CloneSig: Joint inference of intra-tumor heterogeneity and mutational signatures' activity in tumor bulk sequencing data

Judith Abécassis[1,2,3], Fabien Reyal[1,4], Jean-Philippe Vert[5,2*]

[1] Institut Curie, PSL Research University, Translational Research Department,
INSERM, U932 Immunity and Cancer, Residual Tumor & Response
to Treatment Laboratory (RT2Lab), F-75248, Paris, France
[2] MINES ParisTech, PSL Research University,
CBIO - Centre for Computational Biology, F-75006 Paris, France
[3] Institut Curie, PSL Research University, INSERM, U900, F-75005 Paris, France
[4] Department of Surgery, Institut Curie, F-75248 Paris, France
[5] Google Brain, F-75009 Paris, France.

Abstract

The possibility to sequence DNA in cancer samples has triggered much effort recently to identify the forces at the genomic level that shape tumorigenesis and cancer progression. It has resulted in novel understanding or clarification of two important aspects of cancer genomics: (i) intra-tumor heterogeneity (ITH), as captured by the variability in observed prevalences of somatic mutations within a tumor, and (ii) mutational processes, as revealed by the distribution of the types of somatic mutation and their immediate nucleotide context. These two aspects may not always be independent from each other, as different mutational processes could be involved in different stages or regions of the tumor, but current computational approaches to study them largely ignore this potential dependency. In particular, sequential methods that first estimate subclones and then analyze the mutational processes active in each clone can easily miss changes in mutational processes if the clonal decomposition step fails, and conversely information regarding mutational signatures is overlooked during the subclonal reconstruction. To address current limitations, we present CloneSig, a new computational method to jointly infer ITH and mutational processes in a tumor from bulk-sequencing data, including whole-exome sequencing (WES) data, by leveraging their dependency, when it exists. We show through extensive simulations that CloneSig is competitive with state-of-the-art methods for ITH inference and detection of mutational processes when the distribution of mutational signatures is the same across clones, and that CloneSig outperforms other methods when the distribution changes. We then apply CloneSig to a large cohort of 8,951 tumors with WES data from the cancer genome atlas (TCGA), and on a pan-cancer dataset of 2,632 whole-genome sequencing (WGS) tumor samples (PCAWG), where we obtain results coherent with previous studies, as well as new promising findings. This validates the applicability of CloneSig to WES and WGS data, paving the way to its use in a clinical setting where sequencing is increasingly deployed nowadays.

1 Introduction

The advent and recent democratization of high-throughput sequencing technologies has triggered much effort recently to identify the genomic forces that shape tumorigenesis and cancer progression. In particular, they have begun to shed light on evolutionary principles happening during cancer progression, and responsible for intra-tumor heterogeneity (ITH). Indeed, as proposed by Nowell in the 1970s, cancer cells progressively accumulate somatic mutations during tumorigenesis and the progression of the disease, following similar evolutionary principles as any biological population able to acquire heritable transformations [1]. As new mutations appear in a tumor, either because they bring a selective advantage or simply through neutral

---

evolution, some cancer cells may undergo clonal expansion until they represent the totality of the tumor or a substantial part of it. This may result in a tumor composed of a mosaic of cell subpopulations with specific mutations. Better understanding these processes can provide valuable insights with implications in cancer detection and monitoring, patient stratification and therapeutic strategy [2, 3, 4, 5].

Bulk genome sequencing of a tumor sample allows us in particular to capture two important aspects of ITH. First, by providing an estimate of the proportion of cells harboring each single nucleotide variant (SNV), genome sequencing allows us to assess ITH in terms of presence and proportions of subclonal populations and, to some extent, to reconstruct the evolutionary history of the tumor [6, 7, 8, 9]. This estimation is challenging, both because a unique tumor sample may miss the full extent of the true tumor heterogeneity, and because the computational problem of deconvoluting a bulk sample into subclones is notoriously difficult due to noise and lack of identifiability [6, 10]. Second, beyond their frequency in the tumor, SNVs also record traces of the mutational processes active at the time of their occurrence through biases in the sequence patterns at which they arise, as characterized with the concept of mutational signature [11]. A mutational signature is a probability distribution over possible mutation types, defined by the nature of the substitution and its trinucleotide sequence context, and reflects exogenous or endogenous causes of mutations. Forty-nine such signatures have been outlined [12], and are referenced in the COSMIC database, with known or unknown aetiologies. Deciphering signature activities in a tumor sample, and their changes over time, can provide valuable insights about the causes of cancer, the dynamic of tumor evolution and driver events, and finally help us better estimate the patient prognosis and optimize the treatment strategy [2, 5]. A few computational methods have been proposed to estimate the activity of different predefined signatures in a tumor sample from bulk genome sequencing [12, 13], also known as the signature refitting problem; we refer the reader to the reviewing work of [14, 15] for a more formal overview of existing methods.

These two aspects of genome alterations during tumor development may not always be independent from each other. For example, if a mutation triggers subclonal expansion because it activates a particular mutational process, then new mutations in the corresponding subclone may carry the mark of this process, as observed for APOBEC mutations in human bladder [16]. Alternatively, mutational processes may change over tumor development due to varying exposures to mutagens, or between cells of the same tumor due to different micro-environments and hence selective pressures. This may lead to additional changes in mutational signatures of SNVs not necessarily coinciding with clones, but possibly resulting in different signature activities between different subclones. Consequently, taking into account mutation types in addition to SNV frequencies may benefit ITH methods, although the extent of this dependency in human cancers is still unknown. Furthermore, identifying mutational processes specific to distinct subclones, and in particular detecting changes in mutational processes during cancer progression, may be of clinical interest since prognosis and treatment options may differ in that case. However, current computational pipelines for ITH and mutational process analysis largely ignore the potential dependency between these two aspects, and typically treat them independently from each other or sequentially. In the sequential approach, as for example implemented in Palimpsest [17], subclones are first identified by an ITH analysis, and in a second step mutational signatures active in each subclone are investigated. In such a sequential analysis, however, we can not observe changes in mutational signature composition if the initial clonal decomposition step fails to detect correct subclones, and we ignore information regarding mutational signatures during ITH inference. Recently, TrackSig [18] was proposed to combine these two steps by performing an evolution-aware signature activity deconvolution, in order to better detect changes in signature activity along tumor evolution. However, while TrackSig overcomes the need to rely on a previously computed subclonal reconstruction, it does not leverage the possible association between mutation frequency and mutation type to jointly infer ITH and mutation processes active in the tumor. Furthermore, by design TrackSig can only work if a sufficient number of SNV is available, limiting currently its use to whole genome sequencing (WGS) data. This is an important limitation given the popularity of whole exome sequencing (WES) to characterize tumors, particularly in the clinical setting. Recently, an extension has been proposed to better account for SNV frequencies in the change point detection, TrackSigFreq [19], but is still limited to WGS samples.

In this work, we propose CloneSig, the first method that leverages both the frequency and the mutation type of SNVs to jointly perform ITH reconstruction and decipher the activity of mutational signatures in each subclone. By exploiting the possible association between subclones and mutational processes to increase its statistical power, we show that CloneSig performs accurate estimations with fewer SNVs than competing methods, and in particular that it can be used with WES data. We show through extensive simulations, and

three independent simulated gold-standard datasets [20, 2] that CloneSig reaches state-of-the-art performance in subclonal reconstruction and signature activity deconvolution from WGS and WES data, in the presence or absence of signature activity variations between clones. We then provide a detailed analysis of 8,951 pancancer WES samples from the Cancer Genome Atlas (TCGA) with CloneSig, and on 2,632 WGS samples from the PCAWG cohort [21], where we recover results coherent with published results on WGS [21, 22, 2] as well as novel promising findings of potential clinical relevance. CloneSig is available as a Python package at `https://github.com/judithabk6/clonesig`.

2 Results

2.1 Joint estimation of ITH and mutational processes with CloneSig

We propose CloneSig, a method to jointly infer ITH and estimate mutational processes active in different clones from bulk genome sequencing data of a tumor sample. The rationale behind CloneSig is illustrated in Figure 1, which shows a scatter-plot of all SNVs detected by WES in a sarcoma (TCGA patient TCGA-3B-A9HI) along two axis: horizontally, the mutation type of the SNV, and vertically, its cancer cell fraction (CCF) estimated from WES read counts. Following previous work on mutational processes [11, 23], we consider 96 possible mutation types, defined by the nature of the substitution involved and the two flanking nucleotides. Standard methods for ITH assessment and clonal deconvolution only exploit the distribution of CCF values in the sample, as captured by the histogram on the right panel of Figure 1, while standard methods for mutational signature analysis only exploit the mutation profiles capturing the distribution of mutation contexts, as represented by the histogram on the bottom panel. However, we clearly see in the scatter-plot that these two parameters are not independent, e.g., C>A mutations tend to occur frequently at low CCF, while C>T mutations occur more frequently at high CCF. CloneSig exploit this association by working directly at the 2D scatter-plot level, in order to jointly infer subclones and mutational processes involved in those subclones. Intuitively, working at this level increases the statistical power of subclone detection when subclones are better separated in the 2D scatter-plot than on each horizontal or vertical axis, i.e., when the activity of mutational processes varies between subclones. Additional examples are shown in Supplementary Figures S103 to S108. Note that while the association between the frequency and the type of a mutation may be the consequence of a co-segregation of mutational processes and clones, which itself could occur if a change in a mutational process coincides with a change in the fitness of a cell and subsequently in the clonal composition of the tumour, CloneSig's model does *not* rely on such a strong hypothesis. In particular, an association may exist even if changes in mutational processes do not coincide with clonal evolution, and even if each clone contains heterogeneous populations of cells expressing different mutational processes. What CloneSig assumes is merely that, on average, the proportion of mutations associated to each mutational process may differ between clones. Importantly, if this assumption does not hold, then CloneSig can still be used and behave like a standard method for ITH inference, while we expect the performance of CloneSig to improve in cases where the assumption holds.

More precisely, CloneSig is based on a probabilistic graphical model [24], summarized graphically in Figure 2, to model the distribution of allelic counts and trinucleotidic contexts of SNVs in a tumor. These observed variables are statistically associated through shared unobserved latent factors, including the number of clones in the tumor, the CCF of each clone, and the mutational processes active in each clone. CloneSig infers these latent factors for each tumor from the set of SNVs by maximum likelihood estimation, using standard machinery of probabilistic graphical models. Once the parameters of the model are inferred for a given tumor, we can read from them the estimated number of subclones together with their CCF, as well as the set of mutational processes active in each clone along with their strength. In addition, for each individual SNV, CloneSig allows us to estimate the clone and the signature that generated it, in a fully probabilistic manner; for example, in Figure 1, each SNV in the scatter-plot is colored according to the most likely mutational signature that generated it, according to CloneSig. Finally, we developed a likelihood ratio-based statistical test to assess whether mutational signatures significantly differ between subclones, in order to help characterize the evolutionary process involved in the life of the tumor. We refer the reader to the Material and Methods section for all technical details regarding CloneSig.

Figure 1: **CloneSig analysis of 246 SNVs obtained by WES of a sarcoma sample (patient TCGA-3B-A9HI)**. The main panel displays all SNVs in 2 dimensions: horizontally the mutation type, which describes the type of substitution together with the flanking nucleotides, and vertically the estimated CCF, as corrected by CloneSig with the estimated mutation multiplicity. From these data CloneSig infers the presence of 2 clones and a number of mutational signatures active in the different clones. Each mutation in the main panel is colored according to the most likely mutational signature according to CloneSig. On the right panel, the CCF histogram is represented and colored with estimated clones, and superimposed with mutational signature density. The bottom panel represents the total mutation type profile. The changing pattern of mutation types with CCF is clearly visible, illustrating the opportunity for CloneSig to perform joint estimation of ITH and signature activity, while most methods so far explore separately those data, considering solely the CCF histogram in the right panel for ITH analysis, or the mutation profile of the bottom panel to infer mutational processes.

