Revisiting the recombinant history of HIV-1 group M with dynamic network community detection

Abayomi S Olabode¹, Garway T Ng¹, Kaitlyn E Wade^{1,2}, Mikhail Salnikov⁴, David W Dick³, and Art FY Poon¹⁻⁴

¹Department of Pathology & Laboratory Medicine, Western University, London, Canada;

²Department of Computer Science, Western University, London, Canada;

³Department of Applied Mathematics, Western University, London, Canada;

⁴Department of Microbiology & Immunology, Western University, London, Canada

Abstract

A new abundance of full-length HIV-1 genome sequences provides an opportunity to revisit the standard model of HIV-1/M diversity that clusters genomes into largely non-recombinant subtypes, which is not consistent with recent evidence of deep recombinant histories for SIV and other HIV-1 groups. Here we develop an unsupervised non-parametric clustering approach, which does not rely on predefined non-recombinant genomes, by adapting a community detection method developed for dynamic social network analysis. We show that this method (DSBM) attains a significantly lower mean error rate in detecting recombinant breakpoints in simulated data (quasibinomial GLM, $P < 8 \times 10^{-8}$), compared to other reference-free recombination detection programs (GARD, RDP4 and RDP5). Applied to a representative sample of $n = 525$ actual HIV-1 genomes, we determined $k = 25$ as the optimal number of DSBM clusters, and used change point detection to estimate that at least 95% of these genomes are recombinant. Further, we identified both known and novel recombination hotspots in the HIV-1 genome, and evidence of inter-subtype recombination in HIV-1 subtype reference genomes. We propose that clusters generated by DSBM can provide an informative new framework for HIV-1 classification.

Introduction

Understanding the global epidemiology of human immunodeficiency virus type 1 (HIV-1) is contingent on our ability to accurately reconstruct the origin of the virus [1]. The origin of different lineages of HIV-1 has been traced back to multiple zoonotic transmissions of simian immunodeficiency virus from chimpanzees (SIVcpz) directly or through an intermediate host to human populations in West and Central

1 Africa [2, 3]. SIVcpz, on the other-hand, originated through multiple cross-species transmissions of
 2 SIV from other primates to chimpanzees, with extensive recombination between two or more ancestral
 3 variants [4]. SIVs in other non-human primate species also have evidence of recombinant origins [5], and
 4 at least 13 recombinant breakpoints and 14 host-switch events have been identified among SIV lineages
 5 overall [6]. Furthermore, there is phylogenetic evidence that the ancestor of HIV-1 group N was a mosaic
 6 combination of the HIV-1/M ancestor and an SIV lineage [2]. Taken together, these findings imply
 7 that recombination has been an important evolutionary force in the deep evolutionary history of primate
 8 lentiviruses.

9 Currently, HIV-1 is classified into four major groups (M, N, O, P). Group M, which is responsible for
 10 the global pandemic, is further subdivided into nine ‘pure’ subtypes (A-D, F-H, J and K) that are defined
 11 by substantial genetic divergence (about 10% to 25%) and/or bootstrap support for monophyletic clades
 12 [7]. The first evidence of the existence of recombinant HIV-1 genomes was documented in the late 1980s
 13 [8]. HIV-1 exhibits a high rate of recombination that is driven by obligate template switching during the
 14 reverse transcription stage of its replication cycle [9] and by co-infection at the cellular [10] and host
 15 levels [11]. Consequently, there are presently close to 100 circulating recombinant forms (CRFs) docu-
 16 mented by the Los Alamos National Laboratory (<http://www.hiv.lanl.gov/>), where each CRF comprises a
 17 mosaic of two or more pure HIV-1 subtypes that has been sampled from at least three epidemiologically-
 18 unlinked individuals [12]. There is also a large number of unique recombinant forms (URFs) that do not
 19 meet the latter criterion [13]. Nevertheless, the prevailing view of HIV-1 diversity is that the majority
 20 of virus genomes can be classified into one of the subtypes or a relatively small number of CRFs that
 21 have reached a high level of global or regional prevalence, such as CRF01_AE (southeast Asia [14]) or
 22 CRF07_BC (China [15]). Thus, inter-subtype recombination is considered to be an infrequent event,
 23 such that new infections can be routinely classified into one of the subtypes or CRFs by sequencing a
 24 specific region of the genome, such as the HIV-1 *pol* gene that encodes the major targets of antiretroviral
 25 therapy.

26 Increasing deployment of whole-genome sequencing technologies for HIV-1 around the world is provid-
 27 ing a clearer picture of the prevalence and diversity of recombinant HIV-1 genomes [16]. For instance,

1 molecular clock estimates of the origin of the HIV-1 pandemic [17], or the time to the most recent ances-
 2 tor (tMRCA), may be inaccurate due to a deep recombinant history, such that a single phylogenetic tree is
 3 not an adequate representation of discordant evolutionary histories relating different regions of the virus
 4 genome. In previous work [18], we used a sliding-window molecular clock analysis of near full-length
 5 HIV-1 genomes to show that different regions of the virus genome yield significantly different estimates
 6 of the tMRCA. This result supports the hypothesis that HIV-1 group M (HIV-1/M) may have a deep re-
 7 combinant origin that involved at least two genome fragments with different evolutionary histories. For
 8 example, this result could be explained by the introgression of a genome fragment from an unsampled
 9 lineage with a more distant common ancestor to the diversity of the present day [19].

10 This phylogenetic evidence, together with other findings on the deep recombinant history of SIV [6],
 11 supports the hypothesis that the evolutionary history of HIV-1/M can be improved if the conventional
 12 framework of established ‘pure’ reference subtypes with limited subsequent recombination is revisited.
 13 One of the obstacles to evaluating this hypothesis is that most tools for classifying HIV-1 genomes rely
 14 on a reference set of predefined subtypes, such as COMET [20] or SCUEAL [21]. In other words,
 15 these are supervised classification methods. We propose that the best way to test whether the early
 16 evolutionary history of HIV-1/M is reticulate is to apply an unsupervised clustering method to the global
 17 diversity of virus genomes. For example, GARD (genetic algorithm for recombination detection [22])
 18 is a divisive unsupervised method that attempts to find the optimal partition of the sequence alignment
 19 that maximizes the joint likelihood. Fitting GARD to large numbers of near full-length genomes is
 20 computationally intensive because phylogenies must be reconstructed for each genome segment defined
 21 by every candidate partition. Furthermore, a post-processing step is required to determine support for
 22 a recombination breakpoint from the topological discordance of phylogenies. In addition, there are a
 23 number of recombination detection heuristic methods that do not require reference sequences, several of
 24 which are implemented in the software package RDP (recombination detection program) [23]. However,
 25 this package is only distributed as binaries compiled for the Windows operating system.

