV were taken from patientA at five time points (T1 to T5); reconstructed haplotypes are indicated with H0 and H1. The two haplotypes from the original sample and the consensus sequence from the re-sequenced sample for time point T1 are indicated with a triangle and the corresponding sequences for T3 with a circle. The two viral populations/strains present in the tree are coloured in blue and pink.




**Figure 6. Multi-dimensional scaling (MDS) of HaROLD reconstructed haplotypes for patient infected with norovirus.** Pairwise differences between haplotypes were calculated and used for MDS clustering. The plot shows the first two components. Reference GenBank strain is coloured in grey.



**Figure 7. Maximum likelihood phylogenetic tree of HaROLD reconstructed haplotypes for patient infected with norovirus.** Twelve samples were available for this patient (S1-S12) and were coloured differently using a continuous scale representing time (from blue S1 to red S12). The tips' size indicates the frequency of the haplotype. The black sequence is the GenBank strain used for mapping (tip size set as 50% frequency). Grey transparent circles represent the bootstrap values (1000 bootstraps). Two viral populations were identified represented in orange and grey transparent circles.





**Figure 8. Bar plot of estimated abundances overtime of HaROLD reconstructed haplotypes for patient infected with norovirus. Two viral populations were present coloured in orange and grey, matching the colours on the maximum likelihood phylogenetic tree (Figure 7).**




Table 1.

Norovirus							
Set	Sample composition		Similarity between haplotypes (Percentage identity)				
2 haplotypes Low similarity 5 time points	Original samples		<table border="1"> <tr> <td></td> <td>A</td> </tr> <tr> <td>B</td> <td>98.6</td> </tr> </table>		A	B	98.6
		A					
	B	98.6					
	Synthetic Samples	A: KJ19283    B: MH218631					
	1	0%    100%					
	2	10%    90%					
	3	20%    80%					
4	40%    60%						
5	50%    50%						

2 haplotypes High similarity 5 time points	Original samples			<table border="1"> <tr><td></td><td>A</td></tr> <tr><td>B</td><td>99.7</td></tr> </table>		A	B	99.7													
		A																			
	B	99.7																			
	Synthetic Samples	A: KC175323	B: KJ196279																		
	1	0%	100%																		
	2	10%	90%																		
3	20%	80%																			
4	40%	60%																			
5	50%	50%																			
3 haplotypes 5 time points	Original samples			<table border="1"> <tr><td></td><td>A</td><td>B</td></tr> <tr><td>B</td><td>99.3</td><td></td></tr> <tr><td>C</td><td>98.9</td><td>98.6</td></tr> </table>		A	B	B	99.3		C	98.9	98.6								
		A	B																		
	B	99.3																			
	C	98.9	98.6																		
	Synthetic Samples	A: KC631827	B: KJ196283		C: MH218631																
	1	20%	30%		50%																
2	40%	30%	30%																		
3	60%	30%	10%																		
4	80%	20%	0%																		
5	70%	30%	0%																		
4 haplotypes 5 time points	Original samples				<table border="1"> <tr><td></td><td>A</td><td>B</td><td>C</td></tr> <tr><td>B</td><td>99.7</td><td></td><td></td></tr> <tr><td>C</td><td>99.4</td><td>99.0</td><td></td></tr> <tr><td>D</td><td>99.3</td><td>98.9</td><td>98.6</td></tr> </table>		A	B	C	B	99.7			C	99.4	99.0		D	99.3	98.9	98.6
		A	B	C																	
	B	99.7																			
	C	99.4	99.0																		
	D	99.3	98.9	98.6																	
	Synthetic Samples	A: KC176323	B: KJ196279	C: KJ196283		D: MH218631															
1	0%	0%	30%	70%																	
2	0%	20%	30%	50%																	
3	0%	30%	30%	40%																	
4	30%	20%	40%	10%																	
5	40%	20%	40%	0%																	

Table 2.

HCMV						
Set	Sample composition	Similarity between haplotypes (Percentage identity)				
2 haplotypes Low similarity 3 time points	Original samples					
	Synthetic Samples	A: KP745652.1    B: KP745644.1				
	1	0%    100%				
	2	10%    90%				
		<table border="1"> <tr><td></td><td>A</td></tr> <tr><td>B</td><td>99.1</td></tr> </table>		A	B	99.1
	A					
B	99.1					