2.2 Performance for subclonal reconstruction

We first assess the ability of CloneSig to correctly reconstruct the subclonal organization of a tumor on simulated data, using four different simulators: 1) the DREAM challenge dataset devised in [20], consisting of 5 simulated WGS tumors with different sequencing depths; 2) PhylogicSim500, comprising 500 samples generated using the PhylogicNDT method, and 3) SimClone1000, comprising 972 samples generated with the simulator SimClone, both proposed by [2]; and 4) CloneSigSim, a new simulator we propose which simulates data according to the probabilistic graphical model behind CloneSig. The different simulators differ in the underlying simulation models, as well as other features such as the number of mutations simulated (DREAM, PhylogicSim500 and SimClone1000 simulate typical WGS samples while CloneSigSim simulates both WES and WGS samples), the depth of sequencing, the presence of subclonal copy number events (in DREAM and SimClone1000). Importantly, the SNVs in the DREAM dataset were simulated with a fixed distribution of signature activities across subclones, so this is a case where we do not expect a benefit for CloneSig compared to other methods. The PhylogicSim500 and SimClone1000 datasets were generated without concern about mutational signatures, so we can add a mutation type to each simulated SNV (among the 96 possibilities) using either constant or varying signature activities across clones. Similarly, for CloneSigSim, we can simulate SNV under both scenario. In order to assess the performance of CloneSig under various scenarios, we therefore consider two sets of simulations: 1) the *unfavourable* scenario, comprising DREAM, PhylogicSim500, SimClone1000 and CloneSigSim simulated with constant signature activities across subclones, and 2) the *favourable* scenario, comprising PhylogicSim500, SimClone1000 and CloneSigSim simulated with varying signature activities across subclones,

In order to measure the correctness of the subclonal reconstruction we use four different metrics adapted from [20] and described in details in the Material and Method section. Briefly, score1B measures how similar

Figure 2: **Probabilistic graphical model for CloneSig**. This plot summarizes the structure of the probabilistic graphical model underlying CloneSig. Each node represents a random variable, shaded ones being observed, and edges between two nodes describe a statistical dependency encoded as conditional distribution in CloneSig. For a given tumor we observe p, the tumor purity of the sample, and for each SNV, B_n and D_n are respectively the variant and total read counts, C_n is the copy number state, and T_n is the trinucleotide context. Unobserved latent variable include U_n, the clone or subclone where the SNV occurs, S_n, the clone-dependent mutational process that generates the mutation, and M_n, the number of chromosomal copies harboring the mutation. See the main text for details about the distributions and parameters of the model.

the true and the estimated number of clones are, score1C assesses in addition the correctness of frequency estimates for each subclone, score2A measures the adequacy between the true and predicted co-clustering matrices, and score2C the classification accuracy of clonal and subclonal mutations. We also assess the performance of eight other state-of-the-art methods for ITH estimation and compare them to CloneSig. First we evaluate TrackSig [18], that reconstructs signature activity trajectory along tumor evolution by binning mutations in groups of 100 with decreasing CCFs, and for each group performs signature activity deconvolution using an expectation-maximization (EM) algorithm. A segmentation algorithm is then applied to determine the number of breakpoints, from which we obtain subclones with different mutational processes. A recent extension integrating CCF in the segmentation algorithm to also perform subclonal reconstruction, TrackSigFreq [19] is also considered. Because of this rationale, the authors recommend to have at least 600 observed mutations to apply TrackSig or TrackSigFreq. For sake of completeness, however, we also apply TrackSig with fewer mutations in order to compare it with other methods in all settings. Third, we test Palimpsest [17], another method which associates mutational signatures and evolutionary history of a tumor. In Palimpsest, a statistical test based on the binomial distribution of variant and reference read counts for each mutation is performed, with correction for copy number, in order to classify mutations as clonal or subclonal. Then, for each of the two groups, signature activity deconvolution is performed using non-negative matrix factorization (NMF). Those methods are representative of the main approaches to the signature refitting problem: NMF-based approaches, and probability-based approaches [14, 15]. This limitation to two populations can induce a bias in the metrics 1B, 1C and 2A that are inspired from [20], so we introduce the metric 2C to account for the specificity of Palimpsest. Finally, we test five other popular methods for ITH reconstruction which do not model mutational processes: PyClone [7], PhylogicNDT [25] and DPClust [6], which are Bayesian clustering models optimized with a Markov Chain Monte Carlo (MCMC) algorithm, Ccube [8], another Bayesian clustering model, optimized with a variational inference method, and SciClone [26], also a Bayesian clustering model, optimized with a variational inference method, that only focuses on mutation in copy-number neutral regions.

Figure 3 summarizes the subclonal reconstruction performance of CloneSig and other ITH reconstruction

methods, under both the unfavourable and favourable scenarios. Each radar plot shows the average scores (1B, 1C, 2A and 2C) reached by each method on a set of simulated data. Under the unfavourable scenario, we see that CloneSig is on par with or better than the best ITH methods (PhylogicNDT, DPClust and Ccube) on all scores and across all simulators, while SciClone, TrackSig and TrackSigFreq have overall poorer performance. This confirms that CloneSig reaches state-of-the-art performance even if the main assumption underlying its model is not met in the data, which is an important property since the question of how often subclonal populations with different signature activities emerge in tumors remains open. Moving to the favourable setting, where signature activities vary across subclones, we see that, as expected, CloneSig, TrackSig and TrackSigFreq improve their performance, at least on the PhylogicSim500 and CloneSigSim simulations. Interestingly, CloneSig outperforms all methods on all simulations in that setting, confirming the potential benefits of using CloneSig on tumors where signature activities vary across subclones.

Figure 3: **Subclonal reconstruction performance of CloneSig and eight ITH reconstruction methods on simulated data.** Each radar plot summarizes the performance of the nine ITH reconstruction methods according to four performance measures on one set of simulated data. In short, score1B evaluates the number of clones found by the method, score1C the resulting mutation CCF distribution, score2A the co-clustering of mutations in the defined clones, and score2C the classification of subclonal versus clonal mutations. All scores are normalized between 0 and 1, with 1 being the best and 0 the worst. We have ensured that all scores are comparable by averaging them only on simulations where all methods successfully produced an estimate under reasonable computation time and memory limits; PyClone results are not shown on the DREAM, PhylogicSim500 and SimClone1000 datasets because it failed to produce an output too often. The four simulated sets on the left follow the unfavourable scenario, where signature activites are the same across all subclones, while the three simulated sets on the right follow the favourable scenario, where signature activites vary between subclones.

While Figure 3 shows average performances across hundreds of simulations with different parameters, it is also interesting to dig deeper into how the performance of each method fluctuates as a function of simulation parameters such as the number of clones in the tumor or the number of mutations simulated. Figure 4 summarizes this for the PhylogicSim500 simulations (in both the favourable and unfavourable settings, corresponding respectively to a "varying" or "constant" signature setting), and we refer the reader to Supplementary Note S2 for similar analyses for other simulators. Regarding the estimation of the

Figure 4: **Detailed subclonal reconstruction performance on the PhylogicSim500 simulations.** This figure shows the four performance scores (rows) of the different ITH reconstruction methods, when we vary one simulation parameter such as the number of clones simulated (left), the signature activity pattern to mimic the favourable or unfavourable scenario (middle), and the number of mutations observed rounded to the closest thousand (right). Each point represents the average of the scores over all simulated samples with a given parameter value. Bootstrap sampling of the scores was used to compute 95% confidence intervals, which are not visible if smaller than the dot. We have ensured that all scores are comparable by keeping only simulations where all methods output a prediction; PyClone results were excluded because they led to exclude too many samples.

number of clones (score1B), CloneSig, Ccube, DPClust and PhylogicNDT exhibit the best performances in most settings except for 2 clones where they are outperformed by Palimpsest, which however systematically predicts two clones by design. In the unfavourable scenario ("constant" signature activity), PhylogicNDT has the best score, which may reflect the fact that this particular dataset is simulated according to PhylogicNDT's model and might therefore be biased towards that method by construction. As already observed in Figure 3, CloneSig, TrackSig and TrackSigFreq see their performance increase in the favourable scenario, which is expected as they can leverage extra information to distinguish clones in cases where signature activity varies between distinct clones. SciClone tends to find a high number of clones, explaining its relatively good score1B for heterogeneous samples with 5 or 6 clones. However, other scores do not have the same positive evolution, revealing SciClone's limits. Regarding the impact of the number of mutations on score1B, we

see that CloneSig and TrackSig outperform all other comparable methods when the number of mutations is lower, illustrating the advantage of considering extra information in the ITH reconstruction process (we have already made clear that Palimpsest's estimation of the number of clones is irrelevant, and TrackSig's score has a very large error bar). As expected, all methods improve when the number of SNVs increases. Regarding score1C, which focuses not on the number of clones estimated by the ITH methods but on their ability to correctly recapitulate the distribution of CCF values, we see that all methods except SciClone, TrackSig and TracksigFreq have almost a perfect performance in all settings. TrackSigFreq performs slightly worse than TrackSig, but this may be explained by its poor performance when the number of mutations is too low, as performance is closer to the other methods as the number of mutations increases. Finally, SciClone is clearly the worst performing method for score1C. Besides the ability of different methods to reconstruct the correct number of subclones and their CCF, as assessed by score1B and score1C, we measure with score2A their ability to correctly assign individual mutations to their clones, an important step for downstream analysis of mutations in each subclone. Similarly to other scores, Ccube, DPClust, PhylogicNDT and CloneSig have the best (and similar performances) in the majority of settings. For all methods, score2A decreases when the number of clones increases, and increases with more observed mutations. Again, when comparing "constant" and "varying" settings for signature activity, PhylogicNDT appears as the best performing method over all "constant" samples, and CloneSig dominates in the "varying" setting. Finally, when we assess the capacity of each method to simply discriminate clonal from subclonal mutations using score2C, a measure meant not to penalize Palimpsest which only performs that task, we see again that CloneSig is among the best methods in all scenarios, in particular with fewer observed mutations, and very close to by Ccube, DPClust and PhylogicNDT in other cases. Palimpsest is a bit below these methods, as well as TrackSig and TrackSigFreq. Again, CloneSig, TrackSig and TrackSigFreq benefit from the "varying" signature activity setting.

To further illustrate the interplay between signature change and ability to detect clones, we now test CloneSig, TrackSig, TrackSigFreq and Palimpsest on simulations with exactly two clones, and where we vary how the clones differ in terms of CCF, on the one hand, and in terms of mutational processes, on the other hand (quantified in terms of cosine distance between the two profiles of mutation type). Figure 5 shows the area under the ROC curve (AUC) for the correct classification of clonal and subclonal mutations by CloneSig as a function of these two parameters. We see an increased AUC as the distance between the mutation type profiles increases, for a constant CCF difference between the clones. For example, when two clones have similar signatures (small cosine distance), the AUC is around 0.7 when the difference between their CCF is around 0.2; when their signatures are very different (large cosine distance), the same performance can be achieved when their CCF only differ by 0.1 or slightly below. We show in Supplementary Figure S110 how other parameters (number of mutations, sequencing depth, diploid proportion of the genome, choice of input signature) also impact the performance of CloneSig in this setting. We also present similar results for other methods that both account for mutation frequency and mutational context in Supplementary Figure S111, as well as the influence of the number of observed mutations for each of them in Supplementary Figure S112.

2.3   Performance for signature activity deconvolution

In addition to ITH inference in terms of subclones, CloneSig estimates the mutational processes involved in the tumor and in the different subclones. We now assess the accuracy of this estimation on simulated data, using five performance scores detailed in the Material and Methods section. In short, score_sig_1A is the Euclidean distance between the normalized mutation type counts and the reconstructed profile (activity-weighted sum of all signatures); score_sig_1B is the Euclidean distance between the true and the reconstructed profile; score_sig_1C measures the identification of the true signatures; score_sig_1D is the proportion of signatures for which the true causal signature is correctly identified; and score_sig_1E reports the median of the distribution of the cosine distance between the true and the predicted mutation type profile that generated each mutation. We compare CloneSig to the three other methods that perform both ITH and mutational process estimation, namely, TrackSig, TrackSigFreq and Palimpsest, and add also deconstructSigs [13] in the benchmark, a method that optimizes the mixture of mutational signatures of a sample through multiple linear regressions without performing subclonal reconstruction.

Figure 6 shows the performance of the different methods according to the different metrics on the PhylogicSim500 dataset. For Score_sig_1A and Score_sig_1B, there is a clear advantage for CloneSig, TrackSig and

Figure 5: Accuracy of correctly estimating the presence of two clones by CloneSig as a function of the difference in the CCF between the two clones (vertical axis), and of the cosine distance between their mutational profiles. The accuracy denotes the proportion of runs where CloneSig rightfully identifies two clones.