26 This study endeavours to reconstruct the extent of recombination across the entire HIV-1/M genome.
 27 Here we describe a new unsupervised non-parametric approach to this problem based on adapting a

community detection method that was developed for the analysis of dynamic social networks [24]. Community detection consists of grouping or partitioning nodes (vertices) of a network graph into the same community based on their relative edge density [25, 26]. We note that a connected component in the network, which is often referred to as a ‘cluster’ in the context of genetic epidemiology [27], may comprise multiple communities. Likewise, network communities are often referred to as clusters. To avoid confusion, we will explicitly refer to subgraphs in which there are no edges to external nodes as connected components instead of clusters.

A network in which the distribution of edges among nodes remains constant over the observational time period is a static network [28]. However, real networks are dynamic as relationships evolve over time, represented by the addition and/or removal of network edges [29, 30]. A dynamic network can be represented by a series of static networks that each capture the state of the system at a given point in time. In the context of a dynamic network, communities are a group of nodes that are stably connected by a relatively high density of edges over time [31]. The development of methods to detect network communities is an active area of research with a broad domain of application, including biological sciences, for extracting patterns from complex relational data [26].

In this study, we use genetic distances between sequences to generate a network or graph, where each node represents a virus genome, and each edge indicates that the respective genomes have a distance below some threshold. We assume that communities in the resulting graph correspond to phylogenetic clusters such as subtypes. When distances are calculated from different regions of the genome, a node may switch membership from one community to another, which can correspond to a recombination event. This is analogous to an individual in a social network switching affiliations from one community to another. We use stochastic block modelling to detect the community structure of the genetic similarity graph, and employ a recently described [32] expectation maximization algorithm to estimate the ‘migration’ rates between communities. Stochastic block models (SBMs) are one of the most widely utilized class of models for community detection in networks [24]. Individuals belong to one of K latent communities, and the probability of an edge between individuals is determined only by community membership, *e.g.*, $P_{ii} > P_{ij}$ for $i, j \in \{1, \dots, K\}$. SBMs are designed to uncover hidden structural features in complex

networks by clustering nodes based on similar or shared attributes [33], while dynamic stochastic block models (DSBMs) uncover these hidden data structures as a dynamic network changes over discrete time [34]. We propose to adapt DSBMs as an unsupervised method to characterize the effect of recombination on the evolutionary history of HIV-1 genomes.

Methods

Data processing

A total of $n = 3,900$ near full length ($>8,000$ nt) HIV-1/M genomes, manually curated from our previous study [18] were used in this study. Briefly, we constructed a reduced multiple sequence alignment based on the pairwise alignment of sequences against a consensus genome sequence, discarding insertions relative to this reference to filter out regions of relatively low evolutionary homology. To minimize the computing time of subsequent steps in our analysis, we selected $n = 550$ sequences from this alignment by progressively removing genomes with the shortest genetic distances to other genomes in the dataset. Next, we realigned the remaining sequences using MAFFT (version 7.271) [35], then partitioned the alignment into sliding windows of 500nt at steps of 100nt, resulting to a total of 82 alignment subsets covering HXB2 nucleotide coordinates 790 to 9465. (We arrived at these parameter settings after some preliminary tests varying window and step sizes.) Sequences that had deletions spanning more than 20% in one or more subsets were excluded, resulting in a final total of 525 sequence subsets. We used the SCUEAL (Subtype Classification Using Evolutionary ALgorithms) method implemented in HyPhy to classify these genomes with respect to the HIV-1/M subtype reference sequences obtained from the Los Alamos National Laboratory (LANL) HIV Sequence Database (<http://www.hiv.lanl.gov>).

Community detection

The aim of this study was to characterize the extent of recombination over a series of networks as snapshots of the dynamic system over time, throughout the evolutionary history of HIV-1/M. We are adapting DSBMs to this problem by drawing an analogy between time and the length of the HIV-1 genome, with recombination breakpoints as discrete events. If recombination is relatively infrequent, then genomes should largely cluster into pure subtypes that diverge over time from a common ancestor. With abun-

1 dant recombination, however, we would expect sequences to switch frequently between communities. To
 2 construct a graph, we used the *dist.dna* function of the R package *ape* [36] to compute the Tamura-Nei
 3 [37] (TN93) distances between every pair of sequences within each of the 82 windows extracted from the
 4 alignment. We imported the resulting distance matrices into R to generate undirected graphs using the
 5 package *igraph* [38].

6 We decided not to use a fixed threshold distance for all windows because rates of molecular evolution
 7 vary substantially from one part of the HIV-1 genome to another. Instead, we applied different quantile
 8 thresholds (*e.g.*, the upper 25%) to the observed distributions of TN93 distances for each window, and
 9 used simulation experiments (see below) to assess which quantile yielded the most informative graphs.
 10 Finally, to reconstruct the distribution of recombination events in the full alignment, we used the R
 11 package *dynsbm* [39] to fit a DSBM to the series of graphs. We used the integrated completed likelihood
 12 (ICL) criterion [40] to determine the optimal number of latent groups (K). The ICL criterion is often
 13 used in model-based clustering applications to select the appropriate number of communities [41, 42].

14 To predict recombination breakpoints from the DSBM outputs, we used the R package *change point* to
 15 perform change point detection on cluster assignments along the sequence of windows for each genome.
 16 We dropped missing values from each sequence up to a maximum tolerance of 10 missing values. Fur-
 17 thermore, we evaluated both the default ‘at most one change’ (AMOC) method [43] and the more recently
 18 described pruned exact linear time (PELT) algorithm [44] with a minimum segment length of 10 win-
 19 dows. In both cases, we used the default modified Bayesian information criterion (MBIC) to penalize the
 20 addition of change points [45].

21 **Recombination simulation**

22 Simulating sequence evolution is an important tool for validating new methods in phylogenetic analysis
 23 by evaluating their accuracy on data with known parameters. To calibrate our DSBM method, we evalu-
 24 ated its prediction accuracy on simulations under varying parameter settings. In addition, we compared
 25 its performance to other recombination detection methods that do not require reference genomes, *i.e.*,
 26 GARD and RDP. First, we created a basic and simplified recombination scenario by selecting four ref-

1 erence sequence each of ‘pure’ HIV-1/M subtypes curated by LANL <http://www.hiv.lanl.gov/> A, B, C,
2 and D (n=16) and aligning them. For each breakpoint category a uniform breakpoint location and frag-
3 ment was selected and concatenated for both parent and child/children (recombinant) trees as follows:
4 4500nt, 3000nt, 2250nt for 1, 2 and 3 breakpoints respectively. We then evaluated our ability to recon-
5 struct these simulated recombination breakpoints. This ‘post-hoc’ simulation method is similar to the one
6 used to validate the subtyping and recombination detection program COMET [20], which is reference-
7 dependent. Because one of our objectives was to screen the full diversity of HIV-1 for recombination, we
8 repeated this first simulation method for a larger set ($n = 37$) of subtype reference sequences including
9 representatives of subtypes A-H, J and K.

