

METHODOLOGY ARTICLE

Open Access

Inferring clonal evolution of tumors from single nucleotide somatic mutations

Wei Jiao^{1,2†}, Shankar Vembu^{3†}, Amit G Deshwar⁴, Lincoln Stein^{1,2} and Quaid Morris^{1,3,4,5,6*}

Abstract

Background: High-throughput sequencing allows the detection and quantification of frequencies of somatic single nucleotide variants (SNV) in heterogeneous tumor cell populations. In some cases, the evolutionary history and population frequency of the subclonal lineages of tumor cells present in the sample can be reconstructed from these SNV frequency measurements. But automated methods to do this reconstruction are not available and the conditions under which reconstruction is possible have not been described.

Results: We describe the conditions under which the evolutionary history can be uniquely reconstructed from SNV frequencies from single or multiple samples from the tumor population and we introduce a new statistical model, *PhyloSub*, that infers the phylogeny and genotype of the major subclonal lineages represented in the population of cancer cells. It uses a Bayesian nonparametric prior over trees that groups SNVs into major subclonal lineages and automatically estimates the number of lineages and their ancestry. We sample from the joint posterior distribution over trees to identify evolutionary histories and cell population frequencies that have the highest probability of generating the observed SNV frequency data. When multiple phylogenies are consistent with a given set of SNV frequencies, *PhyloSub* represents the uncertainty in the tumor phylogeny using a “partial order plot”. Experiments on a simulated dataset and two real datasets comprising tumor samples from acute myeloid leukemia and chronic lymphocytic leukemia patients demonstrate that *PhyloSub* can infer both linear (or chain) and branching lineages and its inferences are in good agreement with ground truth, where it is available.

Conclusions: *PhyloSub* can be applied to frequencies of any “binary” somatic mutation, including SNVs as well as small insertions and deletions. The *PhyloSub* and partial order plot software is available from <https://github.com/morrislab/phylosub/>.

Background

Cancer is a complex disease often associated with a characteristic series of somatic genetic variants [1,2]. Substantial effort has been devoted to genetic profiling of tumors in hopes of identifying these driver mutations and studying how they drive tumor development and resistance to treatment [3]. Tumors often contain multiple, genetically diverse subclonal populations of cells [4], and in some cases it is possible to reconstruct the evolutionary history of the tumor, thereby aiding in the identification of driver mutations, by computing the population

frequencies of mutations that distinguish the subclonal populations [5-13].

Somatic mutations can be detected, and roughly quantified, using exome and whole genome sequencing of a sample from a bulk tumor [14]. However, recent attempts to reconstruct subclonal phylogenies have employed much deeper targeted sequencing [15] of tumor-associated single nucleotide variants (SNVs) to achieve higher accuracy in estimated SNV frequencies [9,10,16,17]. These SNV frequencies were then used to partially reconstruct the evolutionary history of tumors based on a single [10,16] or multiple [9] samples of same tumor. However, due to short read sequencing, the frequencies of different SNVs are measured independently, so linkage between the SNVs in subclones is unavailable and standard phylogenetic methodology can not be used to construct evolutionary histories (as done in [18] or [17]). However, if one

*Correspondence: quaid.morris@utoronto.ca

†Equal contributors

¹Department of Molecular Genetics, University of Toronto, Toronto, Canada

³Donnelly Center for Cellular and Biomolecular Research, University of Toronto, Toronto, Canada

Full list of author information is available at the end of the article

makes the *infinite sites assumption* about tumor evolution, namely that every SNV only appeared once, then it is possible to use SNV frequencies to automatically reconstruct full or partial subclonal phylogenies while also inferring the multiple SNV genotypes of the major subclonal lineages in the tumor.

Here we describe a new method that automatically performs this phylogenetic reconstruction. First, we demonstrate that an unambiguous reconstruction is possible by describing *topological constraint rules* that are sufficient conditions to infer whether a triplet of SNV frequencies is consistent with only a chain or a branching phylogeny. We then describe a new method, *PhyloSub*, that automatically infers tumor phylogenies from SNV allele frequencies measured in single or multiple tumor samples. *PhyloSub* is based on a generative probabilistic model, inference in which implicitly implements the two rules by inferring the hidden phylogenies that have high probabilities of generating the observed SNV frequencies. It uses Bayesian inference, based on Markov Chain Monte Carlo (MCMC) sampling, to infer a distribution over phylogenies that incorporates uncertainty due to multiple phylogenies being consistent with the SNV frequencies and also noise in the measurement of the SNV frequencies. *PhyloSub* uses a Dirichlet process prior over phylogenies [19] to group SNVs into major subclonal lineages.

Model assumptions

We assume that the tumor evolution proceeds according to the clonal evolution theory, namely that all tumor cells are derived from ancestors that gain growth advantages over normal tissue and begins to expand [18]. Subsequent mutations can provide a further fitness or survival advantage to their subclonal lineage [20] which subsequently increases in frequency compared to cells containing only the SNVs in the parental lineage. A given tumor sample is a snapshot of this evolutionary process and may contain, at non-negligible frequency, cells from multiple major subclonal lineages, each containing a different assortment of these advantageous mutations. We make the *infinite sites assumption* [21,22], namely that each SNV appears only once and furthermore that once it appears, it does not revert back to its original state. As we describe below and illustrate in Figure 1, in some circumstances, this assumption highly constrains the phylogenies that are consistent with the SNV allele frequency data, especially if SNV frequencies from multiple samples from the same tumor are available. Finally, to make our model robust to low tumor cellularity, we assume that each tumor is derived from a single clone, however, this assumption is not critical in modeling tumor evolution and we revisit this assumption in the discussion section where we describe how to generalize our model to multicentric tumors (e.g., [23]).

To simplify our initial discussion, we will assume that the exact population frequencies of the cells containing each SNV (i.e., the SNV population frequency) are available before discussing how we estimate these frequencies from deep sequencing data of the SNV locus. Note, we assume that the copy number of a locus is available as per [10]. In the datasets that we considered, most SNVs are heterozygous at a normal copy number locus and the population frequency of other SNVs is easily inferred from their allele frequencies. In more complex situations, a number of tools are available to infer copy number changes associated with specific subclonal lineages from whole genome sequencing data [11,13].

An important consequence of the infinite sites assumption is that if SNV B occurred in a cell that contained SNV A, then all cells that have B also have A and thus the population frequency of A must always be greater than or equal to that of B, regardless of where and when the tumor sample was taken. However, a given set of three SNV population frequencies can still be consistent with two different phylogenies: a linear phylogeny or a branching phylogeny (see Figure 1A).

Topological constraint rules

One can distinguish linear or branching descent under some circumstances. For example, if we have already established that SNV A is ancestral to both B and C (i.e., that all cells with B or C also contain A), then if the population frequency of B plus the population frequency of C is greater than the population frequency of A, then the phylogeny must be linear. This is true because in a branching phylogeny, there are no cells that contain both B and C, so the population frequency of A must be at least as large as sum of the frequencies of B and C (see Figure 1B). We call this the “sum rule”. However, because a linear phylogeny is consistent with any set of SNV frequencies from a single sample, without making any further assumptions about the tumor evolution process, one needs at least two tumor samples to be able to rule out a linear phylogeny. However, given two samples and again assuming that SNV A is ancestral to both B and C, if the population frequency of B is larger than that of C in one sample, and vice versa in the other, then neither B nor C can be ancestral to the other, and the only phylogeny consistent with both sets of SNV frequencies is the branching one. We call this the “crossing rule” because the frequencies of B and C cross (see Figure 1C for an example). However, there is no guarantee that one can apply either rule to any set of SNV frequencies for all triplets of SNVs, although increasing the number of tumor samples does make it more likely that either the sum or crossing rule will be applicable for one or a pair of tumor samples, respectively. Furthermore, one needs to also consider the possibility of estimation error in the SNV population frequencies because these are

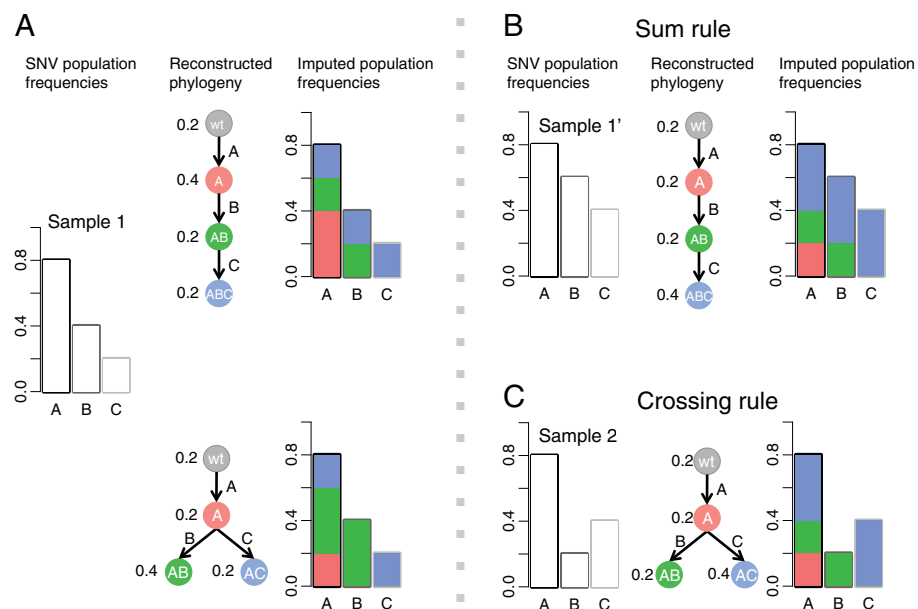


Figure 1 Visualization of topological constraint rules. (A) A, B, C are three SNVs, each of which represents a set of SNVs with similar SNV population frequencies. When the SNV population frequencies are 0.8, 0.4, 0.2 (left panel), there might be two possible phylogenies that are consistent with these frequencies (middle panel). The two solutions estimate same number of clonal populations but different genotypes for each clone. The decomposition of the clonal population frequencies are shown on the right panel. (B) Because of the sum rule, for this given set of SNV population frequencies 0.8, 0.6, 0.4, a chain structure may be the only possible phylogeny to explain the frequency changes. (C) Under the crossing rule, when multiple samples from the same patient are taken, we would expect the phylogenies are shared between samples. When another set of frequencies are observed 0.8, 0.2, 0.4, the branching structure is the only possible phylogeny to explain the frequencies changes for both sample 1 and 2.

inferred from discrete read counts. Note that these two rules also apply where SNV A is a mock SNV representing the wildtype state and having population frequency of 100%; as such these two rules also apply for multicentric tumors.