TrackSigFreq over Palimpsest and deconstructSigs in all scenarios except when the number of mutation is the smallest, in which case all methods behave similarly. For Score_sig_1C, CloneSig, TrackSig and TrackSigFreq exhibit the best AUC to detect present signatures in all scenarios. This may be related to a better sensitivity as CloneSig and TrackSig perform signature activity deconvolution in subsets of mutations with less noise. All methods perform similarly with respect to Score_sig_1D. The median cosine distance (Score_sig_1E) is slightly better for CloneSig compared to other methods in all settings, particularly when there are three clones or more, and when signature activity varies across subclones.

Overall, as for ITH inference, we conclude that CloneSig is as good as or better than all other tested methods in all tested scenarios. Additional results where we vary other parameters in each methods, notably the set of mutations used as inputs or the set of signatures used as prior knowledge, can be found in Supplementary Note S2; they confirm the good performance of CloneSig in all tested settings. We also present a thorough evaluation of signature activity deconvolution on the CloneSigSim, SimClone1000 and DREAM simulated datasets in Supplementary Note S2, which overall confirm the results observed on the PhylogicSim500 dataset.

2.4   Pan-cancer overview of signature changes

We now use CloneSig on real data, to analyze ITH and mutational process changes in a large cohort of 8,951 tumor WES samples from the TCGA cohort spanning 31 cancer types.

For each sample in the cohort, we estimate with CloneSig the number of subclones present in the tumor, the signatures active in each subclone, and test for the presence of a significant signature change between clones. Based on samples exhibiting a significant signature change, we attempt to identify the signatures that are the most variant for each cancer type. To that end, we compute the absolute difference in signature activity between the largest subclone and the set of clonal mutations, neglecting cases where the absolute difference is below 0.05. Figure 7 shows a global summary of the signature changes found in the cohort. For each cancer type with at least 100 patients (see Supplementary Note S3 for the analysis over all cancers), it shows the proportion of samples where a signature change is found, and a visual summary of the proportion of samples where each individual signature is found to increase or to decrease in the largest subclone, compared to the clonal mutations. The thickness of each bar, in addition, indicates the median change of each signature. We retained only signatures found variant in more than 10% of each cohort samples.

Overall, CloneSig detects a significant change in signature activity from the protected set of mutations in 39% of all samples, and in 14% when it is trained on the public set of mutations, although these proportions vary between cancer types. In terms of signature changes, we recover patterns already observed in other cohorts, usually using WGS, which confirms that CloneSig is able to detect patterns of ITH and signature activity change using WES data. We compare our findings in the TCGA to the cohort of 2,778 WGS tumors analyzed by the International Cancer Genome Consortium's Pan-Cancer Analysis of Whole Genomes

Figure 6: **Comparison of CloneSig, TrackSig, Palimpsest, and deconstructSigs for signature activity deconvolution on the PhylogicSim500 dataset.** Several metrics have been implemented, and are detailed in the main text. Scores_sig_1A and 1B are distances between the estimated mutation type profile (as defined by the signature activity proportions) to the true mutation profile (defined using the parameters used for simulations, 1A), and the empirical observed mutation profile (defined using available observed mutations, 1B), and is better when close to 0. Score_sig_1C is the area under the ROC curve for the classification of signatures as active or inactive in the sample, and is better when close to 1. Score_sig_1D is the proportion of mutations for which the correct signature was attributed, and is better when close to 1. Finally, Score_sig_1E is the median distance to the true mutation type profile of the clone to which a mutation was attributed from the true distribution of its original clone in the simulation, and is better when close to 0. The results are presented depending on several relevant covariates: the true number of clones (left), the signature activity setting (middle), and the number of mutations (right). Each point represents the average of the score over all available simulated samples. We used bootstrap sampling of the scores to compute 95% confidence intervals.

(PCAWG) initiative which represents the largest dataset of cancer WGS data to date, analyzed with two different methods: TrackSig [2], and the careful distinction of early and late (and clonal and subclonal

Figure 7: **Mutational signature changes in the TCGA cohort.** Each plot corresponds to one cancer type, indicates the number of samples with a significant signature change compared to the total number of samples, and shows on the right panel an increase of a signature in the largest subclone, compared to clonal mutations, and on the left panel a decrease. The length of each bar corresponds to the number of patients with such changes, and the thickness to the median absolute observed change. Only types with more than 100 patients are represented, full results are available in Supplementary Note S3.

mutations) in [21]. For example, similarly to the PCAWG, we observe that signature 5, of unknown aetiology, varies in almost all cancer types, and can be both increasing or decreasing. Lifestyle-associated signatures associated with tobacco-smoking (signature 4) and UV light exposure (signature 7) decrease systematically in lung tumors and oral cancers and skin melanoma respectively.

More precisely, patterns of change detected by CloneSig on the TCGA are similar to what was described on the PCAWG cohort for cervical squamous cell carcinoma and endocervical adenocarcinoma (CESC), glioblastoma multiforme (GBM), uterine carcinosarcoma (UCS) and uterine corpus endometrial carcinoma (UCEC), lung adenocarcinoma (LUAD), lung squamous cell carcinoma (LUSC) and skin cutaneous melanoma (SKCM). In addition, CloneSig detects several new patterns of variations. In stomach adenocarcinoma (STAD), the results are close, except for signature 18, reported decreasing in the most recent analysis of the PCAWG [21], but previously increasing [2] in the same data; with CloneSig and the TCGA, it varies in both directions. In bladder carcinoma (BLCA), signature 3, related to defective homologous recombination-based DNA damage repair is found increasing and was not previously reported; besides, signatures 2, 13 and 5 are found decreasing in [21], and varying in both directions in [2], and in the current study. In breast cancer (BRCA), CloneSig fails to detect the increase in signature 17 reported in [21], but not in [2], and detects three new signature variation patterns: signature 8 and 37 are increasing, and signature 30 is varying in both directions, while signatures 1 (deamination of 5-methylcytosine to thymine) and 18 (possibly damage by reactive oxygen species) tend to be preferentially decreasing and increasing respectively, instead of varying in both directions according to [2]. In prostate adenocarcinoma (PRAD), CloneSig finds signature 3 to be varying in both direction, contrary to solely increasing in [2, 21], but similarly to the findings of [22]. Signature 37 is found to vary in both directions instead of increasing. Additionally to changes identified in [2, 21], but already described in [22], CloneSig detects variations in signatures 8, 9 and 16. A new signature seems to exhibit variations along tumor evolution: signature 15 (defective DNA mismatch repair), which was not previously described in PRAD to the best of our knowledge. In esophageal carcinoma (ESCA), we do not observe the important decrease of signature 17 [2, 21], however, we describe an increase of signature 18 and a decrease of signature 16, and variations of signature 3 in both directions. For head-neck squamous cell carcinoma (HNSC), we observe similar patterns for signatures 5, 2 and 13 (related with APOBEC enzymes activity), and 18, but an undescribed two-way variation of signature 3 [2, 21], and a decrease of signature 4 (related to tobacco smoking), identified in [21] but not in [2], probably in relation to the fact that this cohort includes oral tumors. In ovary tumors (OV), increase of signature 40 and decrease of signature 5 are coherent with the findings of [2], however, CloneSig finds an important number of samples with an increase of signature 8, while a decrease of this signature was reported in [2], and an absence in [21]. CloneSig does not detect the incease of signature 39 reported in both PCAWG studies [2, 21]; regarding signature 3, it is reported increasing in [2], but found varying in both directions in our analysis, and in [21]. For thyroid carcinoma (THCA), the variations of signatures found are different, in particular for signature 2 and 13 found increasing [2, 21], while CloneSig identifies variations in both directions, and and undescribed increase of signature 16 activity, however the number of samples with a significant change of signature activity is small. In liver hepatocellular carcinoma (LIHC), and pancreatic adenocarcinoma (PAAD), we report important differences between patterns, in particular with signature 12 reported to decrease systematically in LIHC [2, 21] while we observe an increasing trend, and no variation of signatures 2, 13 and 17 in PAAD [2, 21]. In colorectal cancer (COADREAD), we observe as described in [2, 21] an strong increase in signature 40, but not 17, but an unreported increase in signature 37, a variation of signature 5 in both direction, and not only an increase, and no variation of signature 1. We also observe a slight increase in signature 26, observed in one of the three samples analyzed with single cells in [27]. In adrenocortical carcinoma (ACC), we observe variations in signature 36 (associated to defective base excision repair) and in signature 3. In kidney renal papillary cell carcinoma (KIRP) and kidney renal clear cell carcinoma (KIRC), signature 40 is strongly decreasing, and signature 1 increasing, as reported in [21]. Additionally CloneSig uncovers variations in signature 3 in about half of samples with a signature change in KIRC; activity of signature 3 in KIRC was previously outlined in [28], but undetected in the PCAWG cohort. In Kidney Chromophobe tumors (KICH), CloneSig finds a strong increase of the activity of signature 12, not reproduced in the PCAWG [2, 21], where an increase of signatures 2 and 13 are described.

To complement the analysis of the TCGA, we provide heatmaps to delineate an overview of each cancer type in Figures S38 to S68. For each type, the first panel represents the difference between subclonal and clonal signature activities (in case of a significant change in activity), and the bottom panel represents the

absolute values of each signature activity for clonal SNVs (belonging to the clone of largest CCF estimated by CloneSig), and in the main subclone (in terms of number of SNVs). This allows researchers to fully explore CloneSig's results on the TCGA, and further compare their results in future studies. For each panel, we have added several clinical variables, in particular, the patient's age at diagnosis, the stage of the tumor, the size class of the primary tumor, and the patient's sex. Overall, we found no trend of association between signature activities or change in activities and those clinical characteristics, as previously observed in the particular case of prostate cancer [22].

In most types, like CESC (Figure S41), HNSC (Figure S47) and others, we observe groups of patients with different patterns of signature activity. The clinical significance of such groups remains to be further explored.

Altogether, we note a good agreement between those new findings on the TCGA WES dataset using CloneSig, and previous similar analyses of the PCAWG WGS cohort, but also some differences which may be due to several factors including (i) the fact that both cohorts differ in size, patient clinical profiles and treatments, (ii) the fact that the algorithms used for the analysis are different, (iii) the fact that the TCGA cohort focuses on exonic mutations while PCAWG is based on WGS, and (iv) the fact that we and [2] focus on signatures with large activity changes in absolute difference while [21] focuses more on a general description of the cohort, report changes in log-fold change, and go beyond SNVs by integrating doublet-based substitutions (DBS) and small insertions and deletions (ID) signature changes. To mitigate these different effects, we provide an analysis of CloneSig's results on the TCGA using log-fold change as a metric in Supplementary Figure S37, that gives a slightly different view of the results. We also analyze the full PCAWG cohort of WGS profiles with CloneSig, and report complete results in Supplementary Note S4. We find good agreement between CloneSig's results on PCAWG and the results obtained previously on the PCAWG cohort, which confirms that CloneSig is able to detect patterns of ITH and signature activity change using both WES and WGS data, and which confirms that the differences observed between PCAWG and TCGA are primarily due to differences between cohorts or between the focus on exonic or whole genomic mutations.

3 Discussion

In recent years, a large number of methods have been developed to unravel ITH in tumors [7, 26, 8, 29, 6], and have been applied to different cohorts, including the TCGA. Recent analyses illustrate limits encountered when applying those methods to bulk WES [30, 10], as the number of observed mutations is small, the variance in read counts can be high, and a unique sample may miss the heterogeneity of the tumor. As sequencing costs are continuously decreasing, WGS, multi-sample sequencing and single cell sequencing will constitute relevant alternatives and simplify the study of ITH. However, to date a much larger number of tumor samples with sufficient clinical annotation is available with WES compared to other more advanced technologies, and can lead to interesting insights. Beyond the number of clones present in a tumor, another relevant aspect of tumor evolution is the presence of changes in mutational signatures activities [5], which could have clinical implications in cancer prevention and treatment, and unravel the evolutionary constraints shaping early tumor development. To the best of our knowledge, TrackSig [18], TrackSigFreq [19] and Palimpsest [17] are the only methods addressing the problem of systematic detection of signature changes, but they present serious limitations: Palimpsest first detects ITH, and then performs signature activity deconvolution, which has the major drawback that if this first step fails, no signature change can be detected. Moreover, Palimpsest simply aims to distinguish subclonal from clonal mutations, thus ignoring more complex patterns. TrackSig and TrackSigFreq are only applicable to WGS data, and though avoiding the caveat of relying on a previous detection of ITH for TrackSig, the final step of associating signature changes to the subclonal reconstruction is manual. TrackSigFreq is meant to be an extension of TrackSig to detect also clones not distinguished by a change in signature activity, but shows mitigated results on our benchmark compared to TrackSig. Finally, none of these methods efficiently leverages the changes in signature activity to inform and improve the ITH detection step. To overcome these limitations, we have developed CloneSig, the first method to offer joint inference of both subclonal reconstruction and signature activity deconvolution, which can be applied to WGS as well as to WES data.