10 Furthermore, we generated more realistic simulations in which recombination events could be distributed
11 throughout the evolutionary history of the ‘observed’ sequences. First, we used IQ-TREE (version
12 1.3.11.1) [46] to reconstruct a maximum likelihood tree relating the subtype reference sequences used in
13 the previous method. We rooted this tree using the sample collection dates by root-to-tip regression using
14 the *rtt* function in the R package *ape* [36]. Next, we used BEAST (version 1.10.4) [47] to sample time-
15 scaled phylogenies from the posterior distribution, using the Tamura-Nei (TN93) nucleotide substitution
16 model, with rate distribution modeled by a gamma distribution with 4 rate categories; an uncorrelated
17 lognormal clock model ($\mu = 1$, $\sigma = 0.33$); and a skyline tree prior with 10 population sizes. In addition,
18 we set the prior distribution for the time to the most recent common ancestor to a normal distribution with
19 $\mu = 82$ years and $\sigma = 4.1$, based on recent estimates of the origin of HIV-1/M [1]. We ran a single chain
20 sample for 10^8 steps, discarded the first 10^7 steps as burn-in, and generated a maximum clade credibility
21 (MCC) tree from the remainder using TreeAnnotator (version 1.10.4).

22 We used a Python script to update the MCC tree with recombination events by switching random nodes
23 that span a randomly selected time. To generate 1, 2 or 3 recombination breakpoints, we selected a ran-
24 dom point in time to prune and regraft subtrees on two randomly selected extant branches. We assumed
25 that the leftmost portion of the alignment was always related by the original tree, and that the portion
26 of the alignment past the first breakpoint was related by the modified tree. We repeated this process
27 for additional breakpoints by progressively modifying the tree to the immediate left of the breakpoint to

1 relate sequence fragments to the right. In all, we generated 100 simulations of 1, 2 and 3 breakpoints for
 2 a total of 300 trees. The evolution of a 9,000 nt sequence at the root was simulated along each tree using
 3 INDELible (v1.0.3) [48] under a codon substitution model with rate variation modeled by a gamma dis-
 4 tribution ($\alpha = 1.5$, $\beta = 3$) discretized into 50 rate categories, and a transition/transversion ratio $\kappa = 8.0$.
 5 The input tree was rescaled such that the expected number of substitution events per codon was 3.22,
 6 which we derived from the total length of the maximum clade credibility tree (643.5 years), clock rate
 7 estimate (0.00167 substitutions/nt/year) and an adjustment of 3 nucleotides per codon.

8 We used the sequence alignments produced by both simulation methods to evaluate different recombina-
 9 tion detection methods, including the DSBM community detection method. We ran GARD (genetic algo-
 10 rithms for recombination detection) method in HyPhy (version 2.3.11) [22] with the HKY85 nucleotide
 11 substitution model and no rate variation. Since this program required a message passing interface (MPI)
 12 parallel computing environment, we ran this analysis with 8 threads. In addition, we evaluated both ver-
 13 sions 4 and 5 of the program RDP (recombination detection program) [23, 49] using the default settings
 14 and with the sequence type set to ‘linear’.

15 To measure the computing times of these different methods, we generated test sets by randomly sampling
 16 sequences with replacement from the larger set of 37 reference genomes, and added random mutations
 17 at 0.1% of positions in each genome on average. Through this process, we generated 10 replicates of 50,
 18 100 and 200 sequences for a total of 30 test sets, and then ran each set through DSBM, GARD, RDP4 and
 19 RDP5. Since the RDP programs are only released as Windows binary executables, and GARD requires
 20 an MPI-enabled environment, it was not feasible to run all tests in the same computing environment.
 21 Hence, we focused on characterizing the time complexity of each method (relative change in time with
 22 increasing data).

23 Results

24 Our hypothesis is that community detection with a dynamic stochastic block model (DSBM) [32] can be
 25 adapted for the unsupervised detection of recombination events from variation among HIV-1 genomes.
 26 We draw an analogy between latent communities in a social network and genetic clusters, such as the

clusters that are presently labeled as the HIV-1 subtypes. Hence, a recombination event is analogous to moving from one network community to another. To briefly summarize our approach, we first generate a series of networks by partitioning a multiple alignment of HIV-1 genomes into ‘sliding windows’ of a fixed width and step size, and then compute a genetic distance (TN93) between all pairs of sub-sequences within each window. A network (undirected graph) can be derived from each distance matrix by applying a distance threshold, below which an edge is drawn between the respective nodes (sequences). Finally, we applied the implementation of DSBMs in R [39] to the resulting series of graphs to parameterize the transition matrix modelling the movement of individuals between network communities. For brevity, we will refer to this method as DSBM.

Simulation analysis

To evaluate the accuracy of this approach, we first used a simple method to generate alignments of recombinant sequences by combining fragments from one or more HIV-1/M subtype reference genomes at pre-determined positions (breakpoints). We also used this simulation experiment to optimize the threshold for converting TN93 distances to graphs. DSBM performed best at a TN93 threshold of 40%, where the recombinant fragments were correctly identified in 12 out of 15 replicate simulations with an average accuracy of 99% (Supplementary Figure S1). The remaining 3 of the 15 simulations consist of recombination between subtypes B and D. DSBM did not detect recombination events between subtypes B and D, simply because the method tends to classify subtypes B and D sequences as one group. Next, we generated more realistic simulations by grafting recombination events into a time-scaled phylogeny reconstructed from HIV-1 genomes with prior information [1]. Starting from the left of the alignment, we pruned and regrafted subtrees by selecting two extant branches at a random at given time points to produce a new tree for sequence fragments past the next breakpoint. We tested the ability of DSBM to accurately capture the recombination breakpoints in these simulations (Figure 1A). Overall, DSBM was significantly more accurate (measured by error percentage) than the three other methods on these data (quasibinomial GLM, $t > 5.4$, $P < 7.5 \times 10^{-8}$). Error rates increased significantly with the number of actual recombination breakpoints in the data ($t = 15.6$, $P < 10^{-12}$).