The PhyloSub algorithm

To explicitly model uncertainty in estimates of the SNV population frequencies and the precise tumor phylogeny, we have developed the PhyloSub model that we describe here. PhyloSub attempts to explain the observed read counts in terms of a latent phylogeny that associates SNVs with particular subclonal lineages. We provide software that takes as input a set of read counts for a set of SNVs and the copy number status of each SNV, performs inference in the PhyloSub model to estimate the number of major subclonal lineages, the mutational profile of each lineage, and the proportion of each lineage within the tumor cell population from which the read data was drawn. PhyloSub implements the parsimony assumptions detailed above using a non-parametric prior over tree structures. It is “generative” in that it attempts to explain the observed SNV frequencies in terms of an unobserved phylogeny; our model is also “Bayesian” in that it infers a posterior distribution over phylogenies and associated

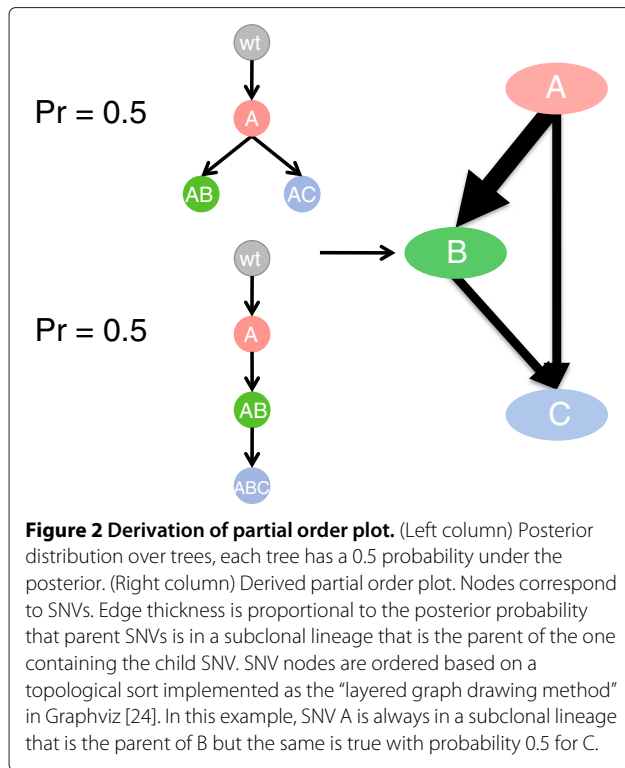
subclonal lineage frequencies. We introduce a new visualization, the *partial order plot*, to represent the posterior uncertainty in the phylogeny when the SNV frequencies alone do not provide sufficient information to uniquely reconstruct the phylogeny (Figure 2). The sum and crossing rule described above are implicitly incorporated into our generative model – our model assigns very low probability to any read counts that reflect deviations from either rule.

In the following, we provide a brief introduction to the PhyloSub model (see Section “Methods” for the full model) and we demonstrate its application to datasets where a single sample is profiled [17] and those where multiple samples are profiled [9]. We also report the application of the model on a simulated dataset to show that its prior parameterization allows it to represent a wide variety of phylogenies.

Results and discussion

PhyloSub

PhyloSub represents the major subclonal lineages and their evolutionary relationships using a directed tree in which each node is associated with a subclonal lineage and the edges connect parental lineages to their direct child lineage. Each subclonal lineage is associated with a



distinct subset of the SNVs input to the model, we call this subset the genotype of the lineage. Each node is also associated with (i) a set of SNVs that are present in this lineage but not its parent lineage and (ii) the population frequency of cells with the lineage genotype (and with no other SNVs from the input set). A subclonal lineage contains all of the SNVs associated with its parent, so its full genotype can be reconstructed by taking the union of the SNVs associated with its node and all of its ancestral nodes. Similarly, the population frequency of an SNV is the sum of the subclonal lineage frequencies of the lineage it appeared in and all of its descendent lineages. So, the subclonal lineage tree can be used to compute the population frequencies of each SNV and the genotype of each subclonal lineage. Associated with each SNV is a variable that indicates its zygosity and copy number in the cells that it appears (e.g., Aa indicating heterozygous and normal copy number), we assume all cells with the SNV have the same zygosity and copy number, and that all other cells have normal copy number at the SNV locus. The SNV genotype variable along with the population frequency is used to compute the allelic frequency of the SNV i , p_i . The data input to the model for each SNV is the number of reads mapping to the SNV locus, d_i , and the number of these reads that do not contain the SNV, a_i . We evaluate the likelihood of a given subclonal lineage tree (including the lineage population frequencies and the SNV genotype variables) by taking the product of the read count probabilities for each

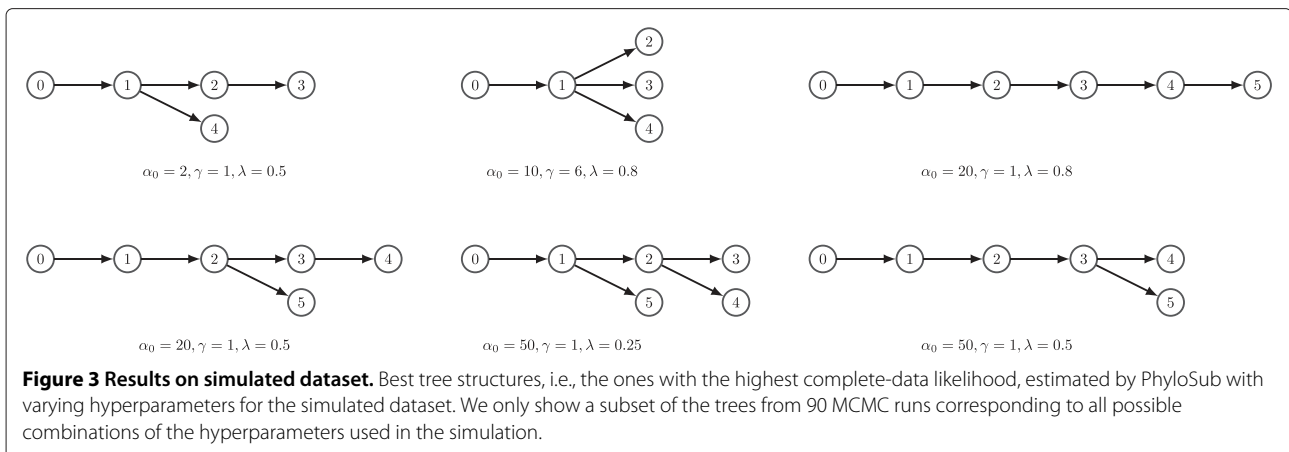
SNV, where the probability for the locus of SNV i is computed using a binomial distribution whose parameter is derived from p_i and an estimate of the error rate of the sequencer. PhyloSub also contains a vague prior over tree structures that is parameterized by three hyperparameters ($\alpha_0, \gamma, \lambda$) (see Section “Methods”) that govern how the prior scores trees with more or fewer nodes, and different average numbers of siblings. We use ranges for these hyperparameters that in simulations have a slight preference for trees with fewer nodes but a limited preference for sibling numbers (see below for details).

Simulations

PhyloSub’s Dirichlet process prior over tree structures depends on three hyperparameters: α_0 , γ , and λ . The hyperparameters α_0 and λ determine the number of nodes (subclones) in the tree, λ also affects the height of the tree and γ affects the number of siblings in the tree which in turn affects the width of the tree. In all the experiments, we sample these hyperparameters [19] as part of the MCMC sampling from a range whose upper and lower bounds we establish in this section.

To establish the ranges that we use for the hyperparameters in PhyloSub, we simulated read counts from clusters with an average of nine SNVs per cluster with SNV population frequencies {1.0, 0.85, 0.6, 0.35, 0.2, 0.08}, with a read depth of $\approx 10,000\times$ which is a typical read depth for the targeting deep sequencing data that PhyloSub is designed for. We simulated heterozygous SNVs at loci with normal copy number and sample read counts for each SNV from a Binomial distribution (see Section “Methods”). The hyperparameter settings we used in the simulations are all possible combinations of $\alpha_0 \in \{1, 2, 4, 10, 20, 50\}$, $\gamma \in \{1, 2, 4, 6, 8\}$ and $\lambda \in \{0.25, 0.5, 1\}$. The SNV population frequencies are consistent with many different tree structures and Figure 3 shows that the tree structures with highest complete-data likelihoods varies in the expected way for different settings of the tree prior hyperparameters. Although the preferred structure varies, the inferred SNV frequencies remain well-correlated with the baseline values (Pearson correlation > 0.99) for these hyperparameter ranges, so the prior is not over-regularizing the SNV frequencies for these settings. To allow a range of tree structures, we integrate over these ranges by placing a uniform prior on the choice of these settings in our MCMC simulations (c.f., [19]).

Although we focused on high read depths in the above simulation, we found that PhyloSub works well for read depths $\approx 1,000\times$ and was able to recover the clusters similar to the ones reported above and the SNV frequencies are well-correlated with the baseline values (Pearson correlation > 0.99). However, we found that the performance of the model degrades slightly at a read depth of $\approx 200\times$,



due to merging of clusters whose nearby SNV frequencies could not be distinguished. Nonetheless, we note that the inferred SNV population frequency estimates remain well-correlated (Pearson correlation > 0.96) and that the majority of the clusters were recovered at read depth $\approx 200\times$.

The simulation as described above has no clear phylogeny by design. The SNV frequencies were consistent with multiple phylogenies and the main goal of this simulation was to establish the ranges for our hyperparameters that permit a wide variety of tree structures. We integrate over these parameter ranges on the real data in order to remove any prior bias towards particular structures. To determine whether PhyloSub can correctly recover the phylogenies from a single sample of SNV frequencies, we simulated read data from a chain phylogeny with SNV population frequencies $0.9 \rightarrow 0.75 \rightarrow 0.55 \rightarrow 0.4 \rightarrow 0.25$. By the sum rule, these frequencies are only consistent with a chain phylogeny. PhyloSub was able to reliably recover this chain. The real datasets described in the later sections are representative of the types of problems that our methodology could be applied to as they contain single and multiple samples, some of which have clear phylogenies and some do not.