3.1 Improved ITH and signature detection in WES

CloneSig is a generative probabilistic graphical model that considers somatic mutations as derived from a mixture of clones where different mutational signatures are active. We demonstrated with a thorough simulation study the benefits of the joint inference in detecting ITH, both in WES and WGS samples. We showed that CloneSig is competitive with or outperforms state-of-the art ITH methods, even in the absence of signature activity change between the clones, and is particularly efficient for the detection of samples with one or a few subclones. Interestingly, several other methods we considered including PyClone [7], SciClone [26], DPClust [6], PhylogicNDT [25] and Ccube [8], are fully Bayesian and choose the number of clones by maximizing of the posterior probability of the data. In those methods the prior has a regularizing role, and they exhibit a decrease of accuracy as the number of observed mutations increases. This may be related to the fact that the regularizing prior is less influential as more mutations are taken into account. We instead developed a specific adaptive criterion to estimate the number of clones, as we observed that standard statistical tools for model selection performed poorly in preliminary experiments.

Regarding the signature activity deconvolution problem, results on simulations (Score_sig_1C) suggest that CloneSig exhibits an improved sensitivity. Application to the TCGA also indicates such increased sensitivity: in the TCGA pancreatic ductal adenocarcinoma cohort (PAAD), the original study using deconstructSigs could not detect signature 3 activity in samples with somatic subclonal mutations in genes BRCA1 and BRCA2 [31], while CloneSig reports signature 3 exposure in some PAAD tumors.

3.2 Limits of CloneSig

It is important to explicitly state some of CloneSig's limits, a number of them largely shared by most ITH methods. CloneSig is currently limited to SNVs, and does not account for indels or structural variants. Regarding the copy number profile of the tumor, CloneSig only considers clonal segments, but provides a complete framework to estimate a mutation's multiplicity, which is an improvement over many existing methods. Regarding the range of application of CloneSig, the results are slightly below some pre-existing methods when fewer than 100 mutations are observed, as more stability is needed to fully benefit for robust estimates of signature activity, but exhibits very good performances for a number of mutations ranging from 100 to 10,000, as in the DREAM, SimClone1000, and PhylogicSim500 datasets, which makes it the most flexible method considered in this study. Regarding the runtime and scalability, resorting to subsampling above 10,000 observed mutations, as implemented in DPClust might be advisable. Another point to be vigilant about when using CloneSig, as illustrated in simulations, and based on our experience with the TCGA and the PCAWG cohorts, is that the choice of the input signatures is key to CloneSig's optimal performances. This is related to the unidentifiability of the signature activity deconvolution problem. Several solutions have been proposed: use of a pre-defined cancer-specific matrix [12, 18], selection of signatures based on other genomic information, such as patterns of indels or structural variants, or strand biases [12], or with other molecular or clinical covariates [32]. We have tested several such strategies (using all available signatures, or a well-chosen subsets), and CloneSig's results highly depend on such a choice; we made similar observations for TrackSig, TrackSigFreq and Palimpsest. The probabilistic framework of CloneSig is well suited to integrate other mutation types (indels, structural variants), as well as prior knowledge on signature co-occurrence, and a prior based on other molecular and clinical covariates. The difficulty of this approach is the possibility to learn such association patterns. Another direction for further development would be to use CloneSig's model to learn the signatures, or to allow some variations in the pre-defined signatures, as suggested in [33].

3.3 Clinical relevance of signature variations

To assess the clinical relevance of signature changes, we jointly analyzed clinical features available for the TCGA cohort and CloneSig results. We found no evidence of association between survival and the number of clones or the occurrence of a change of signature activity between subclones in any cancer type considered. We have also systematically analyzed whether we could identify an association between the exact pattern of signature change, i.e., the increase or decrease of a particular mutational signature and clinical variables, but found no significant association. However, more refined or complete analyses may be necessary to uncover the full significance of signature activity changes. Previous studies report important signature activity differences between early and metastatic tumors in endometrial and breast cancers [34, 35], with impact on

the survival in the breast cancer study [35]. We could not perform a similar analysis using the TCGA with only untreated primary tumors, but this constitutes new directions and opportunities of research using CloneSig on metastatic cohorts, for instance to refine findings of [35], that compares signatures deconvoluted from the whole metastasis, and could benefit from subclonal analysis to distinguish early and late mutations.

A final potential clinical application could be usage as a marker for personalized treatment. For example, signature 3 is associated with homologous recombination repair defect (HRD), and a targeted therapy, PARP inhibitors, can successfully target cells with such defect. We may therefore use the detection of signature 3 to identify patients that can benefit from such therapy [36], and CloneSig exhibits better identification of active signatures, as illustrated in the simulation studies. Indeed, several mutations in genes like BRCA1 and 2, RAD51 are known to cause HRD, but some other mutations are less frequent, or other events may result in HRD and be undetectable using regular genome sequencing, such as epigenetic inactivation [37]. In addition, the intensity of HRD mutational process may be predictive of the treatment response. Pursuing this line of thought, the change in signature activity could also be exploited as an indicator of the current driver status of HRD in tumor development. Of course, these remain at this stage hypotheses that require further in silico and experimental exploration. In particular relating the presence or chance of a signature to the fact that it represents a driving process amenable to therapeutic intervention remains challenging, particularly in the absence of longitudinal data about tumor evolution. As the underlying processes of signatures keeps being uncovered, we nevertheless expect more examples of such clinical applications to arise.

4 Materials and methods

4.1 CloneSig model

CloneSig is a probabilistic graphical framework, represented in Figure 2, to model the joint distribution of SNV frequency and mutational context using several latent variables to capture the subclonal composition of a tumor and the mutational processes involved in each clone. For a given SNV it assumes that we observe the following variables: D, the total number of reads covering the SNV; $B \leq D$, the number of mutated reads; $T \in \{1, \ldots, 96\}$ the index of the mutation type (i.e., the mutation and its flanking nucleotides, up to symmetry by reverse complement); and $C = (C_{normal}, C_{tumor}^{major}, C_{tumor}^{minor})$ the allele-specific copy number at the SNV locus, as inferred using existing tools such as ASCAT [38]. Here C_{normal} is the total copy number in normal cells, and $(C_{tumor}^{major}, C_{tumor}^{minor})$ are respectively the copy number in the cancer cells of the major and minor allele, respectively. We therefore also observe $C_{tumor} = C_{tumor}^{major} + C_{tumor}^{minor}$, the total copy number in cancer cells. Finally, we assume observed the tumor sample purity p, i.e., the fraction of cancer cells in the sample.

In addition to those observed variables, CloneSig models the following unobserved variables: $U \in \{1, \ldots, J\}$, the index of the clone where the SNV occurs (assuming a total of J clones); $S \in \{1, \ldots L\}$ the index of the mutational signature that generated the SNV (assuming a total of L possible signatures, given a priori); and $M \in \{1, \ldots, C_{tumor}^{major}\}$, the number of chromosomes where the SNV is present. Note that here we assume that SNVs can only be present in one of the two alleles, hence the upper bound of M by C_{tumor}^{major}.

Denoting for any integer d by $\Sigma_d = \{u \in \mathbb{R}_+^d, \sum_{i=1}^d u_i = 1\}$ the d-dimensional probability simplex, and for $u \in \Sigma_d$ by $\text{Cat}(u)$ the categorical distribution over $\{1, \ldots, d\}$ with probabilities $u_1, \ldots, u_d$ (i.e., $X \sim \text{Cat}(u)$ means that $P(X = i) = u_i$ for $i = 1, \ldots, d$), let us now describe the probability distribution encoded by CloneSig for a single SNV; its generalization to several SNVs is simply obtained by assuming they are independent and identically distributed (i.i.d.) according to the model for a single SNV. We do not model the law of C and D, which are observed root nodes in Figure 2, and therefore only explicit the conditional distribution of (U, S, T, M, B) given (C, D).

Given parameters $\xi \in \Sigma_J$, $\pi \in (\Sigma_L)^J$ and $\mu \in (\Sigma_{96})^L$, we simply model U, S and T as categorical variables:

$$U \sim \text{Cat}(\xi),$$
$$S \,|\, U \sim \text{Cat}(\pi_U),$$
$$T \,|\, S \sim \text{Cat}(\mu_S).$$

Conditionally on C, we assume that the number of mutated chromosomes M is uniformly chosen between 1

and C_{tumor}^{major}, i.e.,

$$M \,|\, C \sim \text{Cat}(1/C_{tumor}^{major}),$$

where $1/C_{tumor}^{major} \in \Sigma_{C_{tumor}^{major}}$ represents the vector of constant probability. Finally, to define the law of B, the number of mutated reads, we follow a standard approach in previous studies that represent ITH as a generative probabilistic model [7, 9, 8, 26] where the law of the mutated read counts for a given SNV must take into account the purity of the tumor, the proportion of cells in the tumor sample carrying that mutation (cancer cell fraction, CCF), as well as the various copy numbers of the normal and tumor cells. More precisely, as reviewed by [6], one can show that the expected fraction of mutated reads (variant allele frequency, VAF) satisfies

$$\text{VAF} = \frac{p \times \text{CCF} \times M}{p \times C_{tumor} + (1-p) \times C_{normal}}.$$

Note that this only holds under the classical simplifying assumption that all copy number events are clonal and affect all cells in the sample. If we now denote by $\phi \in [0,1]^J$ the vector of CCF for each clone, and introduce a further parameter $\rho \in \mathbb{R}_+^*$ to characterize the possible overdispersion of mutated read counts compared to their expected values, we finally model the number of mutated reads using a beta binomial distribution as follows:

$$B \,|\, D, U, C, M \sim \text{BetaBinomial}\left(D, \rho\phi_U\eta(M,C), \rho(1 - \phi_U\eta(M,C))\right)$$

$$\text{with } \eta(M,C) = \frac{p \times M}{p \times C_{tumor} + (1-p) \times C_{normal}}.$$

4.2 Parameter estimation

Besides the tumor purity p, we assume that the matrix of mutational processes $\mu \in (\Sigma_{96})^L$ is known, as provided by databases like COSMIC and discussed below in Section 4.8. We note that we could consider μ unknown and use CloneSig to infer a new set mutational signatures from large cohorts of sequenced tumors, but prefer to build on existing work on mutational processes in order to be able to compare the results of CloneSig to the existing literature. Besides p and μ, the free parameters or CloneSig are J, the number of clones, and $\boldsymbol{\theta} = (\xi, \phi, \pi, \rho)$ which define the distributions of all random variables. On each tumor, we optimize $\boldsymbol{\theta}$ separately for $J = 1$ to $J_{max} = 8$ clones to maximize the likelihood of the observed SNV data in the tumor. The optimization is achieved approximately by an expectation-maximization (EM) algorithm [39] detailed in Supplementary Section S1.1. The number of clones $J^* \in [1, J_{max}]$ is then estimated by maximizing an adaptive model selection criterion, detailed in Supplementary Section S1.2.

4.3 Test of mutational signature changes

We use a likelihood ratio test to determine the significance of a signature change, by comparing a regular CloneSig fit to a fit with a single mixture of signatures common to all clones. To adapt the test, the parameter of the chi-squared distribution needs a calibration, that we perform on simulated data under the null hypothesis (without change of signatures between clones). We obtain the optimal parameter using a ridge regression model with the number of clones and the degree of freedom of the input signature matrix as covariates. The coefficient values are averaged over 10-fold cross-validation to ensure robustness. We provide more details about this test in Supplementary Section S1.3.

4.4 Simulations

We use several simulation strategies to evaluate the performance of CloneSig and other methods in various situations. We also use simulations to adjust several aspects of CloneSig, in particular the setting of a custom stopping criterion and the calibration of the statistical test to detect a significant signature change along tumor evolution.