In addition, we evaluated the sensitivity of DSBM to varying the sizes of windows (250, 500, 750, 1000nt)

Figure 1: (A) Box-and-whisker plots summarizing the percent error of recombination detection methods on sequences simulated from a time-scaled HIV-1/M phylogeny. Percent error was calculated by mapping outputs from each program to the same outcome space of cluster assignments by sequence and nucleotide position. (B) Slopegraphs summarizing the computing times of different recombination detection programs on alignments of 50, 100 and 200 sequences, derived from $n = 37$ HIV-1 subtype reference sequences with random mutations. We evaluated 10 replicate simulations per sample size. DSBM-8 and -12 correspond to dynamic stochastic block model analyses with 8 and 12 clusters, respectively. Note that DSBM computing times do not include the fast calculation of Tamura-Nei (TN93) distances. GARD was run in an MPI environment with 8 cores.

1 and steps (50, 100, 200, 300nt), which determines the number and stability of the graphs generated from
2 the genome alignment. Error percentage was not significantly associated with window size (quasibino-
3 mial GLM, $t = -0.33$, $P = 0.74$) or step size ($t = 0.75$, $P = 0.46$), although we noted a slight tendency
4 for error to increase with window size for simulations with 1 breakpoint, and to decrease with window
5 size for 3 breakpoints (Supplementary Figure S2).

6 Finally, we measured the runtimes for processing varying numbers of sequences derived from HIV-1
7 subtype reference genomes with DSBM, GARD, RDP4 and RDP5 (Figure 1B). Since we were more
8 interested in computing time than accurately inferring unknown recombination events in this experiment,
9 we arbitrarily set the number of DSBM clusters to 8 and 12. Computing times for DSBM were compara-
10 ble to RDP4, whereas GARD and RDP5 were substantially faster; however, it is difficult to compare the
11 latter two since GARD requires a parallel computing (MPI) environment. Overall, these results indicate
12 that the expected computing times tend to increase faster than linearly with the number of sequences, and
13 increase with the number of clusters for DSBM. We also noted an unusual trend in GARD computing

times, with a faster times obtained for sample sizes above 50 sequences. On examining the output files, we determined that GARD efficiently rejected all models with recombination breakpoints for the larger datasets. In contrast, GARD runs on datasets with 50 sequences were unable to reject the presence of recombination and were burdened with 677.8 candidate models (partitions of the alignment among trees) on average.

Application to HIV-1/M genomes

Having validated the DSBM method on simulated data in comparison to other unsupervised recombination detection methods, we next applied our method to $n = 525$ full-length HIV-1/M genome sequences to characterize the role of recombination in the evolutionary history of this virus (Figure 2). These sequences were selected from all available full-length HIV-1 genomes to maximize the representation of global diversity in a reduced data set. The integrated completed likelihood (ICL) criterion, used to determine the optimal number of clusters for DSBM, was maximized at 25 clusters (Supplementary Figure S4).

The results of our analysis are summarized in Figure 2. Figure 2A displays a heatmap summarizing the transition rate matrix from the DSBM analysis. This rate matrix was not constrained to be symmetric, since the outflow from cluster 1 to 2 is not necessary balanced by a return flow from cluster 2. Some clusters have substantial transition rates (darker shades) between them, such as clusters 8 and 18, 11 and 17, or 16 and 25. This implies a hierarchical structure to clusters akin to the sub-subtypes of HIV-1. The frequencies of cluster memberships were consistent across the length of the HIV-1 genome (Supplementary Figure S5). To better understand the associations between these clusters, we mapped the original HIV-1 subtype annotations of the genome sequences to these data. We found that windows from genomes labelled as subtype A1 ($n = 33$) were predominantly assigned to clusters 4, 13 and 24; subtype B ($n = 56$) was associated with clusters 3, 8, 15 and 18; and subtype C ($n = 66$) to clusters 11, 16, 17 and 25 (Supplementary Figure S6). Within these major subtypes, cluster assignments tended to become more variable past genome window 50 (Supplementary Figure S7), in association with the *env* gene. In addition, subtypes D, G, F1 and the circulating recombinant form (CRF) 01AE associated strongly with clusters 14, 10, 2 and 7, respectively. Clusters 1 and 9 associated with subtypes J, H and K, as well as

Figure 2: Summary of results from DSBM analysis of $n = 525$ HIV-1/M genomes. (A) A heatmap of the transition rate matrix estimated from the sequence of networks generated by a sliding window analysis of the HIV-1/M genome alignment. The highest rates (darkest shaded cells) corresponded to the rate that an individual remains in the same network community (cluster) on moving from one genomic window to the next. (B) A heatmap displaying the cluster assignments for $n = 525$ genomes, partitioned into $n = 82$ windows of 500nt in steps of 100nt. The colour map corresponds to the stacked barplot legend in (B). White cells indicate that DSBM failed to assign the respective genome to a cluster in that window. Distinct transitions between colours along a row implies recombination. A color-blind accessible version is provided as Supplementary Figure S3. (C) Stepcharts depicting the cluster assignments for the first 16 genomes across 82 windows, where we used the hierarchical clustering permutation order from (A) to minimize the vertical distance between clusters related by higher transition rates. Line segments (red) indicate the location of change points, as determined by a conservative 'at most one change' (AMOC) method.

1 genomes that were annotated as ‘complex’ inter-subtype recombinants.

2 Figure 2B displays the cluster memberships for all 525 genomes in 82 windows of 500nt in steps of 100nt.

3 We mapped cluster memberships to a colour gradient such that clusters with similar colours tend to have

4 relatively higher transition rates between them. Under the standard concept of ‘pure’ HIV-1 subtypes,

5 genomes should tend to maintain the same colouration across its length. Indeed, we see an overall

6 stratification of colours in this heatmap (Figure 2B). However, if we examine the cluster assignments for

7 individual genomes (Figure 2C), we see traces that are consistent with recombination, such as genomes

8 2, 3 and 14. Other traces are too noisy to be interpreted visually. Therefore, we used two different

9 change point detection algorithms to quantify the extent of recombination from the entire set of traces.

10 The AMOC method classified 465 (95.9%) genomes as recombinant (at least one change point) and 20

11 as non-recombinant (40 genomes were excluded due to excessive missing values). This result was robust

12 to varying the minimum segment length; for example, AMOC only classified 47 (9.7%) genomes as

13 non-recombinant at a cutoff of 20 windows.