Results on AML datasets

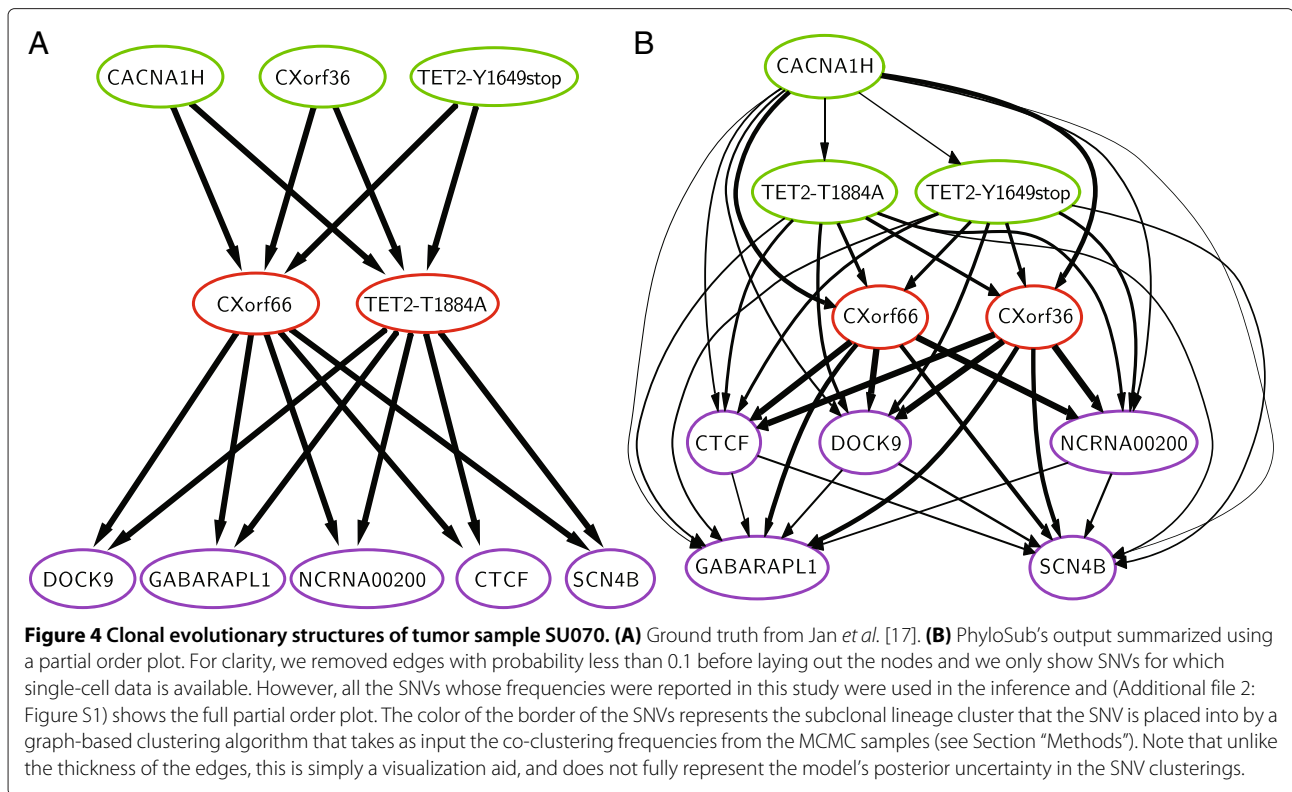
To assess PhyloSub on single samples, we applied it to data from Jan *et al.* [17] who reported the coexistence of multiple subclonal lineages in hematopoietic stem cells (HSC) from acute myeloid leukemia (AML) patient samples. The deep-targeted sequencing of all SNV candidates identified by exome sequencing identified SNVs with differing allelic frequency, suggesting multiple clonal populations in the HSC cells. An independent single-cell assay confirmed the existence of multiple clones, and thus provides a ground truth tree that shows some of the major subclonal lineages within the populations. Here we apply PhyloSub to the two samples profiled by Jan *et al.* that had three

or more SNVs profiled in a single-cell assay. These samples are SU048 and SU070 which have 6 and 10 SNVs in the single-cell assay, respectively. Although this assay confirmed the presence of some of the subclonal lineages, only 100-200 cells were assayed, so lineages with low population frequency in the sample (e.g., < 1%) may not be detected.

We applied PhyloSub providing it with the copy number and zygosity of each SNV (results were similar if we assume normal copy number and have a uniform prior on zygosity). For both SU048 and SU070, a number of different phylogenies were consistent with the SNV read counts, and we developed the “partial order plot” to represent the posterior uncertainty in the phylogeny (see Figure 2 and Section “Methods”).

Figure 4 shows that partial order plot for SU070. The ordering of the nodes in the partial order plot can also be used to infer ancestry via transitivity, for example, in Figure 4, the SNV CXorf66 has high probability of being in the subclonal lineage that is the direct parent of the one that DOCK9 is in, however, because the TET2-T1884A SNV is sorted before CXorf66 (and has a small probability of being a direct parent of it), then in the PhyloSub posterior over lineages, TET2-T1884A has a high probability of being in an ancestral lineage to the one CXorf66 is in.

Furthermore, one can interpret the partial order plot to indicate that both CXorf36 and CXorf66 are in the same lineage because they are both direct parents of DOCK9 (with high probability) and there are no edges between them. For reference, in Figure 4 we have included the results of the single-cell assay for SU070 in the partial order plot representation – Jan *et al.* report three subclonal lineages for SU070, as indicated by the SNV colorings [17]. We note that these plots are largely consistent. Indeed, we assign high posterior probability > 0.96, to two of the three subclonal lineages detected by Jan *et al.* (see Additional file 1: Table S3 for full lineage genotype probabilities). For reference, we also provide the list



of the subclonal lineage trees along with their posterior probabilities in see (Additional file 1: Table S1).

The one major difference between PhyloSub's estimates and the single-cell data from Jan *et al.* is that PhyloSub switches the order of the appearance of SNVs CXorf36 and TET2-T1884A. In fact, there was not a single subclonal lineage that contained CXorf36 but not TET2-T1884A in 5,000 subclonal lineage trees sampled from PhyloSub's posterior. This switch is likely due to the observed SNV frequencies, indeed the 95% confidence intervals of the SNV frequencies of these two SNVs do not overlap (Table 1). One explanation for this difference is a systematic bias in the measurement of one or both of these SNVs; it is also possible that the labels of these two SNVs were switched in Jan *et al.*. We also note, however, that in Jan *et al.*, the existence of the lineage that contains only CXorf36, TET2-Y1649stop, and CACNA1H is only supported by 2 of the 189 clones that they profiled.

For the tumor sample SU048, both the partial order plot and the single-cell assay agree on TET2-E1357stop event occurring early (at the root of the tree), and all other SNVs are secondary mutational events as shown in Figure 5B. Note that the partial order plot shows a large uncertainty in the structure for the rest of the SNVs and this is also reflected in the posterior over subclonal lineage trees and genotypes (see Additional file 1: Tables S2 and S4, respectively). There is no strong evidence for either a linear or

branching lineage or for particular clustering among these SNVs. Also, from Table 2, we see a lot of variation in the allele frequencies of these SNVs suggesting that they may not belong to the same subclonal lineages. The subclonal lineage inferred by Jan *et al.*'s single-cell assay is shown in Figure 5A and only contains two lineages, one with only TET2-E1357stop and the other with the other five SNVs. The TET2-E1357stop lineage genotype has probability 0.81 in our posterior, however the second genotype has a relatively small probability (0.06) under the posterior although we note that the genotype that contains all SNVs but ZMYM3 has a posterior probability of 0.32 (see Additional file 1: Table S4). For reference, we also provide the list of the subclonal lineage trees along with their posterior probabilities in (Additional file 1: Table S2).

In summary, the subclonal lineage trees inferred by PhyloSub on single samples of SNV frequencies are largely consistent with ground truth but there remains substantial uncertainty in SU048 about whether there was a linear or chain lineage. On the other hand, the SNV frequencies in SU070 were only consistent with a linear lineage and PhyloSub almost perfectly reconstructed the results of the single-cell assay with one misordering of the SNVs. This difference may be explained by unmodeled systematic biases in the deep sequencing data or experimental error. Nonetheless, we have shown that in some cases, it is possible to achieve a good estimate of the genotype

Table 1 Allelic counts for tumor sample SU070 from Jan *et al.* [17]

| SNV | Variant allele read counts | Read depth | Allele frequency | Cluster ID |
|----------------|----------------------------|------------|-----------------------------|------------|
| CACNA1H | 12,085 | 24,860 | 0.486 (95% CI: 0.481-0.491) | A |
| TET2-T1884A | 4,220 | 8,772 | 0.481 (95% CI: 0.472-0.490) | B |
| TET2-Y1649stop | 7,792 | 16,211 | 0.481 (95% CI: 0.474-0.487) | A |
| CXorf66 | 3,684 | 8,150 | 0.452 (95% CI: 0.443-0.461) | B |
| CXorf36 | 3,523 | 8,060 | 0.437 (95% CI: 0.428-0.446) | A |
| DOCK9 | 3,391 | 8,676 | 0.391 (95% CI: 0.382-0.400) | C |
| NCRNA00200 | 9,201 | 25,413 | 0.362 (95% CI: 0.357-0.367) | C |
| CTCF | 10,558 | 30,119 | 0.351 (95% CI: 0.346-0.355) | C |
| GABARAPL1 | 1,648 | 4,992 | 0.330 (95% CI: 0.319-0.341) | C |
| SCN4B | 5,113 | 16,386 | 0.312 (95% CI: 0.306-0.318) | C |

of multiple subclonal lineages as well as their evolution from a single, targeted deep sequencing sample of SNV frequencies.

Results on CLL datasets

To evaluate PhyloSub on a multiple sample dataset, we used data from a study of chronic lymphocytic leukemia (CLL) by Schuh *et al.* [9] which quantified SNV frequencies of a set of SNVs during different time points spanning the patient therapy cycle. The candidate SNVs were identified by exome sequencing and then subjected to targeted resequencing. The tumor samples from the three patients in the study labeled CLL077, CLL006 and CLL003 have 11, 16 and 20 SNVs respectively with SNV frequencies for five different time points. Originally, Schuh *et al.* reconstructed the evolutionary histories of each cancer by a semi-manual procedure in which they first automatically grouped SNVs into subclonal lineages using *k*-means

clustering on the allele frequencies and the differences in allele frequencies between the time points for each patient and then reconstructed the evolutionary structure of those lineages using a procedure that they do not describe in the paper. In PhyloSub, we model multiple samples from the same cancer as sharing the same evolutionary history but we allow subclonal frequencies to change between samples.

We applied PhyloSub to the SNV read count data, providing the algorithm with the likely zygosity estimates – in most cases, SNVs appeared to be heterozygous with normal copy number but in a few cases, SNVs appeared to be hemizygous and were input to the model as such. For these data, because of the multiple samples per tumor, there is very little posterior uncertainty in the best fitting tree – as such, we only show the best single tree structure corresponding to the MCMC sample with the highest complete-data likelihood [19].