4.4.1 Default simulations

We implemented a class `SimLoader` to perform data simulation in CloneSig package. The user sets the number of clones J, the number of observed mutations N, and the matrix of L possible signatures μ. She

can also specify the desired values for the CCF of each clone $\phi \in [0,1]^J$, the proportion of each clone $\xi \in \Sigma_J$, the exposure of each signature in each clone $\pi \in (\Sigma_L)^J$, and the overdispersion parameter $\rho \in \mathbb{R}^{+*}$ for the beta-binomial distribution, as well as the proportion of the genome that is diploid. If the user does not provide values for one or several parameters, we generate them randomly as follows:

π the number of active signatures follows a $Poisson(7) + 1$ distribution, and the signatures are chosen uniformly among the L available signatures. Then for each subclone, the exposures of active signatures follow a Dirichlet distribution of parameter 1 for each active signature;

ϕ the cancer cell fraction of each clone is set such that the largest clone has a CCF of 1, and each subsequent CCF is uniformly drawn in decreasing order to be greater than 0.1, and at a distance at least 0.05 from the previous clone;

ξ the proportions of clones are drawn from a Dirichlet distribution of parameter 1 for each clone. The proportions are repeatedly drawn until the minimal proportion of a clone is greater than 0.05;

ρ follows a normal distribution of mean 60 and of variance 5.

The same strategy is used for random initialization of the parameters for the EM algorithm.

The total copy number status is drawn for a user-set diploid proportion of the genome with a bell-like distribution centered in 2, and skewed towards the right (see Supplementary Figure S109 for examples), or from a rounded log-normal distribution of parameters 1 and 0.3. The minor copy number is then drawn as the rounded product between a beta distribution of parameters 5 and 3 and the total copy number. The multiplicity of each mutation n is uniformly drawn between 1 and $C_{n,tumor_{major}}$. The purity is drawn as the minimum between a normal variable of mean 0.7 and of variance 0.1, and 0.99. The other observed variables (T, B, D) are drawn according to CloneSig probabilistic model.

4.4.2 Simulations for comparison with other ITH and signature methods

To calibrate the custom stopping criterion and for further evaluation of CloneSig, we simulated a dataset comprising $7,812$ samples using the previously described setting, with a few adjustments: we set the minimal proportion of each clone to 0.1, the minimal difference between 2 successive clone CCFs to 0.1, and we chose the active signatures among the active signatures for each of the 31 cancer types from the TCGA, and curated using literature knowledge, as described later in Methods, in section 4.8. We draw the number of active signatures as the minimum of a $Pois(7) + 1$ distribution and the number of active signatures for this cancer type. We required a cosine distance of at least 0.05 between the mutational profiles of two successive clones.

In total, for each of the 31 cancer types, we generated a simulated sample for each combination of a number of mutations from the set $\{20, 50, 100, 300, 600, 1000, 5000\}$ covering the range observed in WES and WGS data, a percentage of the genome that is diploid from the set $\{0\%, 20\%, 40\%, 60\%, 80\%, 100\%\}$ to assess the impact of copy number variations, and finally, between 1 and 6 clones.

4.4.3 Simulations without signature change between clones

We generated a set of simulations similar in all points to the one for comparison with other ITH and signature methods, except that there is a unique signature mixture common to all clones. We used this dataset in two contexts: (i) to evaluate CloneSig in comparison to other methods in the absence of signature change, and (ii) to design a statistical test to assess the significance of a change in mutational signatures. For the latter, the dataset was limited to the first eleven cancer types to avoid unnecessary computations (2,772 samples).

4.4.4 Simulations to assess the separating power of CloneSig

To assess the separating power of CloneSig, we generated a dataset of 12,000 simulated tumor samples with two clones, where each clone represents 50% of the observed SNVs. Our objective was to explore the set of the distance between two clones, in terms of CCF distance, and of cosine distance between the two mutational profiles. For that purpose we first drew ten possible CCF distances evenly on a log scale between 0 and 1, and set to 1 the largest clone CCF. We also generated 50 matrices π with cosine distances covering regularly the

possible cosine distances; to obtain them, we first generated 10,000 such π matrices to estimate an empirical distance distribution, and we implemented a rejection sampling strategy to obtain 50 samples from a uniform distribution. For each pair of CCF distance and π matrix, several samples were generated with the number of mutations varying among $\{30, 100, 300, 1000\}$, the diploid proportion of the genome among $\{0.1, 0.5, 0.9\}$, and the sequencing depth among $\{100, 500\}$.

4.4.5 Simulations to assess the sensitivity of the statistical test

To measure the sensitivity of the statistical test to detect a significant signature change along tumor evolution, we generated a dataset of 3,600 simulated tumor samples with 2 to 6 clones. We used again a rejection sampling strategy to explore the space of the maximal distance between the profiles between any 2 clones, but the target distribution is here a beta distribution of parameters 1.5 and 8 as a target distribution, as the objective was to sample more thoroughly the small cosine distances. We repeated the sampling of 30 π matrices for 2 to 6 clones, and in each case, and generated several samples with the number of mutations varying among $\{30, 100, 300, 1000\}$, the diploid proportion of the genome among $\{0.1, 0.5, 0.9\}$, and the sequencing depth among $\{100, 500\}$.

4.5 Evaluation metrics

We use several evaluation metrics to assess the quality of CloneSig and other comparable methods. Some assess specifically the accuracy of the subclonal decomposition, while others assess the performance of signature activity deconvolution.

4.5.1 Metrics evaluating the subclonal decomposition

The metrics described in this section evaluate the accuracy of the subclonal deconvolution. They are adapted from [20].

Score1B measures the adequacy between the true number of clones J_{true} and the estimated number of clones J_{pred}. It is computed as $\frac{J_{true}+1-\min(J_{true}+1,|J_{pred}-J_{true}|)}{J_{true}+1}$.

Score1C is the Wasserstein similarity, defined as 1 minus the Wasserstein distance between the true and the predicted clustering, defined by the CCFs of the different clones and their associated weights (proportion of mutations), implemented as the function `stats.wasserstein_distance` in the Python package `scipy`.

Score2A measures the correlation between the true and predicted binary co-clustering matrices in a vector form, M_{true} and M_{pred}. It is the average of 3 correlation coefficients:

Pearson correlation coefficient $PCC = \frac{\text{Cov}(M_{true}, M_{pred})}{\sigma_{M_{true}}, \sigma_{M_{pred}}}$, implemented as the function `pearsonr` in the Python package `scipy`,

Matthews correlation coefficient $\text{MCC} = \frac{TP \times TN - FP \times FN}{\sqrt{(TP+FP)(TP+FN)(TN+FP)(TN+FN)}}$, implemented as the function `metrics.matthews_corrcoef` in the Python package `scikit-learn`,

V-measure is the harmonic mean of a homogeneity score that quantifies the fact that each cluster contains only members of a single class, and a completeness score measuring if all members of a given class are assigned to the same cluster [40]; here the classes are the true clustering. We used the function `v_measure_score` in the Python package `scikit-learn`.

Before averaging, all those scores were rescaled between 0 and 1 using the score of the minimal score between two "bad scenarios": all mutations are in the same cluster, or all mutations are in their own cluster ($M_{pred} = \mathbf{1}_{N \times N}$ or $M_{pred} = \mathbb{I}_{N \times N}$). This score follows the strategy initially proposed in the preprint version [41] of [20], and suffers from a high memory usage, making it impractical beyond 20,000 mutations.

Score2C quantifies the accuracy of each method prediction of clonal and subclonal mutations. We report the accuracy, and the area under the ROC curve (implemented in function `metrics.roc_auc_score` in the Python package `scikit-learn`), sensitivity and specificity in Supplementary Note S2.

4.5.2 Metrics evaluating the identification of mutational signatures

The metrics described in this section evaluate the accuracy of the mutational signature activity deconvolution.

Score_sig_1A computes the Euclidean distance between normalized mutation type counts (empirical), and the reconstituted profile. This is the objective function of most signature reconstruction approaches (including deconstructSigs[13] and Palimpsest [17]).

Score_sig_1B is the Euclidean distance between simulated and estimated signature profiles (weighted sum over all clones). This is closer to the objective of CloneSig, TrackSig [18] and TrackSigFreq [19].

Score_sig_1C measures the ability of each method to correctly identify present signatures. For CloneSig, no signature has a null contribution to the mixture, so for each clone, the signatures are considered in the decreasing order of their contribution to the mixture, and selected until the cumulative sum reaches 0.95. This rule is applied to all methods. For that metric, the area under the ROC curve (implemented in function `metrics.roc_auc_score` in the Python package `scikit-learn`) is reported, as well as the accuracy, sensitivity, and specificity in Supplementary Note S2

Score_sig_1D is the percent of mutations with the right signature. For each mutation, the most likely signature is found by taking into account the distribution of each mutation type in each signature, and the contribution of the signature to the mixture.

Score_sig_1E measures for each mutation the cosine distance between the clonal mutation type distribution that generated the mutation and the reconstituted one. We consider a unique global distribution for deconstructSigs. This allows us to measure the relevance of the reconstruction even if the wrong signatures are selected, as several signatures have very similar profiles. The result is a distribution of distances over all mutations, and we report the median of this distribution. We also report in Supplementary Note S2 more results with the minimum, the maximum, and the standard deviation of this distribution (`max_diff_distrib_mut`, `median_diff_distrib_mut`), as well as the proportions of mutations with a distance below 0.05 or 0.1 (`perc_dist_5` and `perc_dist_10`).

4.6 Implementation

CloneSig is implemented in Python, and is available as a Python package at `https://github.com/judithabk6/clonesig`. A wrapper function implements the successive optimization of CloneSig with increasing number of clones. For two clones and more, the model is initialized using results from the precedent run with one fewer clone, by splitting the subclone with the largest contribution to the mixture entropy as described in [42]. This process is stopped when the maximum number of subclones is reached, or when the selection criterion decreases for two successive runs. A class for simulating data according to the CloneSig model is also implemented, as detailed above.

4.7 External Data

4.7.1 TCGA cohort

For the TCGA cohort, we downloaded data from the GDC data portal `https://portal.gdc.cancer.gov/`. We gathered annotated somatic mutations, both raw variant calling output, whose access is restricted and public mutations (data access project 10569 (dbGaP)), from the new unified TCGA pipeline `https://docs.gdc.cancer.gov/Data/Bioinformatics_Pipelines/DNA_Seq_Variant_Calling_Pipeline/`, with alignment to the GRCh38 assembly, and variant calling using 4 variant callers: MuSe, Mutect2, VarScan2 and SomaticSniper. Clinical data were downloaded from the cBioPortal [43]. Instructions for download can be found in the companion Github repository (`https://github.com/judithabk6/CloneSig_analysis`).

We obtained copy number alterations (CNA) data from the ASCAT complete results on TCGA data partly reported on the COSMIC database [38, 44]. We then converted ASCAT results on hg19 to GRCh38 coordinates using the `segment_liftover` Python package [45]. ASCAT results also provide an estimate of purity, which we used as input to ITH methods when possible. Other purity measures are available [46]; however we selected the ASCAT estimate to ensure consistency with CNV data.

filename	md5
consensus.20170119.somatic.cna.icgc.public.tar.gz	b16a44e5b52302136b2dfa2d2f4b2000
consensus.20170119.somatic.cna.tcga.public.tar.gz	0d05a8ec9f07aed1c70c6ebbe1db9aad
consensus.20170217.purity.ploidy.txt.gz	a3cf8c9d0521b23b31a57b0a35379d8e
final_consensus_passonly.snv_mnv_indel.icgc.public.maf.gz	009b9f88ec10d07cf826174c3dea5b78
final_consensus_passonly.snv_mnv_indel.tcga.controlled.maf.gz	a7f054a45e0ec4eab83cfe22b9adee73

Variant calling is known to be a challenging problem. It is common practice to filter variant callers output, as ITH methods are deemed to be highly sensitive to false positive SNVs. We filtered out indels from the public dataset, and considered the union of the 4 variant callers output SNVs. For the protected data, we also removed indels, and then filtered SNVs on the `FILTER` columns output by the variant caller ("PASS" only VarScan2, SomaticSniper, "PASS" or "panel_of_normals" for Mutect2, and "Tier1" to "Tier5" for MuSe). In addition, for all variant callers, we removed SNVs with a frequency in 1000 genomes or Exac greater than 0.01, except if the SNV was reported in COSMIC. A coverage filter was added, and we kept SNVs with at least 6 reads at the position in the normal sample, of which 1 maximum reports the alternative nucleotide (or with a variant allele frequency (VAF) <0.01), and for the tumor sample, at least 8 reads covering the position, of which at least 3 reporting the variant, or a VAF>0.2. The relative amount of excluded SNVs from protected to public SNV sets varied significantly between the 3 cancer types (see Table S3). All annotations are the ones downloaded from the TCGA, using VEP v84, and GENCODE v.22, sift v.5.2.2, ESP v.20141103, polyphen v.2.2.2, dbSNP v.146, Ensembl genebuild v.2014-07, Ensembl regbuild v.13.0, HGMD public v.20154, ClinVar v.201601. We further denote the filtered raw mutation set as "Protected SNVs" and the other one, which is publicly available, as "Public SNVs".