14 Next, we used the less restrictive PELT method to estimate the number of breakpoints per genome (Fig-

15 ure 3, left). As expected, the mean number of breakpoints per genome increased with shorter minimum

16 segment lengths. At a minimum length of 5 windows, for instance, we inferred a mean of 6.8 (IQR 5-

17 12) breakpoints. On the other hand, the distribution of breakpoints across windows was less sensitive to

18 varying the minimum segment length, with significant positive correlations between distributions (Spear-

19 man’s $\rho > 0.58$, $P < 4.7 \times 10^{-7}$; Figure 3, right). Local peaks in the frequencies of breakpoints were

20 robust to varying minimum segment length. Consistent with previous work [50], we observed recombina-

21 tion ‘hotspots’ associated with the 5’ and 3’ ends of the *env* gene. We also observed a distinct and robust

22 peak associated with the 3’ end of the *pol* gene, and a sharp increase in the frequency of breakpoints

23 within or upstream of *gag*. Genomes that were classified as recombinant by SCUEAL had significantly

24 higher numbers of breakpoints assigned by DSBM, and this concordance was robust to varying minimum

25 segment lengths (Wilcoxon rank-sum test, $P < 6.3 \times 10^{-3}$; Supplementary Figure S9). Furthermore, we

26 found no significant association between the number of predicted breakpoints and the year of sample

27 collection, which ranged from 1983 to 2016 with a median of 2006 (linear model $t = 0.86$, $P = 0.39$;

Figure 3: Distribution of DSBM predicted recombination breakpoints in HIV-1/M genomes. (left) Histogram of the number of breakpoints per genome, where breakpoints were extracted from DSBM cluster assignments by the PELT change detection method. We varied the minimum segment length to 3, 5 and 10 windows (see colour legend). (right) Smoothed splines summarizing the distribution of predicted breakpoints among sliding windows of the HIV-1 genome alignment. We used the same colours as the histograms to indicate minimum segment lengths. The raw frequency distributions are drawn as step charts in lighter colours. A diagram of the HIV-1 reading frames, mapped from the HXB2 reference coordinates to our alignment, is displayed at the base of the figure.

1 Supplementary Figure S8).

2 **Revisiting the HIV-1 subtype references**

3 In addition to analyzing a large selection of HIV-1/M genome sequences, we applied the DSBM method
 4 to evaluate the HIV-1 subtype reference genomes curated by the Los Alamos National Laboratory HIV
 5 Sequence Database. These reference genomes have long been used as the ‘gold standard’ against which
 6 other genome sequences are evaluated for evidence of recombination. We assessed a subset of the refer-
 7 ence genomes covering HIV-1/M subtypes A-D, F-H, J and K, and excluded reference genomes corre-
 8 sponding to the circulating recombinant forms (CRFs). For this analysis, the ICL criterion selected $k = 6$
 9 as the optimal number of clusters. This is clearly fewer than the recognized number of HIV-1 subtypes.
 10 We observed that subtypes B and D were assigned to the same cluster. Furthermore, genomic windows
 11 from subtypes F and K tended to be assigned to the same cluster, and subtypes G and J to a third cluster.
 12 These cluster assignments are consistent with the placement of the respective subtypes in a phylogeny of
 13 HIV-1/M subtypes and CRFs [51].

14 Employing the same post-processing method as our previous analysis, we detected evidence of recomb-
 15 nation breakpoints in 18 (48.6%) of 37 reference genomes when the minimum segment length was set to
 16 5 windows. The PELT method of change point detection identified 1, 2 and 4 breakpoints in 2, 13 and 3 of
 17 the remaining genomes, respectively (Figure 4). These breakpoints tended to be concentrated in genomes
 18 classified by SCUEAL into subtypes F (sub-subtypes F1 and F2), J and K. We also observed consistent
 19 patterns of recombination in the subtype G genomes spanning windows 42-51 and 67-78. These putative
 20 breakpoints, which are consistent with previous work [19], were not picked up by the PELT method be-
 21 cause the affected clusters were annotated as being similar, even though we did not observe substantial
 22 clustering in transition rates between the 6 clusters in this analysis. These results highlight an important
 23 limitation of change point detection based on changes in the mean.

24 **Discussion**

25 Recombination is a major contributing factor to the extensive genetic variability observed in RNA viruses
 26 [52]. Therefore, it is pertinent to screen for recombination when studying evolution in these viruses [22].

Figure 4: Detection of recombination in curated HIV-1 subtype reference genomes by community detection with a dynamic stochastic block model (DSBM). Predicted breakpoints (using the PELT change point detection method with a minimum segment length of 5 windows) are marked directly on the heatmap with vertical line segments. We used a different colour palette to emphasize the different number of clusters between this analysis ($k = 6$) and our previous analysis of $n = 525$ HIV-1/M genomes ($k = 25$).

Recent advancements in whole genome sequencing is driving the proliferation of large HIV-1/M genome databases [53, 54], providing more opportunities to reconstruct the role of recombination in the evolutionary history of this virus. Here, we have adapted a community detection method from network science (dynamic stochastic block models, DSBM [24, 32]) to detect the residual evidence of past recombination events without requiring a reference set of non-recombinant genomes. Hence, we characterize DSBM as an unsupervised method for detecting recombination.

Compared to the commonly-used unsupervised recombination detection programs GARD, RDP4 [49] and RDP5 [23]. GARD employs a genetic algorithm to explore the model space of placing one or more recombination breakpoints in the alignment, reconstructing a maximum likelihood phylogeny for each interval between breakpoints [22]. Breakpoints are accepted if the increase in likelihood is sufficient to justify increasing the number of model parameters induced by the addition of a discordant phylogeny. Of the unsupervised methods, it is the closest to a parametric model of the evolutionary consequences of recombination. RDP versions 4 and 5, on the other hand, employ a suite of non-parametric recombination detection methods. For example, BOOTSCAN calculates pairwise distances between sequences in sliding windows. Thus, our initial step of mapping a sequence alignment to DSBM is similar to RDP. On the other hand, DSBM operates on a more comprehensive representation of sequence diversity by converting the pairwise distance matrices into a series of undirected graphs. In addition, DSBM employs a hidden Markov model in which the probability of an edge connecting two individuals is determined by their unobservable (latent) and dynamic cluster memberships. Consequently, the placement of a breakpoint at a particular location in one sequence is informed by observed patterns across all sequences.

DSBM was significantly more accurate for detecting recombination in simulated data. However, DSBM is time-consuming to compute and slower than the other methods, in part due to the quadratic time complexity of parameterizing the matrix of transition rates between clusters (communities). Like GARD, the computing time can be ameliorated by running DSBM in a parallel computing environment. Our extension of DSBM involves a number of tuning parameters — namely, a genetic distance threshold, the window size, and the window step size. Our simulation experiments indicated that results from DSBM are not sensitive to varying the window and step size parameters. However, choosing an appropriate

1 distance threshold, which determines how sequence variation maps to network topologies, was critical
2 to our analysis (Supplementary Figure S1). Another limitation of DSBM is that it assumes the rates
3 of transitions between clusters are ‘time’ homogeneous, *i.e.*, that recombination rates between specific
4 lineages are constant through the genome. Since DSBMs are a recent innovation in network science, it
5 may eventually become possible to fit models with time-heterogeneous rates; however, the current model
6 already estimates a large number of parameters from the data.