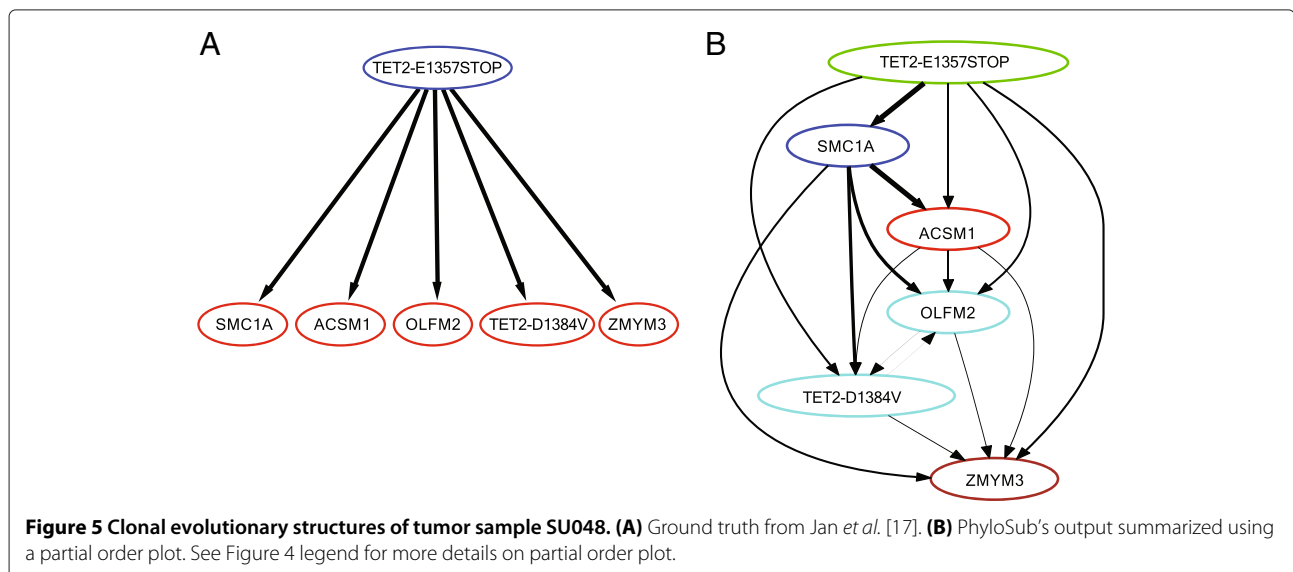


Table 2 Allelic counts for tumor sample SU048 from Jan *et al.* [17]

| SNV | Variant allele read counts | Read depth | Allele frequency | Cluster ID |
|----------------|----------------------------|------------|-----------------------------|------------|
| TET2-E1357stop | 7,436 | 19,553 | 0.380 (95% CI: 0.375-0.386) | A |
| SMC1A | 182,974 | 660,069 | 0.277 (95% CI: 0.276-0.278) | B |
| ACSM1 | 17,149 | 127,236 | 0.135 (95% CI: 0.133-0.136) | B |
| OLFM2 | 13,828 | 122,523 | 0.113 (95% CI: 0.111-0.114) | B |
| TET2-D1384V | 1,833 | 17,687 | 0.104 (95% CI: 0.100-0.107) | B |
| ZMYM3 | 18,536 | 307,346 | 0.060 (95% CI: 0.060-0.061) | B |

For the tumor samples CLL077 and CLL003, the best tree structure estimated by PhyloSub and the tree structure from Schuh *et al.* [9] are in exact agreement and the population frequencies of the subclonal lineages are well-correlated. Figures 6 and 7 compare the PhyloSub estimates with those reported by Schuh *et al.* [9].

For the tumor sample CLL006, PhyloSub inferred a chain structure similar to the chain structure from Schuh *et al.*, but the major difference in PhyloSub's best estimate of the tree structure is the splitting of cluster A into two clusters as shown in Figure 8. However, we found that the complete-data log likelihood of PhyloSub's best estimate of the tree structure is higher than the one for the chain structure of Schuh *et al.* and therefore PhyloSub prefers the splitting of the cluster A into two clusters.

In the CLL dataset, there is no ground truth but to allow the reader to compare the two estimates of the evolutionary history, Figure 9 plots the frequency of each SNV in the three samples, and we have colored SNVs according to

their subclonal lineage assignments by Schuh *et al.* These SNV frequencies are not corrected for copy number, however, the hemizygous SNVs are clear from examination of the figure.

In summary, having multiple samples of SNV frequencies greatly reduces the posterior uncertainty in the evolutionary history of the tumor and PhyloSub is able to reconstruct histories produced by a semi-manual procedure.

Conclusions

We presented a nonparametric Bayesian model called PhyloSub that uses a Dirichlet process prior over trees [19] to model the clonal evolutionary structure of tumors from next generation sequencing data. We also introduced a new visualization method, the partial order plot, to represent the posterior uncertainty in the phylogeny when the clonal frequencies alone do not provide sufficient information to uniquely reconstruct the phylogeny and mutational profiles of each subclonal lineage represented

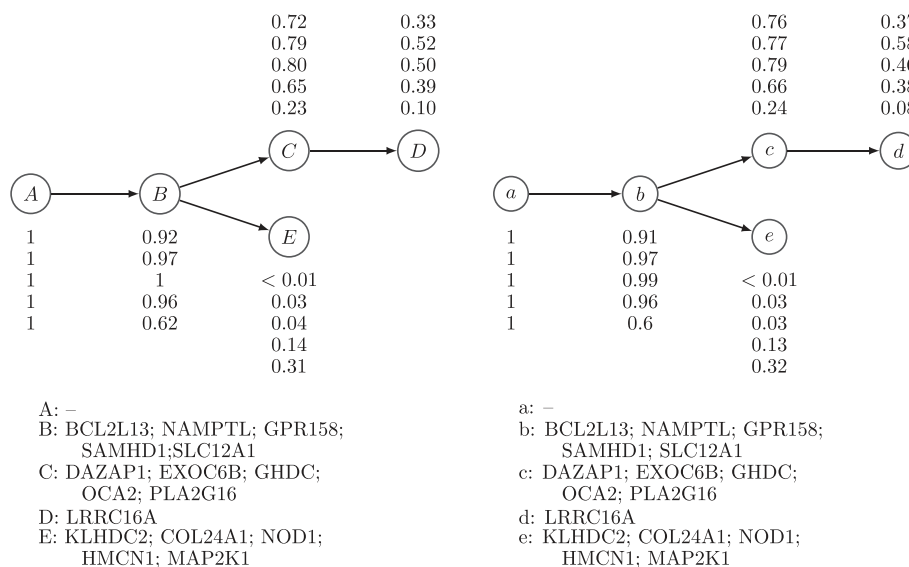


Figure 6 Clonal evolutionary structures of tumor samples from patient CLL077. (Left) Baseline tree structure from Schuh *et al.* [9]. **(Right)** Best tree structure estimated by PhyloSub. The SNV population frequencies and the cluster assignments are also shown in the figure.

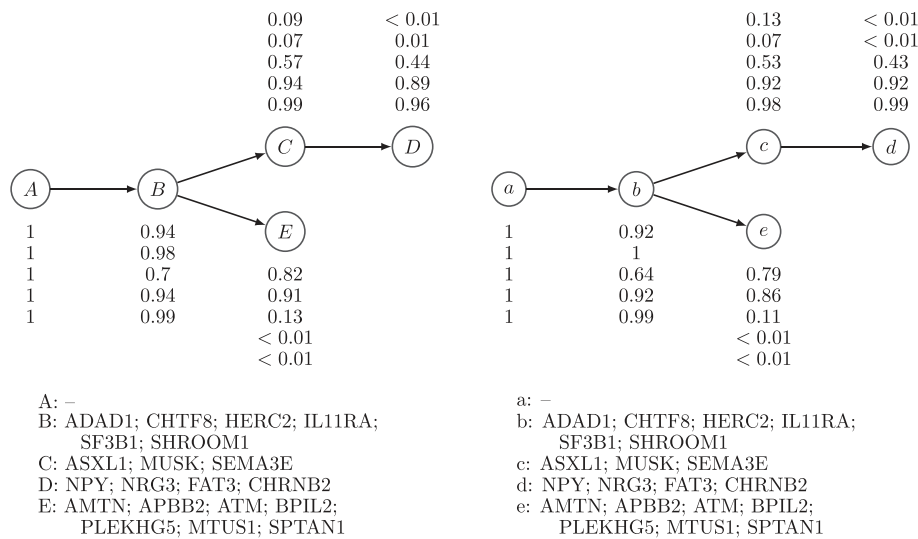


Figure 7 Clonal evolutionary structures of tumor samples from patient CLL003. (Left) Baseline tree structure from Schuh et al. [9]. **(Right)** Best tree structure estimated by PhyloSub. The SNV population frequencies and the cluster assignments are also shown in the figure.

in the tumor. By enforcing a set of structural constraints on the SNV population frequencies using MCMC methods, we were able to infer the phylogenetic relationships between subclones from both single and multiple tumor samples.

We have demonstrated that it is possible, in some cases, to detect a linear lineage from a single, high cellularity sample of the tumor. We have also shown that multiple samples highly constrain the possible lineages that are consistent with the SNV frequency data. PhyloSub's inferred subclonal lineage trees were in good agreement with single cell assays on single sample data and with an expert-driven, semi-manual reconstruction procedure on multiple sample data.

PhyloSub's ability to detect and characterize subclonal lineages depends on the frequency of the lineage in the population (compared to its descendant lineages), the number of SNVs that define the lineage, as well as the accuracy with which the SNV population frequencies are

estimated which depends on both the sequencing depth as well as uncertainty about the copy number of the SNV. Simply put, for lineages defined by a single SNV, the read depth has to be high enough that the uncertainty in the estimated SNV frequency is less than the frequency of the subclonal population. Having more lineage-defining SNVs can relax this hard constraint. As such, the phylogenies of tumors with large numbers of subclonal lineages, each defined by a small number of SNVs (possibly due to a pronounced hypermutability phenotype), will be hard to reconstruct with PhyloSub, or any other method, unless the SNV frequencies are very accurately estimated. Indeed, it is not clear how ground truth could be uncovered in such a case: the gold standard of single cell sequencing would require an exceptionally large number of single cells to survey this highly heterogeneous population, and each of these cells would need to be sequenced deeply in order to ensure precise somatic variant calling.