	3	30%	80%																																	
2 haplotypes High similarity 3 time points	<table border="1"> <thead> <tr> <th colspan="3">Original samples</th> </tr> <tr> <th>Synthetic Samples</th> <th>A: KU221098.1</th> <th>B: KT726952.2</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>0%</td> <td>100%</td> </tr> <tr> <td>2</td> <td>20%</td> <td>80%</td> </tr> <tr> <td>3</td> <td>40%</td> <td>60%</td> </tr> </tbody> </table>			Original samples			Synthetic Samples	A: KU221098.1	B: KT726952.2	1	0%	100%	2	20%	80%	3	40%	60%	<table border="1"> <thead> <tr> <th colspan="2">A</th> </tr> </thead> <tbody> <tr> <td>B</td> <td>99.4</td> </tr> </tbody> </table>	A		B	99.4													
Original samples																																				
Synthetic Samples	A: KU221098.1	B: KT726952.2																																		
1	0%	100%																																		
2	20%	80%																																		
3	40%	60%																																		
A																																				
B	99.4																																			
3 haplotypes Low similarity 3 time points	<table border="1"> <thead> <tr> <th colspan="4">Original samples</th> </tr> <tr> <th>Synthetic Samples</th> <th>A: KP745652.1</th> <th>B: KP745644.1</th> <th>C: KP745670.1</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>20%</td> <td>50%</td> <td>30%</td> </tr> <tr> <td>2</td> <td>40%</td> <td>40%</td> <td>20%</td> </tr> <tr> <td>3</td> <td>60%</td> <td>30%</td> <td>10%</td> </tr> </tbody> </table>			Original samples				Synthetic Samples	A: KP745652.1	B: KP745644.1	C: KP745670.1	1	20%	50%	30%	2	40%	40%	20%	3	60%	30%	10%	<table border="1"> <thead> <tr> <th colspan="2">A</th> <th colspan="2">B</th> </tr> </thead> <tbody> <tr> <td>B</td> <td>99.1</td> <td></td> <td></td> </tr> <tr> <td>C</td> <td>99.0</td> <td></td> <td>99.0</td> </tr> </tbody> </table>	A		B		B	99.1			C	99.0		99.0
Original samples																																				
Synthetic Samples	A: KP745652.1	B: KP745644.1	C: KP745670.1																																	
1	20%	50%	30%																																	
2	40%	40%	20%																																	
3	60%	30%	10%																																	
A		B																																		
B	99.1																																			
C	99.0		99.0																																	
3 haplotypes High similarity 3 time points	<table border="1"> <thead> <tr> <th>Synthetic Samples</th> <th>A: KU221098.1</th> <th>B: KT726952.2</th> <th>C: KJ361953.1</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>80%</td> <td>20%</td> <td>0%</td> </tr> <tr> <td>2</td> <td>60%</td> <td>30%</td> <td>10%</td> </tr> <tr> <td>3</td> <td>40%</td> <td>40%</td> <td>20%</td> </tr> </tbody> </table>			Synthetic Samples	A: KU221098.1	B: KT726952.2	C: KJ361953.1	1	80%	20%	0%	2	60%	30%	10%	3	40%	40%	20%	<table border="1"> <thead> <tr> <th colspan="2">A</th> <th colspan="2">B</th> </tr> </thead> <tbody> <tr> <td>B</td> <td>99.4</td> <td></td> <td></td> </tr> <tr> <td>C</td> <td>99.4</td> <td></td> <td>99.4</td> </tr> </tbody> </table>	A		B		B	99.4			C	99.4		99.4				
Synthetic Samples	A: KU221098.1	B: KT726952.2	C: KJ361953.1																																	
1	80%	20%	0%																																	
2	60%	30%	10%																																	
3	40%	40%	20%																																	
A		B																																		
B	99.4																																			
C	99.4		99.4																																	

Table 3.

Norovirus				
	2 haplotypes Low similarity 5 time points	2 haplotypes High similarity 5 time points	3 haplotypes 5 time points	4 haplotypes 5 time points
<b>Harold</b>	40 sec	1 min 4 sec	48 sec	3 min 46 sec
<b>CliqueSNV</b>	13 min 27 sec	20 min 24 sec	7 min 7 sec	13 min 43 sec
<b>PredictHaplo</b>	5 h 17 min	6 h 27 min	4 h 40 min	5 h 4 min
<b>Evohra</b>	18 min	19 min	16 min	20 min
HCMV				

	<b>2 haplotypes Low similarity 3 time points</b>	<b>2 haplotypes High similarity 3 time points</b>	<b>3 haplotypes Low similarity 3 time points</b>	<b>3 haplotypes High similarity 3 time points</b>
<b>Harold</b>	39 min	36 min 28 sec	25 min	26 min
<b>Evohra</b>	6 min	5 min	8 min	6 min

ACCEPTED MANUSCRIPT