7 An interesting outcome of our analysis is the putative recombination hotspots in association with the
8 *gag* and *nef* gene sequences, which are adjacent to or overlap the 5’ and 3’ long terminal repeats (LTR),
9 respectively (Figure 3). Based on the composition of the CRF genomes documented by the Los Alamos
10 National Laboratory HIV Sequence Database, breakpoints are often found in these regions. For example,
11 CRF01_AE contains a breakpoint at the 5’ end of the 3’ LTR sequence (HXB2 nucleotide coordinate
12 9086), and CRF02_AG at the 3’ end of the 5’ LTR (coordinate 789). Overall, we observed that 57.1%
13 ($n = 77$) and 45.6% ($n = 92$) of the documented CRFs contain breakpoints upstream of HXB2 coordinates
14 1496 and 8757, respectively (adjusting for CRFs without sequence coverage in these regions). Further-
15 more, recombination in these regions of the HIV-1 genome has been observed as non-CRFs [55, 56]. In
16 a comparative analysis of HIV-1 *gag*, Minin *et al.* [57] detected a recombination hotspot in association
17 with an instability element in the region encoding the capsid protein. However, this hotspot was not
18 reproduced in subsequent work by Archer *et al.* [50], who reported hotspots associated with both ends
19 of the *env* gene. Resolving these differences will require the reconciliation of methods and data sets, or
20 more effectively, experimental validation with an *in vitro* system [58].

21 Previous studies have proposed re-classifying some of the defined HIV-1 subtypes and circulating re-
22 combinant forms. In some cases, this was motivated by the availability of new genomic data [59, 60].
23 For example, a recent analysis of the major HIV-1 genes *gag*, *pol* and *env* culminated in a proposal to
24 further partition of subtypes A and D into sub-subtypes, and to merge subtypes B and D into a single
25 subtype [61]. Moreover, a new non-recombinant HIV-1/M subtype L was proposed almost 30 years after
26 the genomes were first sampled in 1983 and 1990 in what is now the Democratic Republic of the Congo
27 [62]. Our analysis indicates that many of the HIV-1/M subtype reference genomes are actually recom-

binant, which is consistent with previous work [18, 19]. These findings support the idea that the current HIV-1 nomenclature [7], which has served as an important framework for our understanding of HIV-1 diversity and evolution, should be revisited in light of new genomic evidence. Our DSBM analysis of a large sample of HIV-1/M genomes identified $k = 25$ as the optimal number of clusters. In this context, a cluster is roughly analogous to a non-recombinant subtype. However, what we recognize as a subtype may also comprise multiple clusters (Supplementary Figure S7). None of the $n = 525$ in our largest analysis was assigned along its entire length to a single cluster. It is unlikely that expanding the scope of our analysis will yield genomes that are each representative of a single cluster. However, it may be feasible to utilize the cluster assignments from DSBM to generate composite ancestral genome reconstructions that can play the same role as the current HIV-1 subtype reference genomes.

Acknowledgements

This work was supported in part by the Government of Canada through Genome Canada and the Ontario Genomics Institute (OGI-131) and by the Canadian Institutes of Health Research (CIHR grants PJT-155990, PJT-156178, FRN-130609, BOP-149562).

References

- [1] Faria NR, Rambaut A, Suchard MA, Baele G, Bedford T, Ward MJ, et al. The early spread and epidemic ignition of HIV-1 in human populations. *science*. 2014;346(6205):56–61.
- [2] Gao F, Bailes E, Robertson DL, Chen Y, Rodenburg CM, Michael SF, et al. Origin of HIV-1 in the chimpanzee *Pan troglodytes*. *Nature*. 1999;397(6718):436.
- [3] Keele BF, Van Heuverswyn F, Li Y, Bailes E, Takehisa J, Santiago ML, et al. Chimpanzee reservoirs of pandemic and nonpandemic HIV-1. *Science*. 2006;313(5786):523–526.
- [4] Bailes E, Gao F, Bibollet-Ruche F, Courgnaud V, Peeters M, Marx PA, et al. Hybrid origin of SIV in chimpanzees. *Science*. 2003;300(5626):1713–1713.
- [5] Jin M, Hui H, Robertson D, Müller M, Barré-Sinoussi F, Hirsch V, et al. Mosaic genome structure of simian immunodeficiency virus from west African green monkeys. *The EMBO journal*. 1994;13(12):2935–2947.
- [6] Bell SM, Bedford T. Modern-day SIV viral diversity generated by extensive recombination and cross-species transmission. *PLoS pathogens*. 2017;13(7):e1006466.
- [7] Robertson DL, Anderson J, Bradac J, Carr J, Foley B, Funkhouser R, et al. HIV-1 nomenclature proposal. *Science*. 2000;288(5463):55–55.
- [8] Li WH, Tanimura M, Sharp PM. Rates and dates of divergence between AIDS virus nucleotide sequences. *Molecular biology and evolution*. 1988;5(4):313–330.
- [9] Cromer D, Grimm AJ, Schlub TE, Mak J, Davenport MP. Estimating the in-vivo HIV template switching and recombination rate. *Aids*. 2016;30(2):185–192.
- [10] Neher RA, Leitner T. Recombination rate and selection strength in HIV intra-patient evolution. *PLoS Comput Biol*. 2010;6(1):e1000660.
- [11] Redd AD, Quinn TC, Tobian AA. Frequency and implications of HIV superinfection. *The Lancet infectious diseases*. 2013;13(7):622–628.