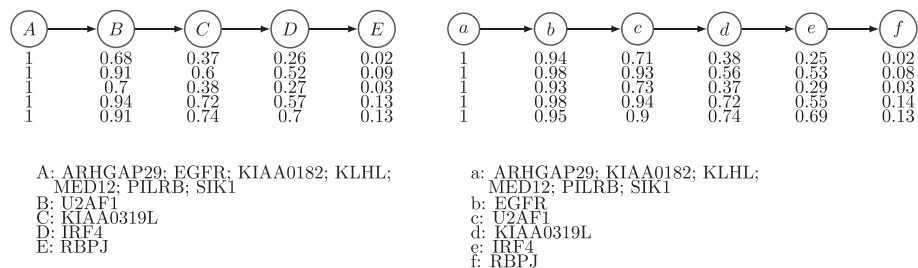
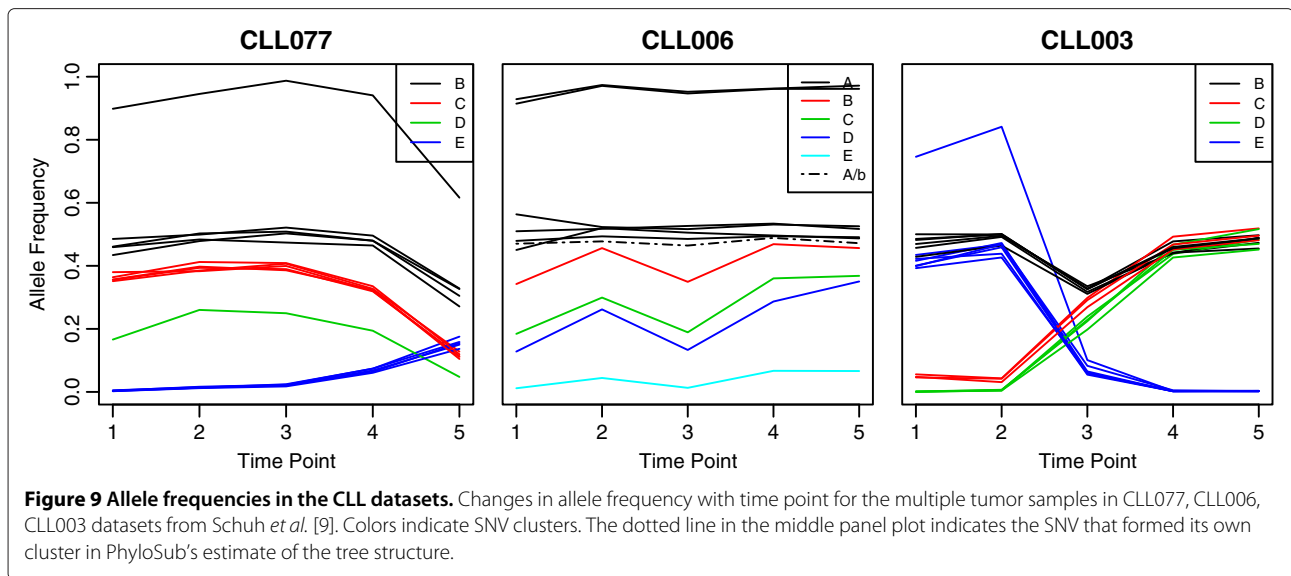


Figure 8 Clonal evolutionary structures of tumor samples from patient CLL006. (Left) Baseline tree structure from Schuh et al. [9]. **(Right)** Best tree structure estimated by PhyloSub. The SNV population frequencies and the cluster assignments are also shown in the figure.



One potential difficulty in scaling our approach to orders of magnitude more SNVs is that the Markov chain may not mix in a timely manner, in other words, may get stuck in local minima. We note that finding suboptimal solutions is an issue for any method based on these data. In our case, the mixing time of the chain would depend largely on the number of subclones represented in the population with less of a dependence on the number of SNVs. There are various techniques for determining whether or not a Markov chain is well-mixed and we refer the reader to a recent excellent review [25].

PhyloSub extends recent work on inferred cellularity and subclonal structure from somatic mutations. ABSOLUTE uses whole genome sequence data or array CGH data to identify regions of copy number change in the tumor and based on this infers cellularity and copy number changes associated with different subclones [11]. THetA [13] also attempts to infer both the copy number profiles and their relative proportions using the whole genome sequencing data based on an infinite sites assumptions. Neither of these algorithms explicitly reconstructs tumor phylogenies. Our work is closest to PyClone [10] which uses a flat Dirichlet process mixture model to group SNVs into subclonal lineages based on their frequencies; PhyloSub extends this work by reconstructing the phylogenetic relationships among these lineages and, in doing so, allows the full SNV genotype of each subclonal population to be reconstructed.

We designed PhyloSub to assume a single clonal origin for the cancerous cells in the sample. We made this decision to increase the applicability of the sum rule for low cellularity tumors (i.e., tumors with high normal contamination). However, removing this assumption would be a simple change to the model, which we have not evaluated.

Another area of future innovation would be in modeling sequencing biases and uncertainty in SNV allele frequencies resulting from them. We did not evaluate replacing our binomial model with a negative binomial one that would have allowed greater variability in the observed read counts for a given SNV allele frequency [26].

Methods

Dirichlet process mixture models

Consider the problem of clustering N objects $\{x_i\}_{i=1}^N$ using a Bayesian finite mixture model of K components (clusters) with the following generative process [27]:

$$\begin{aligned} \omega &\sim \text{Dirichlet}(\alpha/K, \dots, \alpha/K); & z_i &\sim \text{Multinomial}(\omega); \\ \phi_k &\sim H; & x_i &\sim F(\phi_{z_i}), \end{aligned} \tag{1}$$

where ω are the mixing weights such that $\sum_{k=1}^K \omega_k = 1$, α is the concentration parameter of the symmetric Dirichlet prior placed on the mixing weights, $z_i \in \{1, \dots, K\}$ is the cluster assignment variable, H is the prior distribution from which the component parameters $\{\phi_k\}$ are drawn, $F(\phi)$ is the component distribution parameterized by ϕ . The finite mixture model can be extended to a model with an infinite number of mixture components by replacing the Dirichlet prior with a Dirichlet process (DP) prior resulting in what is known as the DP mixture model (DPMM) [28].

Unlike finite mixture models, DPMMs automatically estimate the number of components from the data thereby circumventing the problem of fixing the number of components *a priori*. The stick-breaking construction [29]

given below provides a precise recipe to draw samples from a Dirichlet process:

$$\beta_k \sim \text{Beta}(1, \alpha); \quad \omega_1 = \beta_1; \quad \omega_k = \beta_k \prod_{\ell=1}^{k-1} (1 - \beta_\ell);$$

$$\phi_k \sim H; \quad \mathcal{G} = \sum_{k=1}^{\infty} \omega_k \delta_{\phi_k}, \tag{2}$$

where δ_ϕ is a point mass centered at ϕ and $\mathcal{G} \sim \text{DP}(\alpha, H)$, i.e., \mathcal{G} is a draw from a DP with base distribution H and concentration parameter α . The stick-breaking process can be viewed as recursively breaking sticks of length $\prod_{\ell=1}^{k-1} (1 - \beta_\ell)$, starting with a stick of unit length. The beta variates $\{\beta_k\}$ determine the random location at which the stick is broken. The concentration parameter α determines the number of clusters with high values resulting in large number of clusters. Let $\text{GEM}(\alpha)$ denote the stick-breaking process over ω . Replacing the Dirichlet prior in the finite mixture model 1 with the stick-breaking process prior results in the following generative process for infinite mixture models:

$$\omega \sim \text{GEM}(\alpha); \quad z_i \sim \text{Multinomial}(\omega);$$

$$\phi_k \sim H; \quad x_i \sim F(\phi_{z_i}).$$

An alternative view of the above generative process produces component parameters $\{\tilde{\phi}_i\}$ by drawing samples from \mathcal{G} resulting in the following generative process:

$$\mathcal{G} \sim \text{DP}(\alpha, H); \quad \tilde{\phi}_i \sim \mathcal{G}; \quad x_i \sim F(\tilde{\phi}_i). \tag{3}$$

Note that in the above process every object $\{x_i\}_{i=1}^N$ is associated with a component parameter $\{\tilde{\phi}_i\}_{i=1}^N$ and that all objects assigned to the same cluster will have the same component parameter. In other words, multiple elements in the set $\{\tilde{\phi}_i\}_{i=1}^N$ will take on the same value from the set $\{\phi_k\}_{k=1}^K$ of unique parameters.

Tree-structured stick-breaking process

The stick-breaking construction (2) described above can be used to produce a *flat* clustering of objects, where the clusters are independent of each other. Adams *et al.* [19] extended this construction for hierarchical clustering by interleaving two stick-breaking processes. This construction results in a *relational* clustering of objects where the clusters are connected to form a rooted tree structure. Unlike classical hierarchical clustering algorithms such as agglomerative clustering, this construction allows data to reside in the internal nodes of the tree; a fea-

ture we exploit to model the association of SNVs with subclonal lineages.

We borrow notation from Adams *et al.* [19]. Let $\epsilon = (\epsilon_1, \dots, \epsilon_p)$ denote a sequence of positive integers used to index the nodes of the tree. Let $\epsilon = \kappa$ denote the zero-length string, i.e., the root of the tree. Let $|\epsilon|$ indicate the length of the sequence ϵ and therefore the depth of node ϵ . Let $\epsilon \epsilon_i$ denote the sequence formed by appending ϵ_i to ϵ . The children of node ϵ is the set $\{\epsilon \epsilon_i : \epsilon_i \in 1, 2, \dots\}$ and let the ancestors of ϵ be denoted by the set $\{\epsilon' : \epsilon' < \epsilon\}$. The interleaved, two-layered stick-breaking construction is as follows:

$$v_\epsilon \sim \text{Beta}(1, \alpha(|\epsilon|)); \quad \psi_\epsilon \sim \text{Beta}(1, \gamma); \quad \omega_\kappa = v_\kappa;$$

$$\omega_\epsilon = v_\epsilon \varphi_\epsilon \prod_{\epsilon' < \epsilon} \varphi_{\epsilon'} (1 - v_{\epsilon'}); \quad \varphi_{\epsilon \epsilon_i} = \psi_{\epsilon \epsilon_i} \prod_{j=1}^{\epsilon_i - 1} (1 - \psi_{\epsilon_j}). \tag{4}$$

The v_ϵ and $(1 - v_\epsilon)$ determine the amount of mass allocated to ϵ and its descendants respectively, whereas $\{\varphi_\epsilon\}$ determines the probability of a particular sequence of children. The construction ensures that the mixing weights $\{\omega_\epsilon\}$ sum to one. The parameters α and γ control the height and the width of the tree respectively. Note that the concentration parameter $\alpha(\cdot)$ is a function of the depth of the tree ($\alpha(\cdot) : \mathbb{N} \rightarrow \mathbb{R}^+$) and is defined to be $\alpha(j) = \lambda^j \alpha_0$ with $\alpha_0 > 0$ and $\lambda \in (0, 1]$ [19].