4.7.2 PCAWG cohort

For the PCAWG cohort, we downloaded SNV and copy number calls from the ICGC data portal (data access DACO-1086821). Following indications from the Supplementary materials of [21], we used the following files from the ICGC data portal (`https://dcc.icgc.org/releases/PCAWG`). No further filtering or treatment was applied, except exclusion of indels, and SNVs outside of documented CNV segments (without CNV information). Matching of cancer types to the TCGA types was done using the provided table downloadable at `https://docs.google.com/spreadsheets/d/1NqglNjsU4if6IC6A9yHQ55eImIuSmyTkrfAU7k7-kUI/edit#gid=1199917781`, and referenced in the ICGC data portal, using the second sheet (*Unique List of Tumour Types_August*), as the *Live Version* sheet did not contain the `contributing projects` field, allowing mapping with the TCGA cohorts. For types absent from the TCGA (and hence missing the associated list of present signatures), we manually used the signatures reported in Figure 3 of the associated publication [12].

4.7.3 DREAM simulated cohort

We downloaded processed data including mutation calling by Mutect2 [47], and the Battenberg algorithm [48] for CNV deciphering, and purity and ploidy estimation. We performed no further filtering, except neglecting the subclonal copy number events, as none of the evaluated method considered them (in that case, we considered the event with the largest cellular prevalence as clonal), and excluding indels, and SNVs without CNV information. To run signature methods in "cancertype" mode, that is with considering only a subset of all available 47 signatures, we attributed the tumors to a cancer type that contained all true signatures based on [12], and when several were possible, choosing a frequent one. T2 was considered of the type "Liver-HCC", T3 "Lung-AdenoCA", T4 and T5 "Breast", and T6 "Colorect-AdenoCA". All methods were run on this dataset, on a single CPU node per sample, with a time limit of 48 hours, and a memory limit of 10GB.

4.7.4 PhylogicSim500 and SimClone1000

We obtained the 500 samples in the PhylogicSim500 dataset, and the 972 samples in the SimClone1000 dataset as described in [2], from Stefan Dentro (personal communication). The data comprises the mutations called and CNV information, with subclonal CNV events in the PhylogicSim500 data but not in the SimClone1000 set. We ignored subclonal CNV events, and for each segment, we kept only the one with largest CCF in the input provided to all methods, as none of them deals with those events. As signature activity was not considered in the simulation, we used the ground truth to generate two versions of each sample, a "constant" setting with a signature activity distribution

common to all clones (corresponding to the unfavourable scenario for CloneSig), and a "varying" setting with a different signature activity distribution for each clone (corresponding to the favourable scenario for CloneSig). In each sample, we drew active signatures according to the PCAWG cancer type specified for samples of the PhylogicSim500 dataset, and a random PCAWG cancer type, as it was not set for the SimClone1000 dataset. For the PhylogicSim500 dataset, each method was run on a single CPU node per sample, with a time limit of 48 hours, and a memory limit of 5GB, and for the SimClone1000, each method was run on a single CPU node per sample, with a time limit of 48 hours, and a memory limit of 15GB, only for samples with fewer than 10,000 SNVs to avoid unnecessary computations; this limit in SNV number was set based on observations on the DREAM dataset.

4.8   Construction of a curated list of signatures associated with each cancer type

A very important input for CloneSig is the signature matrix. For application to the TCGA data, we restrict ourselves to signatures known to be active in each subtype. To that end, we downloaded the signatures found in the TCGA using SigProfiler [23] from synapse table syn11801497. The resulting list was not satisfactory as it lacked important known patterns; for instance signature 3, associated with homologous recombination repair deficiency was not found to be active in any tumor of the prostate cohort, while signature 3 in prostate cancer is well described in the literature [22, 49]. We therefore completed the signatures present in each cancer type based on the literature [50, 27, 51, 52, 53, 28, 54, 22, 55, 34, 56, 57, 12], and used the resulting matrix in all CloneSig runs on the TCGA. Our curated list of signatures present in each cancer type is provided in Table S4.

4.9   Quantification of signature changes

For analysis of the TCGA and the PCAWG cohorts, it can be of interest to go beyond the simple assessment of a significant signature change between clones, in particular to observe different patterns of variations for distinct cancer types. Using CloneSig's results, such observations remain qualitative, as the statistical framework does not analyze changes at the individual signature activity level. Qualitatively, it might be difficult to strongly assess whether a signature's activity increases or decreases because both patterns can occur in the case where more than two clones are found, and additionally, there might be complex branching evolution patterns between those clones that go beyond CloneSig's abilities. To provide a first approximation, we measure such variations between the clonal clone (defined by having the highest cellular prevalence) and the largest subclone (defined as having the largest number of mutations). We have used two metrics to report such changes:

absolute difference $A_{subclonal} - A_{clonal}$, where A is defined as the proportion of the activity (between 0 and 1). In that case, we have considered only cases where this difference is greater than 0.05 to avoid noisy reports.

relative difference $\log\left(\frac{A_{subclonal}}{A_{clonal}}\right)$; with exclusion of cases where both $A_{subclonal} < 10^{-4}$ and $A_{clonal} < 10^{-4}$, and otherwise, replacement of smaller values by 10^{-4} to avoid very large or small undue values

Each metric provides a different viewpoint of the data, with the absolute difference stressing variations of some importance, while the relative difference, though providing some focus on important variations on the logarithmic scale between two signature activities below 10^{-2}.

For the graphical representations of Figures 7, S36 and S69, we have additionally applied two thresholds: only absolute differences greater than 0.05 in absolute value were considered, and only signatures varying in more than 10% of the cohort. Those thresholds were not applied in the additional representations as heatmaps (Supplementary Figures S38 to S68, and S71 to S102).

4.10   Survival analysis

We used the package `survival` in R for the survival analysis.

List of abbreviations

WES Whole exome sequencing

WGS Whole genome sequencing

SNV Single nucleotide variant

ITH Intra-tumor heterogeneity

VAF Variant allele frequency

CCF Cancer cell fraction

EM Expectation maximization

TCGA The cancer genome atlas

ICGC International cancer genome consortium

ACC Adrenocortical carcinoma

BLCA Bladder urothelial carcinoma

LGG Brain lower grade glioma

BRCA Breast invasive carcinoma

CESC Cervical squamous cell carcinoma and endocervical adenocarcinoma

CHOL Cholangiocarcinoma

COAD Colon adenocarcinoma

ESCA Esophageal carcinoma

GBM Glioblastoma multiforme

HNSC Head and neck squamous cell carcinoma

KICH Kidney chromophobe

KIRC Kidney renal clear cell carcinoma

KIRP Kidney renal papillary cell carcinoma

LIHC Liver hepatocellular carcinoma

LUAD Lung adenocarcinoma

LUSC Lung squamous cell carcinoma

DLBC Lymphoid neoplasm diffuse large B-cell lymphoma

MESO Mesothelioma

OV Ovarian serous cystadenocarcinoma

PAAD Pancreatic adenocarcinoma

PCPG Pheochromocytoma and paraganglioma

PRAD Prostate adenocarcinoma

READ Rectum adenocarcinoma

SARC Sarcoma

SKCM Skin cutaneous melanoma

STAD Stomach adenocarcinoma

TGCT Testicular germ cell tumors

THYM Thymoma

THCA Thyroid carcinoma

UCS Uterine carcinosarcoma

UCEC Uterine corpus endometrial carcinoma

UVM Uveal melanoma

References

[1] Nowell P. The clonal evolution of tumor cell populations. *Science*, 194(4260):23–28, oct 1976. ISSN 0036-8075. doi:10.1126/science.959840.

[2] Dentro S C, Leshchiner I, Haase K, Tarabichi M, Wintersinger J, Deshwar A G, Yu K, Rubanova Y, Macintyre G, Vazquez-Garcia I, Kleinheinz K, Livitz D G, et al. Portraits of genetic intra-tumour heterogeneity and subclonal selection across cancer types. Technical Report 312041, bioRxiv, 2018. doi:10.1101/312041.

[3] Sottoriva A, Kang H, Ma Z, Graham T A, Salomon M P, Zhao J, Marjoram P, Siegmund K, Press M F, Shibata D, and Curtis C. A Big Bang model of human colorectal tumor growth. *Nature Genetics*, 47(3):209–216, mar 2015. ISSN 1061-4036. doi:10.1038/ng.3214.

[4] Turajlic S, Xu H, Litchfield K, Rowan A, Horswell S, Chambers T, O'Brien T, Lopez J I, Watkins T B, Nicol D, Stares M, Challacombe B, et al. Deterministic evolutionary trajectories influence primary tumor growth: TRACERx Renal. *Cell*, 173(3):595–610.e11, apr 2018. ISSN 00928674. doi:10.1016/j.cell.2018.03.043.

[5] Fittall M W and Van Loo P. Translating insights into tumor evolution to clinical practice: promises and challenges. *Genome Medicine*, 11(1):20, dec 2019. ISSN 1756-994X. doi:10.1186/s13073-019-0632-z.

[6] Dentro S C, Wedge D C, and Van Loo P. Principles of reconstructing the subclonal architecture of cancers. *Cold Spring Harbor Perspectives in Medicine*, 7(8):a026625, aug 2017. ISSN 2157-1422. doi:10.1101/cshperspect.a026625.

[7] Roth A, Khattra J, Yap D, Wan A, Laks E, Biele J, Ha G, Aparicio S, Bouchard-Côté A, and Shah S P. PyClone: statistical inference of clonal population structure in cancer. *Nature Methods*, 11(4):396–398, apr 2014. ISSN 1548-7091. doi:10.1038/nmeth.2883.

[8] Yuan K, Macintyre G, Liu W, Group P E, Working H, and Markowetz F. Ccube: A fast and robust method for estimating cancer cell fractions. Technical Report 484402, bioRxiv, 2018. doi:10.1101/484402.

[9] Deshwar A G, Vembu S, Yung C K, Jang G, Stein L, and Morris Q. PhyloWGS: Reconstructing subclonal composition and evolution from whole-genome sequencing of tumors. *Genome Biology*, 16(1):35, 2015. ISSN 1465-6906. doi:10.1186/s13059-015-0602-8.

[10] Shi W, Ng C K, Lim R S, Jiang T, Kumar S, Li X, Wali V B, Piscuoglio S, Gerstein M B, Chagpar A, Weigelt B, Pusztai L, Reis-Filho J S, and Hatzis C. Reliability of whole-exome sequencing for assessing intratumor genetic heterogeneity. *SSRN Electronic Journal*, 25(6):1446–1457, 2018. ISSN 1556-5068. doi:10.2139/ssrn.3155634.

[11] Alexandrov L B, Nik-Zainal S, Wedge D C, Aparicio S a J R, Behjati S, Biankin A V, Bignell G R, Bolli N, Borg A, Børresen-Dale A L, Boyault S, Burkhardt B, et al. Signatures of mutational processes in human cancer. *Nature*, 500(7463):415–421, aug 2013. ISSN 0028-0836. doi:10.1038/nature12477.

[12] Alexandrov L B, Kim J, Haradhvala N J, Huang M N, Ng A W T, Wu Y, Boot A, Covington K R, Gordenin D A, Bergstrom E N, et al. The repertoire of mutational signatures in human cancer. *Nature*, 578(7793):94–101, 2020.

[13] Rosenthal R, McGranahan N, Herrero J, Taylor B S, and Swanton C. deconstructSigs: delineating mutational processes in single tumors distinguishes DNA repair deficiencies and patterns of carcinoma evolution. *Genome Biology*, 17(1):31, dec 2016. ISSN 1474-760X. doi:10.1186/s13059-016-0893-4.

[14] Baez-Ortega A, Gori K, and Baez-Ortega C A A. Computational approaches for discovery of mutational signatures in cancer. *Briefings in Bioinformatics*, pages 1–12, 2017. ISSN 1467-5463. doi:10.1101/154716.

[15] Omichessan H, Severi G, and Perduca V. Computational tools to detect signatures of mutational processes in dna from tumours: A review and empirical comparison of performance. *PloS one*, 14(9), 2019.