- 1 [12] Robertson D, Anderson J, Bradac J, Carr J, Foley B, Funkhouser R, et al. HIV-1 nomenclature
2 proposal. *Science* (New York, NY). 2000 April;288(5463):55—56. Available from: [https://doi.org/](https://doi.org/10.1126/science.288.5463.55d)
3 10.1126/science.288.5463.55d.
- 4 [13] Rodgers MA, Vallari AS, Yamaguchi J, Holzmayer V, Harris B, Toure-Kane C, et al. ARCHITECT
5 HIV combo Ag/Ab and realtime HIV-1 assays detect diverse HIV strains in clinical specimens.
6 *AIDS research and human retroviruses*. 2018;34(3):314–318.
- 7 [14] Angelis K, Albert J, Mamais I, Magiorkinis G, Hatzakis A, Hamouda O, et al. Global dis-
8 persal pattern of HIV type 1 subtype CRF01_AE: a genetic trace of human mobility related to
9 heterosexual sexual activities centralized in Southeast Asia. *The Journal of infectious diseases*.
10 2015;211(11):1735–1744.
- 11 [15] Zhang M, Foley B, Schultz AK, Macke JP, Bulla I, Stanke M, et al. The role of recombination in
12 the emergence of a complex and dynamic HIV epidemic. *Retrovirology*. 2010;7(1):1–15.
- 13 [16] Grant HE, Hodcroft EB, Ssemwanga D, Kitayimbwa JM, Yebra G, Esquivel Gomez LR, et al.
14 Pervasive and non-random recombination in near full-length HIV genomes from Uganda. *Virus*
15 *Evolution*. 2020;6(1):veaa004.
- 16 [17] Faria NR, Vidal N, Lourenco J, Raghwan J, Sigaloff KC, Tatem AJ, et al. Distinct rates and
17 patterns of spread of the major HIV-1 subtypes in Central and East Africa. *PLoS pathogens*.
18 2019;15(12):e1007976.
- 19 [18] Olabode AS, Avino M, Ng GT, Abu-Sardanah F, Dick DW, Poon AF. Evidence for a Recombinant
20 Origin of HIV-1 group M from Genomic Variation. *Virus evolution*. 2019;5(1):vey039.
- 21 [19] Abecasis AB, Lemey P, Vidal N, de Oliveira T, Peeters M, Camacho R, et al. Recombination
22 confounds the early evolutionary history of human immunodeficiency virus type 1: subtype G is a
23 circulating recombinant form. *Journal of virology*. 2007;81(16):8543–8551.
- 24 [20] Struck D, Lawyer G, Ternes AM, Schmit JC, Bercoff DP. COMET: adaptive context-based model-
25 ing for ultrafast HIV-1 subtype identification. *Nucleic acids research*. 2014;42(18):e144–e144.

- 1 [21] Pond SLK, Posada D, Stawiski E, Chappey C, Poon AF, Hughes G, et al. An evolutionary model-
2 based algorithm for accurate phylogenetic breakpoint mapping and subtype prediction in HIV-1.
3 PLoS Comput Biol. 2009;5(11):e1000581.
- 4 [22] Kosakovsky Pond SL, Posada D, Gravenor MB, Woelk CH, Frost SD. GARD: a genetic algorithm
5 for recombination detection. Bioinformatics. 2006;22(24):3096–3098.
- 6 [23] Martin DP, Varsani A, Roumagnac P, Botha G, Maslamoney S, Schwab T, et al. RDP5: A computer
7 program for analysing recombination in, and removing signals of recombination from, nucleotide
8 sequence datasets. Virus Evolution. 2020.
- 9 [24] Abbe E. Community detection and stochastic block models: recent developments. The Journal of
10 Machine Learning Research. 2017;18(1):6446–6531.
- 11 [25] Yang B, Liu D, Liu J. Discovering communities from social networks: Methodologies and applica-
12 tions. In: Handbook of social network technologies and applications. Springer; 2010. p. 331–346.
- 13 [26] Khan BS, Niazi MA. Network community detection: A review and visual survey. arXiv preprint
14 arXiv:170800977. 2017.
- 15 [27] Poon AF. Impacts and shortcomings of genetic clustering methods for infectious disease outbreaks.
16 Virus evolution. 2016;2(2):vew031.
- 17 [28] Aynaud T, Guillaume JL. Static community detection algorithms for evolving networks. In: 8th
18 International symposium on modeling and optimization in mobile, Ad Hoc, and wireless networks.
19 IEEE; 2010. p. 513–519.
- 20 [29] Cazabet R, Rossetti G, Amblard F. Dynamic community detection; 2017.
- 21 [30] Chen G, Wang Y, Wei J. A new multiobjective evolutionary algorithm for community detection in
22 dynamic complex networks. Mathematical problems in engineering. 2013;2013.
- 23 [31] Ahmed NK, Berchmans F, Neville J, Kompella R. Time-based sampling of social network activity
24 graphs. In: Proceedings of the eighth workshop on mining and learning with graphs. ACM; 2010.
25 p. 1–9.

- 1 [32] Matias C, Miele V. Statistical clustering of temporal networks through a dynamic stochastic block
2 model. *Journal of the Royal Statistical Society Series B*. 2017;79(4):1119–1141.
- 3 [33] Qiao M, Yu J, Bian W, Li Q, Tao D. Adapting stochastic block models to power-law degree distri-
4 butions. *IEEE transactions on cybernetics*. 2018;49(2):626–637.
- 5 [34] Xu K. Stochastic block transition models for dynamic networks. In: *Artificial Intelligence and*
6 *Statistics*; 2015. p. 1079–1087.
- 7 [35] Katoh K, Misawa K, Kuma Ki, Miyata T. MAFFT: a novel method for rapid multiple sequence
8 alignment based on fast Fourier transform. *Nucleic acids research*. 2002;30(14):3059–3066.
- 9 [36] Paradis E, Claude J, Strimmer K. APE: analyses of phylogenetics and evolution in R language.
10 *Bioinformatics*. 2004;20(2):289–290.
- 11 [37] Tamura K, Nei M. Estimation of the number of nucleotide substitutions in the control region of mi-
12 tochondrial DNA in humans and chimpanzees. *Molecular biology and evolution*. 1993;10(3):512–
13 526.
- 14 [38] Csardi G, Nepusz T, et al. The igraph software package for complex network research. *InterJournal,*
15 *complex systems*. 2006;1695(5):1–9.
- 16 [39] Miele V, Matias C. Revealing the hidden structure of dynamic ecological networks. *Royal Society*
17 *open science*. 2017;4(6):170251.
- 18 [40] Biernacki C, Celeux G, Govaert G. Assessing a mixture model for clustering with the inte-
19 grated completed likelihood. *IEEE transactions on pattern analysis and machine intelligence*.
20 2000;22(7):719–725.
- 21 [41] Baudry JP, et al. Estimation and model selection for model-based clustering with the conditional
22 classification likelihood. *Electronic journal of statistics*. 2015;9(1):1041–1077.
- 23 [42] Bertolotti M, Friel N, Rastelli R. Choosing the number of clusters in a finite mixture model using
24 an exact integrated completed likelihood criterion. *Metron*. 2015;73(2):177–199.