PhyloSub model

We follow Shah *et al.* [10,30] to model the allelic count data. For each genetic variant that is detected by high-throughput sequencing methods, cells containing the genetic variant are referred to as variant population and those without the variant as reference population. Let $\Sigma = \{A, C, G, T\}$ denote the set of nucleotides. Let a_i and b_i denote the number of reads matching the reference allele $A \in \Sigma$ and the variant allele $B \in \Sigma$ respectively at position i , and let $d_i = a_i + b_i$. The genotype $g \in \{A, B, AA, AB, BB, AAA, \dots\}$ would depend on the copy number at the variant location. Let $\mu_i^r \in \{\mu_i^A, \mu_i^{AA}, \mu_i^{AAA}, \dots\}$ denote the probability of sampling a reference allele from the reference population. This value depends on the error rate of the sequencer. Let μ_i^v denote a vector whose entries, $\mu_i^{v:g} \in \{\mu_i^B, \mu_i^{AB}, \mu_i^{BB}, \dots\}$, are the probabilities of sampling a reference allele from the variant population with genotype g at position i . Let π_i denote the vector whose entries, $\pi_i^g \in \{\pi_i^B, \pi_i^{AB}, \pi_i^{BB}, \dots\}$, are the probabilities of the variant population at position i to have the genotype g . Let δ_i denote the pseudo-count parameters of the Dirichlet prior over π_i . Let G_i denote the genotype of the variant population at position i . Let $\tilde{\phi}_i$ denote the fraction of cells from the variant

population, i.e., the SNV population frequency at position i , and $1 - \tilde{\phi}_i$ denote the fraction of cells from the reference population at position i . The observation model for allelic counts has the following generative process [10]:

$$\begin{aligned} \mathcal{G} &\sim \text{DP}(\alpha, H); \quad \tilde{\phi}_i \sim \mathcal{G}; \\ \boldsymbol{\pi}_i &\sim \text{Dirichlet}(\boldsymbol{\delta}_i); \quad G_i \sim \text{Categorical}(\boldsymbol{\pi}_i); \\ a_i | d_i, G_i = g, \tilde{\phi}_i, \mu_i^r, \mu_i^v &\sim \text{Binomial}(d_i, (1 - \tilde{\phi}_i)\mu_i^r + \tilde{\phi}_i\mu_i^{v:g}). \end{aligned} \quad (5)$$

The posterior distribution of $\tilde{\phi}_i$ is

$$\begin{aligned} p(\tilde{\phi}_i | a_i, d_i, \mu_i^r, \mu_i^v, \boldsymbol{\pi}_i, \boldsymbol{\delta}_i) &\propto \sum_g p(a_i | d_i, G_i = g, \tilde{\phi}_i, \mu_i^r, \mu_i^v) \\ &\quad \times p(G_i | \boldsymbol{\pi}_i) p(\boldsymbol{\pi}_i | \boldsymbol{\delta}_i) p(\tilde{\phi}_i). \end{aligned}$$

Each of the terms appearing inside the summation over genotypes is the probability distribution of a Dirichlet compound multinomial (with a single draw) [31]. The posterior distribution can thus be rewritten as

$$\begin{aligned} p(\tilde{\phi}_i | \cdot) &\propto \sum_g \left[\frac{\prod_{g' \neq g} \Gamma(\delta_i^{g'}) \times \Gamma(\delta_i^g + 1)}{\Gamma(\sum_{g'} \delta_i^{g'} + 1)} \right] \\ &\quad \times \text{Binomial}(a_i; d_i, (1 - \tilde{\phi}_i)\mu_i^r + \tilde{\phi}_i\mu_i^{v:g}) p(\tilde{\phi}_i), \end{aligned} \quad (6)$$

where $\Gamma(\cdot)$ is the Gamma function.

The Dirichlet process prior $\text{DP}(\alpha, H)$ in the observation model of allelic counts (5) is useful to infer groups of mutations that occur at the same SNV population frequency [10]. Furthermore, being a nonparametric prior, it is useful to avoid the problem of selecting the number of groups of mutations *a priori*. However, it cannot be used to model the clonal evolutionary structure which takes the form of a rooted tree. In order to model this, we propose to use the tree-structured stick-breaking process prior (4) described in the previous section.

The probabilistic graphical model for allelic counts with the tree-structured stick-breaking process prior is shown in Figure 10. Inputs to the model including the hyperparameters are indicated in shaded nodes, whereas the latent variables including the set of SNV population frequencies $\{\tilde{\phi}_i\}$ are indicated in unshaded nodes. The prior/base distribution H of the SNV population frequencies is the uniform distribution $\text{Uniform}(0, 1)$ for the root node and $\text{Uniform}(0, \phi_{\text{par}(v)} - \sum_{w \in \mathcal{S}(v)} \phi_w)$ for any other node v in the tree, where $\text{par}(v)$ denotes the parent node of v and $\mathcal{S}(v)$ is the set of siblings of v . This ensures that

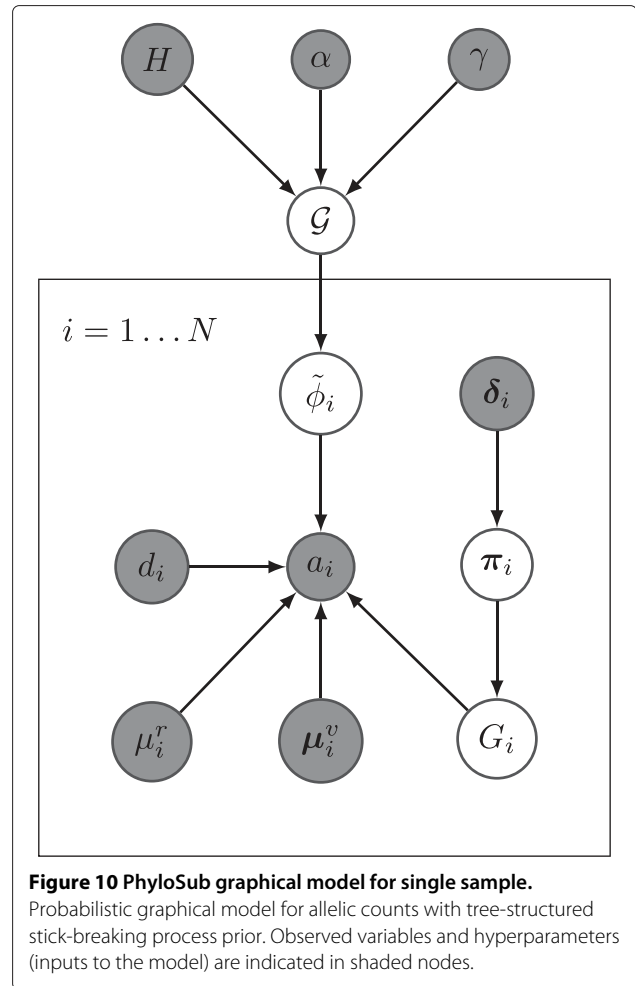


Figure 10 PhyloSub graphical model for single sample. Probabilistic graphical model for allelic counts with tree-structured stick-breaking process prior. Observed variables and hyperparameters (inputs to the model) are indicated in shaded nodes.

the clonal evolutionary constraints (discussed in the next section) are satisfied when adding a new node in the tree. The crucial difference between our model and the model of Shah *et al.* [10] is that we use a tree-structured stick-breaking process instead of a Dirichlet process (cf. 5) to generate the set of SNV population frequencies $\{\tilde{\phi}_i\}$. Given this model and a set of N observations/inputs $\{(a_i, d_i, \mu_i^r, \mu_i^v, \delta_i)\}_{i=1}^N$, the tree structure and the SNV population frequencies $\{\tilde{\phi}_i\}$ are inferred using Markov chain Monte Carlo sampling. In particular, we use Gibbs sampling [19] to generate posterior samples of the SNV population frequencies 6. Each iteration of the Gibbs sampler involves multiple subsampling procedures: sampling cluster assignments $\{z_i\}$, sampling stick lengths ν_ϵ and $\psi_{\epsilon \in \epsilon_i}$, sampling stick-breaking hyperparameters α_0, γ and λ , and sampling the SNV population frequencies $\{\tilde{\phi}_i\}$. Our main algorithmic contribution, described below, is a method to sample SNV population frequencies in such a way that the tumor evolution proceeds according to the assumptions from the clonal evolutionary theory. The rest of the

subsampling procedures follow directly from Adams *et al.* [19] and we refer the reader to it for further technical details.

Sampling SNV population frequencies

Given the current state of the tree structure, we sample SNV population frequencies in such a way that the SNV population frequency ϕ_v of every non-leaf node v in the tree is greater than or equal to the sum of the SNV population frequencies of its children. To enforce this constraint, we introduce a set of auxiliary weights $\{\eta_v\}$, one for each node, that satisfy $\sum_v \eta_v = 1$, and rewrite the observation model for allelic counts 5 explicitly in terms of these weights resulting in the following posterior distribution:

$$p(\tilde{\eta}_i | a_i, d_i, \mu_i^r, \mu_i^v, \pi_i, \delta_i) \propto \sum_g p(a_i | d_i, G_i = g, \tilde{\eta}_i, \mu_i^r, \mu_i^v) \times p(G_i | \pi_i) p(\pi_i | \delta_i) p(\tilde{\eta}_i), \quad (7)$$

where we have used $\{\tilde{\eta}_i\}$ to denote the auxiliary weights for each SNV. The prior/base distribution of the auxiliary weights is defined such that it is 1 for the singleton root node and $\text{Uniform}(0, \eta_{\text{par}(v)})$ for any other node v in the tree, where $\text{par}(v)$ denotes the parent node of v . When a new node w is added to the tree, we sample $\eta_w \sim \text{Uniform}(0, \eta_{\text{par}(w)})$ and update $\eta_{\text{par}(w)} \leftarrow \eta_{\text{par}(w)} - \eta_w$. This ensures that $\sum_v \eta_v = 1$.