[16] Lawson A R J, Abascal F, Coorens T H H, Hooks Y, O'Neill L, Latimer C, Raine K, Sanders M A, Warren A Y, Mahbubani K T A, Bareham B, Butler T M, et al. Extensive heterogeneity in somatic mutation and selection in the human bladder. *Science*, 370(6512):75–82, 2020. ISSN 0036-8075. doi:10.1126/science.aba8347.

[17] Shinde J, Bayard Q, Imbeaud S, Hirsch T Z, Liu F, Renault V, Zucman-Rossi J, and Letouzé E. Palimpsest: an R package for studying mutational and structural variant signatures along clonal evolution in cancer. *Bioinformatics*, 34(19):3380–3381, may 2018. ISSN 1367-4803. doi:10.1093/bioinformatics/bty388.

[18] Rubanova Y, Shi R, Harrigan C F, Li R, Wintersinger J, Sahin N, Deshwar A, and Morris Q. Reconstructing evolutionary trajectories of mutation signature activities in cancer using TrackSig. *Nature communications*, 11(1):1–12, 2020.

[19] Harrigan C F, Rubanova Y, Morris Q, and Selega A. TrackSigFreq: subclonal reconstructions based on mutation signatures and allele frequencies. In *Pacific Symposium on Biocomputing. Pacific Symposium on Biocomputing*, volume 25, page 238. World Scientific, 2020.

[20] Salcedo A, Tarabichi M, Espiritu S M G, Deshwar A G, David M, Wilson N M, Dentro S, Wintersinger J A, Liu L Y, Ko M, et al. A community effort to create standards for evaluating tumor subclonal reconstruction. *Nature Biotechnology*, 38(1):97–107, 2020.

[21] Gerstung M, Jolly C, Leshchiner I, Dentro S C, Gonzalez S, Rosebrock D, Mitchell T J, Rubanova Y, Anur P, Yu K, et al. The evolutionary history of 2,658 cancers. *Nature*, 578(7793):122–128, 2020.

[22] Espiritu S M G, Liu L Y, Rubanova Y, Bhandari V, Holgersen E M, Szyca L M, Fox N S, Chua M L, Yamaguchi T N, Heisler L E, Livingstone J, Wintersinger J, et al. The evolutionary landscape of localized prostate cancers drives clinical aggression. *Cell*, 173(4):1003–1013.e15, may 2018. ISSN 00928674. doi:10.1016/j.cell.2018.03.029.

[23] Alexandrov L B, Kim J, Haradhvala N J, Huang M N, Ng A W, Boot A, Covington K R, Gordenin D A, Bergstrom E, Lopez-Bigas N, Klimczak L J, McPherson J R, et al. The repertoire of mutational signatures in human cancer. Technical Report 322859, bioRxiv, 2018. doi:10.1101/322859.

[24] Koller D and Friedman N. *Probabilistic Graphical Models*. MIT Press, 2009.

[25] Leshchiner I, Livitz D, Gainor J F, Rosebrock D, Spiro O, Martinez A, Mroz E, Lin J J, Stewart C, Kim J, et al. Comprehensive analysis of tumour initiation, spatial and temporal progression under multiple lines of treatment. *bioRxiv*, page 508127, 2019.

[26] Miller C A, White B S, Dees N D, Griffith M, Welch J S, Griffith O L, Vij R, Tomasson M H, Graubert T A, Walter M J, Ellis M J, Schierding W, et al. SciClone: Inferring clonal architecture and tracking the spatial and temporal patterns of tumor evolution. *PLoS Computational Biology*, 10(8):e1003665, aug 2014. ISSN 1553-7358. doi:10.1371/journal.pcbi.1003665.

[27] Roerink S F, Sasaki N, Lee-Six H, Young M D, Alexandrov L B, Behjati S, Mitchell T J, Grossmann S, Lightfoot H, Egan D A, Pronk A, Smakman N, et al. Intra-tumour diversification in colorectal cancer at the single-cell level. *Nature*, 556(7702):457–462, apr 2018. ISSN 0028-0836. doi:10.1038/s41586-018-0024-3.

[28] Warsow G, Hübschmann D, Kleinheinz K, Nientiedt C, Heller M, Van Coile L, Tolstov Y, Trennheuser L, Wieczorek K, Pecqueux C, Gasch C, Kuru T, et al. Genomic features of renal cell carcinoma with venous tumor thrombus. *Scientific Reports*, 8(1):7477, dec 2018. ISSN 2045-2322. doi:10.1038/s41598-018-25544-z.

[29] Turajlic S, McGranahan N, and Swanton C. Inferring mutational timing and reconstructing tumour evolutionary histories. *Biochimica et Biophysica Acta (BBA) - Reviews on Cancer*, 1855(2):264–275, apr 2015. ISSN 0304419X. doi:10.1016/j.bbcan.2015.03.005.

[30] Abécassis J, Hamy A S, Laurent C, Sadacca B, Bonsang-Kitzis H, Reyal F, and Vert J P. Assessing reliability of intra-tumor heterogeneity estimates from single sample whole exome sequencing data. *PLoS One*, 2019.

[31] Raphael B J, Hruban R H, Aguirre A J, Moffitt R A, Yeh J J, Stewart C, Robertson A G, Cherniack A D, Gupta M, Getz G, Gabriel S B, Meyerson M, et al. Integrated genomic characterization of pancreatic ductal adenocarcinoma. *Cancer Cell*, 32(2):185–203.e13, 2017. ISSN 18783686. doi:10.1016/j.ccell.2017.07.007.

[32] Robinson W, Sharan R, and Leiserson M D M. Modeling clinical and molecular covariates of mutational process activity in cancer. *Bioinformatics*, 35(14):i492–i500, jul 2019. ISSN 1367-4803. doi:10.1093/bioinformatics/btz340.

[33] Volkova N V, Meier B, González-Huici V, Bertolini S, Gonzalez S, Vöhringer H, Abascal F, Martincorena I, Campbell P J, Gartner A, et al. Mutational signatures are jointly shaped by dna damage and repair. *Nature Communications*, 11(1):1–15, 2020.

[34] Ashley C W, Da Cruz Paula A, Kumar R, Mandelker D, Pei X, Riaz N, Reis-Filho J S, and Weigelt B. Analysis of mutational signatures in primary and metastatic endometrial cancer reveals distinct patterns of DNA repair defects and shifts during tumor progression. *Gynecologic Oncology*, 152(1):11–19, jan 2019. ISSN 00908258. doi:10.1016/j.ygyno.2018.10.032.

[35] Bertucci F, Ng C K Y, Patsouris A, Droin N, Piscuoglio S, Carbuccia N, Soria J C, Dien A T, Adnani Y, Kamal M, Garnier S, Meurice G, et al. Genomic characterization of metastatic breast cancers. *Nature*, 569(7757):560–564, may 2019. ISSN 0028-0836. doi:10.1038/s41586-019-1056-z.

[36] Davies H, Glodzik D, Morganella S, Yates L R, Staaf J, Zou X, Ramakrishna M, Martin S, Boyault S, Sieuwerts A M, Simpson P T, King T A, et al. HRDetect is a predictor of BRCA1 and BRCA2 deficiency based on mutational signatures. *Nature Medicine*, 23(4):517–525, 2017. ISSN 1546170X. doi:10.1038/nm.4292.

[37] Knijnenburg T A, Wang L, Zimmermann M T, Chambwe N, Gao G F, Cherniack A D, Fan H, Shen H, Way G P, Greene C S, Liu Y, Akbani R, et al. Genomic and molecular landscape of DNA damage repair deficiency across The Cancer Genome Atlas. *Cell Reports*, 23(1):239–254.e6, apr 2018. ISSN 22111247. doi:10.1016/j.celrep.2018.03.076.

[38] Martincorena I, Raine K M, Gerstung M, Dawson K J, Haase K, Van Loo P, Davies H, Stratton M R, and Campbell P J. Universal patterns of selection in cancer and somatic tissues. *Cell*, 171(5):1029–1041.e21, nov 2017. ISSN 00928674. doi:10.1016/j.cell.2017.09.042.

[39] Dempster A P, Laird N M, and Rubin D B. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*, 39(1):1–22, sep 1977. ISSN 00359246. doi: 10.1111/j.2517-6161.1977.tb01600.x.

[40] Rosenberg A and Hirschberg J. V-Measure: A conditional entropy-based external cluster evaluation measure. *EMNLP-CoNLL 2007 - Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, 1(June):410–420, 2007.

[41] Salcedo A, Tarabichi M, Espiritu S M G, Deshwar A G, David M, Wilson N M, Dentro S, Wintersinger J A, Liu L Y, Ko M, Sivanandan S, Zhang H, et al. Creating standards for evaluating tumour subclonal reconstruction. Technical Report 310425, bioRxiv, 2018.

[42] Baudry J P and Celeux G. EM for mixtures: Initialization requires special care. *Statistics and Computing*, 25(4):713–726, 2015. ISSN 15731375. doi:10.1007/s11222-015-9561-x.

[43] Gao J, Aksoy B A, Dogrusoz U, Dresdner G, Gross B, Sumer S O, Sun Y, Jacobsen A, Sinha R, Larsson E, Cerami E, Sander C, and Schultz N. Integrative analysis of complex cancer genomics and clinical profiles using the cBioPortal. *Science Signaling*, 6(269), 2013. ISSN 19450877. doi:10.1126/scisignal.2004088.

[44] Forbes S A, Beare D, Boutselakis H, Bamford S, Bindal N, Tate J, Cole C G, Ward S, Dawson E, Ponting L, Stefancsik R, Harsha B, et al. COSMIC: somatic cancer genetics at high-resolution. *Nucleic Acids Research*, 45(D1):D777–D783, jan 2017. ISSN 0305-1048. doi:10.1093/nar/gkw1121.

[45] Gao B, Huang Q, and Baudis M. segment_liftover : a Python tool to convert segments between genome assemblies. *F1000Research*, 7:319, mar 2018. ISSN 2046-1402. doi:10.12688/f1000research.14148.1.

[46] Aran D, Sirota M, and Butte A J. Systematic pan-cancer analysis of tumour purity. *Nature Communications*, 6(1):8971, dec 2015. ISSN 2041-1723. doi:10.1038/ncomms9971.

[47] Cibulskis K, Lawrence M S, Carter S L, Sivachenko A, Jaffe D, Sougnez C, Gabriel S, Meyerson M, Lander E S, and Getz G. Sensitive detection of somatic point mutations in impure and heterogeneous cancer samples. *Nature Biotechnology*, 31(3):213–219, 2013. ISSN 10870156. doi:10.1038/nbt.2514.

[48] Nik-Zainal S, Van Loo P, Wedge D C, Alexandrov L B, Greenman C D, Lau K W, Raine K, Jones D, Marshall J, Ramakrishna M, Shlien A, Cooke S L, et al. The life history of 21 breast cancers. *Cell*, 149(5):994–1007, 2012. ISSN 00928674. doi:10.1016/j.cell.2012.04.023.

[49] Riaz N, Blecua P, Lim R S, Shen R, Higginson D S, Weinhold N, Norton L, Weigelt B, Powell S N, and Reis-Filho J S. Pan-cancer analysis of bi-allelic alterations in homologous recombination DNA repair genes. *Nature Communications*, 8(1):857, dec 2017. ISSN 2041-1723. doi:10.1038/s41467-017-00921-w.

[50] Nik-Zainal S, Davies H, Staaf J, Ramakrishna M, Glodzik D, Zou X, Martincorena I, Alexandrov L B, Martin S, Wedge D C, Van Loo P, Ju Y S, et al. Landscape of somatic mutations in 560 breast cancer whole-genome sequences. *Nature*, 534(7605):47–54, jun 2016. ISSN 0028-0836. doi:10.1038/nature17676.

[51] Letouzé E, Shinde J, Renault V, Couchy G, Blanc J F, Tubacher E, Bayard Q, Bacq D, Meyer V, Semhoun J, Bioulac-Sage P, Prévôt S, et al. Mutational signatures reveal the dynamic interplay of risk factors and cellular processes during liver tumorigenesis. *Nature Communications*, 8(1):1315, dec 2017. ISSN 2041-1723. doi:10.1038/s41467-017-01358-x.

[52] Shibata T, Arai Y, and Totoki Y. Molecular genomic landscapes of hepatobiliary cancer. *Cancer Science*, 109(5):1282–1291, may 2018. ISSN 13479032. doi:10.1111/cas.13582.