- 1 [43] Hinkley DV. Inference about the change-point in a sequence of random variables. *Biometrika*.
2 1970;57(1):1–17.
- 3 [44] Killick R, Fearnhead P, Eckley IA. Optimal detection of changepoints with a linear computational
4 cost. *Journal of the American Statistical Association*. 2012;107(500):1590–1598.
- 5 [45] Zhang NR, Siegmund DO. A modified Bayes information criterion with applications to the analysis
6 of comparative genomic hybridization data. *Biometrics*. 2007;63(1):22–32.
- 7 [46] Nguyen LT, Schmidt HA, Von Haeseler A, Minh BQ. IQ-TREE: a fast and effective stochas-
8 tic algorithm for estimating maximum-likelihood phylogenies. *Molecular biology and evolution*.
9 2015;32(1):268–274.
- 10 [47] Drummond AJ, Rambaut A. BEAST: Bayesian evolutionary analysis by sampling trees. *BMC*
11 *evolutionary biology*. 2007;7(1):214.
- 12 [48] Fletcher W, Yang Z. INDELible: a flexible simulator of biological sequence evolution. *Molecular*
13 *biology and evolution*. 2009;26(8):1879–1888.
- 14 [49] Martin DP, Murrell B, Golden M, Khoosal A, Muhire B. RDP4: Detection and analysis of recom-
15 bination patterns in virus genomes. *Virus evolution*. 2015;1(1).
- 16 [50] Archer J, Pinney JW, Fan J, Simon-Loriere E, Arts EJ, Negroni M, et al. Identifying the important
17 HIV-1 recombination breakpoints. *PLoS Comput Biol*. 2008;4(9):e1000178.
- 18 [51] Tebit DM, Arts EJ. Tracking a century of global expansion and evolution of HIV to drive under-
19 standing and to combat disease. *The Lancet infectious diseases*. 2011;11(1):45–56.
- 20 [52] Worobey M, Holmes EC. Evolutionary aspects of recombination in RNA viruses. *Journal of General*
21 *Virology*. 1999;80(10):2535–2543.
- 22 [53] Kuiken C, Korber B, Shafer RW. HIV sequence databases. *AIDS reviews*. 2003;5(1):52.
- 23 [54] Pillay D, Herbeck J, Cohen MS, de Oliveira T, Fraser C, Ratmann O, et al. PANGAEA-HIV: phylo-
24 genetics for generalised epidemics in Africa. *The Lancet Infectious Diseases*. 2015;15(3):259–261.

- 1 [55] Blackard JT, Renjifo BR, Mwakagile D, Montano MA, Fawzi WW, Essex M. Transmission of
2 human immunodeficiency type 1 viruses with intersubtype recombinant long terminal repeat se-
3 quences. *Virology*. 1999;254(2):220–225.
- 4 [56] Neogi U, Sood V, Goel N, Wanchu A, Banerjea AC. Novel HIV-1 long terminal repeat (LTR)
5 sequences of subtype B and mosaic intersubtype B/C recombinants in North India. *Archives of*
6 *virology*. 2008;153(10):1961.
- 7 [57] Minin VN, Dorman KS, Fang F, Suchard MA. Phylogenetic mapping of recombination hotspots
8 in human immunodeficiency virus via spatially smoothed change-point processes. *Genetics*.
9 2007;175(4):1773–1785.
- 10 [58] Baird HA, Galetto R, Gao Y, Simon-Loriere E, Abreha M, Archer J, et al. Sequence determi-
11 nants of breakpoint location during HIV-1 intersubtype recombination. *Nucleic acids research*.
12 2006;34(18):5203–5216.
- 13 [59] Yebra G, Frampton D, Gallo Cassarino T, Raffle J, Hubb J, Ferns RB, et al. A high HIV-1 strain
14 variability in London, UK, revealed by full-genome analysis: Results from the ICONIC project.
15 *Plos one*. 2018;13(2):e0192081.
- 16 [60] Rubio-Garrido M, González-Alba JM, Reina G, Ndarabu A, Barquín D, Carlos S, et al. Current
17 and historic HIV-1 molecular epidemiology in paediatric and adult population from Kinshasa in the
18 Democratic Republic of Congo. *Scientific reports*. 2020;10(1):1–13.
- 19 [61] Désiré N, Cerutti L, Le Hingrat Q, Perrier M, Emler S, Calvez V, et al. Characterization update
20 of HIV-1 M subtypes diversity and proposal for subtypes A and D sub-subtypes reclassification.
21 *Retrovirology*. 2018;15(1):1–7.
- 22 [62] Yamaguchi J, Vallari A, McArthur C, Sthreshley L, Cloherty GA, Berg MG, et al. Brief report:
23 complete genome sequence of CG-0018a-01 establishes HIV-1 subtype L. *Journal of acquired*
24 *immune deficiency syndromes (1999)*. 2020;83(3):319.

Figure S2: Box-and-whisker plots showing the percentage error reconstructing simulated recombination breakpoints using DSBM with varying window (left) and step (right) sizes.

Figure S3: Color-blind accessible (viridis palette) version of heatmap summarizing DSBM cluster assignments for HIV-1/M genome sequences (Figure 2B).

Figure S4: Profile of integrated composite likelihood (ICL) values as a function of k , the number of network communities, in DSBM analyses of graphs derived from the HIV-1/M genome windows. The maximal value ($k = 25$) is highlighted in red. Lower values of k were truncated from the plot to clarify differences in ICL in the neighbourhood of the maximal value.

Figure S5: This stacked barplot summarizes the frequencies of cluster memberships across genome windows. We mapped a colour gradient to the $k = 25$ DSBM clusters using the permutation order from the hierarchical clustering used to generated the heatmap, such that clusters connected by higher transition rates are assigned similar colours. The bars have different heights because a varying number of genomes could not be assigned to any cluster (index 0) in a given window.

Figure S6: Heatmap summarizing the distribution of cluster assignments among HIV-1 subtypes. Subtype designations were generated using the SCUEAL (subtype classification using evolutionary algorithms) method as implemented in the software package *HyPhy*. White indicates higher frequencies of cluster assignments, whereas dark red indicates lower frequencies.

Figure S7: Plot showing community transitions in 3 pure subtypes (A1, B and C).

Figure S8: Distribution of DSBM-predicted recombination breakpoints in HIV-1/M genomes over time. This plot summarizes the distribution of genomes with respect to the number of breakpoints and year of sample collection with a kernel density, mapped to a heat colour palette with dark red for higher densities. In addition, the individual data points are mapped to the same plot region with added random noise with respect to both variables to reduce overlap.

Figure S9: Frequency distribution of breakpoints between non-recombinant and recombinant genomes as classified by SCUEAL.