This change is crucial as it allows us to design a Markov chain that converges to the stationary distribution of $\{\eta_v\}$. The SNV population frequency for any node v can then be computed via

$$\phi_v = \eta_v + \sum_{w \in \mathcal{D}(v)} \eta_w = \eta_v + \sum_{w \in \mathcal{C}(v)} \phi_w, \quad (8)$$

where $\mathcal{D}(v)$ and $\mathcal{C}(v)$ are the sets of all descendants and children of node v respectively. This construction ensures that the SNV population frequencies of mutations appearing at the parent node is greater than or equal to the sum of the frequencies of all its children. The procedure to generate a random sample of SNV population frequencies is given in Algorithm 1 where we generate (η_v, ϕ_v) for every node v by traversing the tree in a breadth-first fashion. The input to this algorithm is the current state of the tree $\mathcal{T} = (V, E)$ where V is the set of vertices and E is the set of edges, and the output is a multi-dimensional sample of SNV population frequencies $\phi = \{\phi_1, \phi_2, \dots, \phi_{|V|}\}$ (where $|V| = K$) and the corresponding auxiliary weights $\eta = \{\eta_1, \eta_2, \dots, \eta_{|V|}\}$. A sample from this algorithm is shown in Figure 11.

Algorithm 1 Algorithm to generate SNV population frequencies satisfying the assumptions from clonal evolutionary theory.

Input: Rooted tree $\mathcal{T} = (V, E)$ with root node r

Output: $\eta = \{\eta_1, \eta_2, \dots, \eta_{|V|}\}$, $\phi = \{\phi_1, \phi_2, \dots, \phi_{|V|}\}$

```

1: create a queue Q
2: Q.enqueue(r)
3: while Q is not empty do
4:   v = Q.dequeue()
5:   if v is root then
6:      $\phi_v = 1$ 
7:      $s_v \sim \text{Uniform}(0, 1)$ 
8:      $\eta_v = \phi_v \cdot s_v$ 
9:   end if
10:   $m_v = \phi_v - \eta_v$  {mass assigned to children of v}
11:  for c in children of v do
12:     $r_c \sim \text{Uniform}(0, 1)$  {distribute mass}
13:  end for
14:   $r = m_v \cdot r / \sum_c r_c$ 
15:  for c in children of v do
16:     $\phi_c = r_c$ 
17:     $s_c \sim \text{Uniform}(0, 1)$ 
18:     $\eta_c = \phi_c \cdot s_c$ 
19:    Q.enqueue(c)
20:  end for
21: end while
22: return  $\eta, \phi$ 
    
```

We use Metropolis-Hastings algorithm to sample from the posterior distribution of the auxiliary weights $\{\tilde{\eta}_i\}$ as shown in Algorithm 2 and derive the SNV population frequencies from these samples. We use an asymmetric Dirichlet distribution as the proposal distribution. This ensures that the Markov chain converges to the stationary distribution of $\{\tilde{\eta}_i\}$. The inputs to the sampling algorithm are the current state of the tree $\mathcal{T} = (V, E)$, a scaling factor

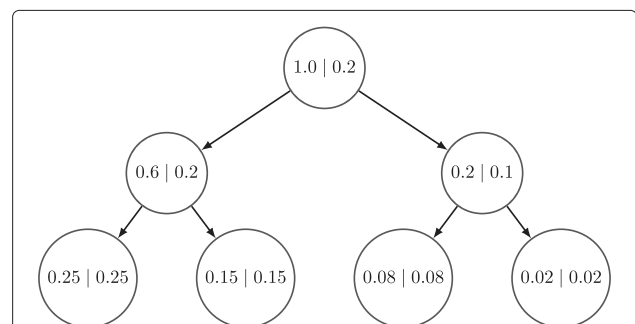


Figure 11 Example of SNV population frequencies generated using Algorithm 1. The labels of the nodes are its corresponding SNV population frequencies and auxiliary weights ($\phi | \eta$). Note that $\sum_{v \in V} \eta_v = 1$ and $\phi_v \geq \sum_{w \in \mathcal{C}(v)} \phi_w$ for every non-leaf node v .

σ , and the number of iterations T . The output is a sample from the posterior distribution of $\eta = \{\eta_1, \eta_2, \dots, \eta_{|V|}\}$ and its corresponding $\phi = \{\phi_1, \phi_2, \dots, \phi_{|V|}\}$.

Algorithm 2 Metropolis-Hastings algorithm to sample from the posterior distribution of the auxiliary weights $\{\eta_v\}$ and compute the SNV population frequencies $\{\phi_v\}$.

Input: Rooted tree $\mathcal{T} = (V, E)$, σ , T

Output: $\eta = \{\eta_1, \eta_2, \dots, \eta_{|V|}\}$, $\phi = \{\phi_1, \phi_2, \dots, \phi_{|V|}\}$

- 1: Initialize $\eta^{(0)}$ using Algorithm 1
- 2: **for** $t = 1 : T$ **do**
- 3: //draw a proposal state from the Dirichlet distribution with density function $Q(\cdot)$
- 4: $\eta' \sim \text{Dirichlet}(\sigma \eta^{(t-1)} + 1)$
- 5: // accept/reject state
- 6: $a = \log p(\eta' | \cdot) - \log p(\eta^{(t-1)} | \cdot) + \log Q(\eta^{(t-1)}; \eta', \sigma) - \log Q(\eta'; \eta^{(t-1)}, \sigma)$
- 7: $r \sim \text{Uniform}(0, 1)$
- 8: **if** $\log(r) < a$ **then**
- 9: $\eta^{(t)} \leftarrow \eta'$
- 10: **else**
- 11: $\eta^{(t)} \leftarrow \eta^{(t-1)}$
- 12: **end if**
- 13: **end for**
- 14: Compute ϕ from η
- 15: **return** $\eta^{(T)}$, $\phi^{(T)}$

Extension to multiple tumor samples

PhyloSub (cf. Figure 10) can be easily extended for multiple tumor samples. We allow the tree-structured stick-breaking process prior 4 to be shared across all the samples. Let a_i^t and b_i^t denote the number of reads matching the reference and the variant allele respectively at position i for sample $t \in \{1, \dots, S\}$, and let $d_i^t = a_i^t + b_i^t$. Let $\tilde{\phi}_i^t$ denote the fraction of cells from the variant population, i.e., the SNV population frequency at position i for sample t , and $\tilde{\eta}_i^t$ denote its corresponding auxiliary weight. The graphical model of PhyloSub for multiple tumor samples is shown in Figure 12. The main technical difference between the single and the multiple sample models lies in the sampling procedure for SNV population frequencies. In the multiple sample model, we ensure that the clonal evolutionary constraints described in the previous section are satisfied separately for each tumor sample and then make a global Metropolis-Hastings move based on the distribution $\prod_{t=1}^S p(\eta^t | \cdot)$, where $\{\eta^1, \eta^2, \dots, \eta^S\}$ is the set of auxiliary weights for all the tumor samples.

Partial order plot

We construct a partial order plot to summarize and visualize the trees from all the posterior MCMC samples. It is

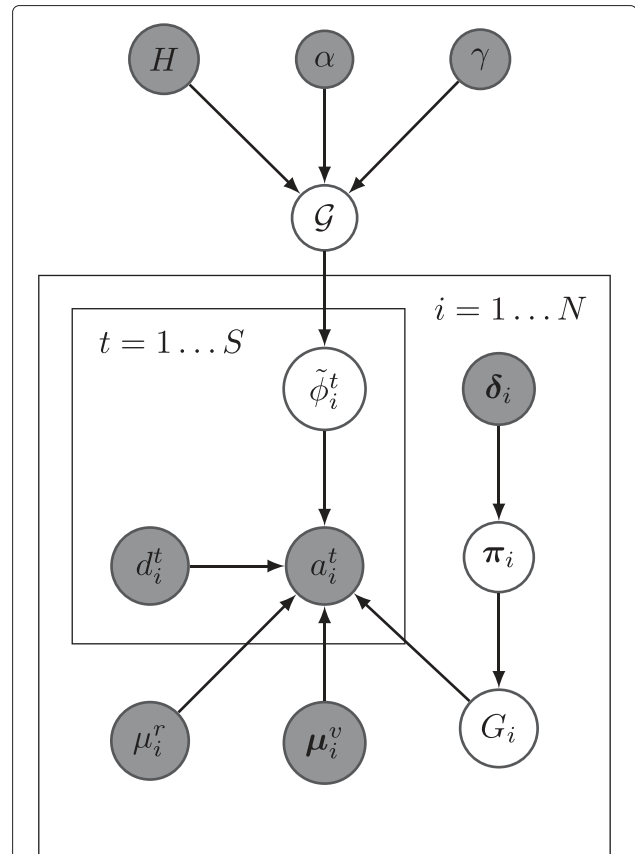


Figure 12 PhyloSub graphical model for multiple samples.

Probabilistic graphical model for allelic counts from multiple samples with a shared tree-structured stick-breaking process prior. Observed variables and hyperparameters (inputs to the model) are indicated in shaded nodes.

important to note that the nodes of this partial order plot are the SNVs and not the SNV clusters. The thickness of a directed edge $P \rightarrow Q$ in the tree is proportional to the fraction of MCMC samples in which SNV P first appears in a subclonal lineage that was the parent of the subclonal lineage that Q first appears in. The color of the border of the SNVs represents the subclonal lineage cluster that the SNV is placed into post hoc using an algorithm called correlation clustering [32]. Note that the main purpose of this clustering algorithm is only to color the nodes in the partial order plot by aggregating the clustering information from all the MCMC samples obtained from our model; this clustering is a summary but does not represent any (possibly quite large) uncertainty in the cluster assignments. The input to this algorithm is a symmetric $N \times N$ co-clustering matrix C , whose elements C_{ij} is the difference between the number of samples in which i and j were assigned to the same SNV cluster and the number of samples in which i and j were assigned to different SNV clusters. The algorithm estimates a symmetric $N \times N$

cluster indicator matrix Y , whose elements $Y_{ij} = 1$ if i and j are assigned to the same cluster and $Y_{ij} = 0$ otherwise. This cluster indicator matrix Y has all the information about the number of SNV clusters as well as the SNVs assigned to each of them.

MCMC settings

In all the experiments, we fix the number of MCMC iterations to 5,000 with a burn-in of 100 samples. We also fix the number of iterations in the Metropolis-Hastings algorithm to 5,000 and set the scaling factor for the Dirichlet proposal distribution to $\sigma = 100$. We run the MCMC samplers multiple times with different random initializations and pick a single run based on the complete-data likelihood trace and its auto-correlation function. We use all the 5,000 samples without thinning [33] to construct the partial order plots. We use the CODA R package [34] for MCMC diagnostics to monitor the convergence of the samplers. The complete-data log likelihood traces and the corresponding autocorrelation function plots after the burn-in period of 100 samples for all the experiments on AML and CLL datasets are shown in (Additional file 2: Figures S3 to S7).

Datasets and inputs to PhyloSub

All datasets used in the experiments, including details about the inputs to PhyloSub, i.e., the set of observations $\{(a_i, d_i, \mu_i^r, \mu_i^v, \delta_i)\}_{i=1}^N$, are provided in the (Additional file 1: Tables S5 – S10).

Additional files

Additional file 1: Supplementary tables. This file contains supplementary Tables S1 to S10.

Additional file 2: Supplementary figures. This file contains supplementary Figures S1 to S7.

Competing interests

The authors declare that they have no competing interests.

Authors' contributions

QM and LS conceived of the project. WJ collected and analyzed the data, performed all the experiments. SV, QM and AGD designed the PhyloSub model. SV implemented the machine learning algorithms and models, and wrote the Section "Methods" of the paper. QM and LS reviewed and edited the paper. All authors read and approved the final manuscript.

Acknowledgements

This work was funded by a National Science and Engineering Research Council (NSERC) operating grant and a Early Researcher Award from the Ontario Research Fund to QM. WJ and LS were supported by The Ministry of Research and Innovation, Province of Ontario.

We thank Andrew Roth and Sohrab Shah for sharing a pre-publication version of the PyClone software with us; helping us to install it and to duplicate their published results; and for providing unpublished details of the PyClone observation model in their user manual.

Author details

¹Department of Molecular Genetics, University of Toronto, Toronto, Canada.

²Ontario Institute for Cancer Research, Toronto, Canada. ³Donnelly Center for

Cellular and Biomolecular Research, University of Toronto, Toronto, Canada.

⁴Edward S. Rogers Sr. Department of Electrical and Computer Engineering, University of Toronto, Toronto, Canada. ⁵Banting and Best Department of Medical Research, University of Toronto, Toronto, Canada. ⁶Department of Computer Science, University of Toronto, Toronto, Canada.

Received: 21 May 2013 Accepted: 24 January 2014

Published: 1 February 2014

References

1. Hanahan D, Weinberg RA: **The hallmarks of cancer.** *Cell* 2000, **100**:57–70.
2. Hanahan D, Weinberg RA: **Hallmarks of cancer: The next generation.** *Cell* 2011, **144**(5):646–674.
3. Yap TA, Gerlinger M, Futreal PA, Pusztai L, Swanton C: **Intratumor heterogeneity: Seeing the wood for the trees.** *Sci Transl Med* 2012, **4**(127):127ps10.
4. Visvader JE: **Cells of origin in cancer.** *Nature* 2011, **469**(7330):314–322.
5. Navin NE, Hicks J: **Tracing the tumor lineage.** *Mol Oncol* 2010, **4**(3):267–283.
6. Mullighan CG, Phillips LA, Su X, Ma J, Miller CB, Shurtleff SA, Downing JR: **Genomic analysis of the clonal origins of relapsed acute lymphoblastic leukemia.** *Science* 2008, **322**(5906):1377–1380.
7. Gerlinger M, Rowan AJ, Horswell S, Larkin J, Endesfelder D, Gronroos E, Martinez P, Matthews N, Stewart A, Tarpey P, Varela I, Phillimore B, Begum S, McDonald NQ, Butler A, Jones D, Raine K, Latimer C, Santos CR, Nohadani M, Eklund AC, Spencer-Dene B, Clark G, Pickering L, Stamp G, Gore M, Szallasi Z, Downward J, Futreal PA, Swanton C: **Intratumor heterogeneity and branched evolution revealed by multiregion sequencing.** *N Engl J Med* 2012, **366**(10):883–892.
8. Marusyk A, Polyak K: **Tumor heterogeneity: Causes and consequences.** *Biochim Biophys Acta* 2010, **1805**:105–117.
9. Schuh A, Becq J, Humphray S, Alexa A, Burns A, Clifford R, Feller SM, Grocock R, Henderson S, Khrebtkova I, Kingsbury Z, Luo S, McBride D, Murray L, Menju T, Timbs A, Ross M, Taylor J, Bentley D: **Monitoring chronic lymphocytic leukemia progression by whole genome sequencing reveals heterogeneous clonal evolution patterns.** *Blood* 2012, **120**(20):4191–4196.
10. Shah SP, Roth A, Goya R, Oloumi A, Ha G, Zhao Y, Turashvili G, Ding J, Tse K, Haffari G, Bashashati A, Prentice LM, Khattra J, Burleigh A, Yap D, Bernard V, McPherson A, Shumansky K, Crisan A, Giuliany R, Heravi-Moussavi A, Rosner J, Lai D, Birol I, Varhol R, Tam A, Dhalla N, Zeng T, Ma K, Chan SK, et al.: **The clonal and mutational evolution spectrum of primary triple-negative breast cancers.** *Nature* 2012, **486**(7403):617–656.
11. Carter SL, Cibulskis K, Heltman E, McKenna A, Shen H, Zack T, Laird PW, Onofrio RC, Winckler W, Weir BA, Beroukhir R, Pellman D, Levine DA, Lander ES, Meyerson M, Getz G: **Absolute quantification of somatic DNA alterations in human cancer.** *Nat Biotechnol* 2012, **30**(5):413–421.
12. Landau DA, Carter SL, Stojanov P, McKenna A, Stevenson K, Lawrence MS, Sougnez C, Stewart C, Sivachenko A, Wang L, Wan Y, Zhang W, Shukla SA, Vartanov A, Fernandes SM, Saksena G, Cibulskis K, Tesar B, Gabriel S, Hacohen N, Meyerson M, Lander ES, Neuberger D, Brown JR, Getz G, Wu CJ: **Evolution and impact of subclonal mutations in chronic lymphocytic leukemia.** *Cell* 2013, **152**(4):714–726.
13. Oesper L, Mahmoody A, Raphael BJ: **THetA: Inferring intra-tumor heterogeneity from high-throughput DNA sequencing data.** *Genome Biol* 2013, **14**(7):R80.
14. Marusyk A, Almendro V, Polyak K: **Intra-tumour heterogeneity: A looking glass for cancer?** *Nat Rev Cancer* 2012, **12**(5):323–334.
15. Meyerson M, Gabriel S, Getz G: **Advances in understanding cancer genomes through second-generation sequencing.** *Nat Rev Genet* 2010, **11**(10):685–696.
16. Nik-Zainal S, Loo PV, Wedge DC, Alexandrov LB, Greenman CD, Lau KW, Raine K, Jones D, Marshall J, Ramakrishna M, Shlien A, Cooke SL, Hinton J, Menzies R, Stebbings LA, Leroy C, Jia M, Rance R, Mudie LJ, Gamble SJ, Stephens PJ, McLaren S, Tarpey PS, Papaemmanuil E, Davies HR, Varela I, McBride DJ, Bignell GR, Leung K, Butler AP, et al.: **The life history of 21 breast cancers.** *Cell* 2012, **149**:994–1007.
17. Jan M, Snyder TM, Corces-Zimmerman MR, Vyas P, Weissman IL, Quake SR, Majeti R: **Clonal evolution of preleukemic hematopoietic stem**

- cells precedes human acute myeloid leukemia. *Sci Transl Med* 2012, **4**(149):149ra118.
18. Campbell PJ, Pleasance ED, Stephens PJ, Dicks E, Rance R, Goodhead I, Follows GA, Green AR, Futreal PA, Stratton MR: **Subclonal phylogenetic structures in cancer revealed by ultra-deep sequencing.** *Proc Natl Acad Sci* 2008, **105**(35):13081–13086.
 19. Adams RP, Ghahramani Z, Jordan MI: **Tree-structured stick breaking for hierarchical data.** In *Proceedings of the 24th Annual Conference on Neural Information Processing Systems*; 2010.
 20. Brosnan J, Iacobuzio-Donahue C: **A new branch on the tree: Next-generation sequencing in the study of cancer evolution.** *Semin Cell Dev Biol* 2012, **23**(2):237–242.
 21. Kimura M: **The number of heterozygous nucleotide sites maintained in a finite population due to steady flux of mutations.** *Genetics* 1969, **61**(4):893.
 22. Hudson RR: **Properties of a neutral allele model with intragenic recombination.** *Theor Popul Biol* 1983, **23**(2):183–201.
 23. Shattuck TM, Westra WH, Ladenson PW, Arnold A: **Independent clonal origins of distinct tumor foci in multifocal papillary thyroid carcinoma.** *N Engl J Med* 2005, **352**(23):2406–2412.
 24. **Graphviz - Graph visualization software.** <http://graphviz.org>.
 25. Levin DA, Peres Y, Wilmer EL: *Markov chains and mixing times*: AMS Bookstore; 2008.
 26. Robinson MD, Oshlack A: **A scaling normalization method for differential expression analysis of RNA-seq data.** *Genome Biol* 2010, **11**(3):R25.
 27. Teh YW: **Dirichlet processes.** In *Encyclopedia of Machine Learning*: Springer; 2010.
 28. Antoniak CE: **Mixtures of Dirichlet processes with applications to Bayesian nonparametric problems.** *Ann Stat* 1974, **2**(6):1152–1174.
 29. Sethuraman J: **A constructive definition of Dirichlet process.** *Stat Sinica* 1994, **4**:639–650.
 30. **PyClone.** <http://compbio.bccrc.ca/software/pyclone/>.
 31. Minka TP: **Estimating a Dirichlet distribution.** Tech. rep., Microsoft Research Cambridge, 2012, <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.220.175>
 32. Bansal N, Blum A, Chawla S: **Correlation clustering.** *Mach Learn* 2004, **56**(1–3):89–113.
 33. Link WA, Eaton MJ: **On thinning of chains in MCMC.** *Methods Ecol Evol* 2012, **3**:112–115.
 34. Plummer M, Best N, Cowles K, Vines K: **CODA: Convergence diagnosis and output analysis for MCMC.** *R News* 2006, **6**:7–11.

doi:10.1186/1471-2105-15-35

Cite this article as: Jiao et al.: Inferring clonal evolution of tumors from single nucleotide somatic mutations. *BMC Bioinformatics* 2014 **15**:35.

Submit your next manuscript to BioMed Central
and take full advantage of:

- Convenient online submission
- Thorough peer review
- No space constraints or color figure charges
- Immediate publication on acceptance
- Inclusion in PubMed, CAS, Scopus and Google Scholar
- Research which is freely available for redistribution

Submit your manuscript at
www.biomedcentral.com/submit