[53] Ren W, Ye X, Su H, Li W, Liu D, Pirmoradian M, Wang X, Zhang B, Zhang Q, Chen L, Nie M, Liu Y, et al. Genetic landscape of hepatitis B virus–associated diffuse large B-cell lymphoma. *Blood*, 131(24):2670–2681, jun 2018. ISSN 0006-4971. doi:10.1182/blood-2017-11-817601.

[54] Royer-Bertrand B, Torsello M, Rimoldi D, El Zaoui I, Cisarova K, Pescini-Gobert R, Raynaud F, Zografos L, Schalenbourg A, Speiser D, Nicolas M, Vallat L, et al. Comprehensive genetic landscape of uveal melanoma by whole-genome sequencing. *The American Journal of Human Genetics*, 99(5):1190–1198, nov 2016. ISSN 00029297. doi:10.1016/j.ajhg.2016.09.008.

[55] Macintyre G, Goranova T E, De Silva D, Ennis D, Piskorz A M, Eldridge M, Sie D, Lewsley L A, Hanif A, Wilson C, Dowson S, Glasspool R M, et al. Copy number signatures and mutational processes in ovarian carcinoma. *Nature Genetics*, 50(9):1262–1270, sep 2018. ISSN 1061-4036. doi:10.1038/s41588-018-0179-8.

[56] Liu Y, Sethi N S, Hinoue T, Schneider B G, Cherniack A D, Sanchez-Vega F, Seoane J A, Farshidfar F, Bowlby R, Islam M, Kim J, Chatila W, et al. Comparative molecular analysis of gastrointestinal adenocarcinomas. *Cancer Cell*, 33(4):721–735.e8, apr 2018. ISSN 15356108. doi:10.1016/j.ccell.2018.03.010.

[57] Verhagen C V, Vossen D M, Borgmann K, Hageman F, Grénman R, Verwijs-Janssen M, Mout L, Kluin R J, Nieuwland M, Severson T M, Velds A, Kerkhoven R, et al. Fanconi anemia and homologous recombination gene variants are associated with functional DNA repair defects in vitro and poor outcome in patients with advanced head and neck squamous cell carcinoma. *Oncotarget*, 9(26):18198–18213, apr 2018. ISSN 19492553. doi:10.18632/oncotarget.24797.

[58] Bertsekas D P. Projected Newton methods for optimization problems with simple constraints. *SIAM Journal on Control and Optimization*, 20(2):221–246, mar 1982. ISSN 0363-0129. doi:10.1137/0320018.

[59] Martinez E Z, Achcar J A, and Aragon D C. Estimação dos parâmetros da distribuição beta-binomial: Uma aplicação usando o software SAS. *Ciência e Natura*, 37(3):12–19, sep 2015. ISSN 2179-460X. doi:10.5902/2179460X17512.

[60] Schwartz G. Estimating the dimension of a model. *Annals of Statistics*, 6(2):461–464, 1978. doi:10.1214/aos/1176344136.

[61] Akaike H. Information theory and an extension of the maximum likelihood principle. In *2nd International Symposium on Information Theory*, pages 199–213. Springer, 1998. doi:10.1007/978-1-4612-1694-0_15.

[62] Biernacki C, Celeux G, and Govaert G. Assessing a mixture model for clustering with the integrated completed likelihood. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(7):719–725, jul 2000. ISSN 01628828. doi:10.1109/34.865189.

[63] Maugis C and Michel B. A non asymptotic penalized criterion for Gaussian mixture model selection. *ESAIM: Probability and Statistics*, 15:41–68, jan 2011. ISSN 1292-8100. doi:10.1051/ps/2009004.

[64] Arlot S. Minimal penalties and the slope heuristics: a survey. *Journal de la Société Française de Statistique*, 160(3):1–160, 2019.

[65] Neyman J and Pearson E S. On the problem of the most efficient tests of statistical hypotheses. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 231(694-706):289–337, jan 1933. ISSN 1364-503X. doi:10.1098/rsta.1933.0009.

[66] Wilks S S. The large-sample distribution of the likelihood ratio for testing composite hypotheses. *The Annals of Mathematical Statistics*, 9(1):60–62, mar 1938. ISSN 0003-4851. doi:10.1214/aoms/1177732360.

Note S1   Supplementary methods

S1.1   EM algorithm for parameter estimation

In this section we detail the EM algorithm used to estimate the parameters $\boldsymbol{\theta} = (\xi, \phi, \pi, \rho)$ of CloneSig, for a given number of clones J. To lighten notations, we use in this section the notation $M_{max_n} = (C_{tumor}^{major})_n$ for the maximum value that M_n can take. We do not model the distributions of the observed variables C_n (copy number information) and D_n (total read count), and therefore only consider the following complete conditional log-likelihood:

$$\mathcal{L}(\boldsymbol{\theta}) = \log \left[\prod_{n=1}^{N} \mathbb{P}(B_n, T_n, U_n, S_n, M_n | D_n, C_n; \boldsymbol{\theta}, p, \mu) \right]$$

$$= \log \left[\prod_{n=1}^{N} \prod_{u=1}^{J} \prod_{s=1}^{L} \prod_{m=1}^{M_{max_n}} (\mathbb{P}(U_n = u; \boldsymbol{\theta})\mathbb{P}(S_n = s|U_n = u; \boldsymbol{\theta})\mathbb{P}(T_n = t|S_n = s; \mu) \right.$$

$$\left. \mathbb{P}(M_n = m|C_n)\mathbb{P}(B_n|D_n, C_n, M_n = m, U_n = u; \boldsymbol{\theta}, p))^{\mathbb{I}(S_n=s, U_n=u, M_n=m)} \right]$$

$$= \sum_{n=1}^{N} \sum_{u=1}^{J} \sum_{s=1}^{L} \sum_{m=1}^{M_{max_n}} \mathbb{I}(S_n = s, U_n = u, M_n = m) \log \left[\xi_u \pi_{us} \mu_{st} M_{max_n}^{-1} \, \mathrm{BB}(B_n \, ; \, D_n, \rho \phi_u \eta_{nm}, \rho(1 - \phi_u \eta_{nm})) \right] ,$$

where BB is the beta-binomial density:

$$\mathrm{BB}(k \, ; \, n, \alpha, \beta) = \binom{n}{k} \frac{\Gamma(k+\alpha)\Gamma(n-k+\beta)\Gamma(\alpha+\beta)}{\Gamma(n+\alpha+\beta)\Gamma(\alpha)\Gamma(\beta)} ,$$

and

$$\eta_{nm} = \frac{pm}{p \times (C_{tumor})_n + (1-p) \times (C_{normal})_n} .$$

To maximize $\mathcal{L}(\boldsymbol{\theta})$, we introduce the auxiliary function $\mathcal{Q}(\boldsymbol{\theta}, \boldsymbol{\theta}')$ as the expected value of the loglikelihood function of $\boldsymbol{\theta}$ when the latent variables follow the law with parameters $\boldsymbol{\theta}'$, that will be alternatively computed and maximized in the two steps of the EM algorithm. For that purpose, let us denote by $\mathbf{X}_n = (C_n, T_n, B_n, D_n)$ the set observed variables for the n-th SNV, and $\mathcal{D} = (\mathbf{X}_1, \ldots, \mathbf{X}_N)$ the totality of observed variables. Then we define:

$$\mathcal{Q}(\boldsymbol{\theta}, \boldsymbol{\theta}') = \mathbb{E}(\mathcal{L}(\boldsymbol{\theta})|\mathcal{D}; \boldsymbol{\theta}', p, \mu)$$

$$= \sum_{n=1}^{N} \sum_{u=1}^{J} \sum_{s=1}^{L} \sum_{m=1}^{M_{max_n}} q_{nu} r_{nus} v_{mnu} \log \left[\xi_u \pi_{us} \mu_{st} M_{max_n}^{-1} \, \mathrm{BB}(B_n \, ; \, D_n, \rho \phi_u \eta_{nm}, \rho(1 - \phi_u \eta_{nm})) \right] , \tag{1}$$

with

$$q_{nu} = \mathbb{P}(U_n = u|\mathbf{X}_n; \boldsymbol{\theta}') , \tag{2}$$

$$r_{nus} = \mathbb{P}(S_n = s|U_n = u, \mathbf{X}_n; \boldsymbol{\theta}') , \tag{3}$$

$$v_{mnu} = \mathbb{P}(M_n = m|U_n = u, \mathbf{X}_n; \boldsymbol{\theta}') . \tag{4}$$

The EM algorithm iteratively builds a sequence of estimate $\boldsymbol{\theta^1}, \boldsymbol{\theta^2}, \ldots$ by solving recursively

$$\boldsymbol{\theta^i} = \mathrm{argmax}_{\boldsymbol{\theta}} \, \mathcal{Q}(\boldsymbol{\theta}, \boldsymbol{\theta^{i-1}}) .$$

For that purpose, at each iteration i, the expectation (E) step first consists in computing the function $\mathcal{Q}(\boldsymbol{\theta}, \boldsymbol{\theta^{i-1}})$ with the current parameters $\boldsymbol{\theta^{i-1}}$. In other words, we must estimate the variables (2)-(4). Given the conditional independence relationships encoded in the graphical model (Figure 2), one easily gets:

$$q_{nu} = \frac{\sum_{s=1}^{L} \sum_{m=1}^{M_{max_n}} \xi_u^{i-1} \, \mathrm{BB}(B_n|D_n, \rho^{i-1}\phi_u^{i-1}\eta_{nm}^{i-1}, \rho^{i-1}(1-\phi_u^{i-1}\eta_{nm}^{i-1}))\mu_{sT_n}\pi_{us}^{i-1}}{\sum_{u'=1}^{J} \sum_{s=1}^{L} \sum_{m=1}^{M_{max_n}} \xi_{u'}^{i-1} \, \mathrm{BB}(B_n|D_n, \rho^{i-1}\phi_{u'}^{i-1}\eta_{nm}^{i-1}, \rho^{i-1}(1-\phi_{u'}^{i-1}\eta_{nm}^{i-1}))\mu_{sT_n}\pi_{u's}^{i-1}} , \tag{5}$$

$$r_{nus} = \frac{\mu_{sT_n}\pi_{us}^{i-1}}{\sum_{s'=1}^{M} \mu_{s'T_n}\pi_{us'}^{i-1}} , \tag{6}$$

$$v_{mnu} = \frac{\sum_{s=1}^{L} \xi_u^{i-1} \, \mathrm{BB}(B_n|D_n, \rho^{i-1}\phi_u^{i-1}\eta_{nm}^{i-1}, \rho^{i-1}(1-\phi_u^{i-1}\eta_{nm}^{i-1}))\mu_{sT_n}\pi_{us}^{i-1}}{\sum_{s=1}^{L} \sum_{m'=1}^{M_{max_n}} \xi_u^{i-1} \, \mathrm{BB}(B_n|D_n, \rho^{i-1}\phi_u^{i-1}\eta_{nm'}^{i-1}, \rho^{i-1}(1-\phi_u^{i-1}\eta_{nm'}^{i-1}))\mu_{sT_n}\pi_{us}^{i-1}} . \tag{7}$$

In the maximization (M) step, we compute $\boldsymbol{\theta}^i$ by plugging the estimates of the E-step (5)-(7) onto (1) and maximizing $\mathcal{Q}(\boldsymbol{\theta}, \boldsymbol{\theta}^{i-1})$ separately for each component of $\boldsymbol{\theta}$. The maximization in ξ and π are easily obtained as:

$$\forall u \in (1 \ldots J), \xi_u^i = \sum_{n=1}^{N} \frac{q_{nu}}{N},$$

$$\forall u \in (1 \ldots J), \forall s \in (1 \ldots L), \pi_{us}^i = \frac{\sum_{n=1}^{N} r_{nus} q_{nu}}{\sum_{n'=1}^{N} q_{n'u}}.$$

The optimization of ϕ and ρ inside the beta-binomial density term are not computable using a close formula. We therefore resort to numerical optimization and use a projected Newton method, with line search to set the Newton step at each iteration [58], in order to compute approximations of ϕ^i and ρ^i that respect constraints on their domain. Indeed, ρ must be non-negative and ϕ is a proportion so in the unit interval. For that purpose, we now compute the first and second derivatives of $\mathcal{Q}$ with respect to ϕ and $\tau = 1/\rho$:

$$\mathcal{Q}(\boldsymbol{\theta}, \boldsymbol{\theta}^{i-1}) = \sum_{n=1}^{N} \sum_{u=1}^{J} \sum_{s=1}^{L} \sum_{m=1}^{M_{max_n}} r_{nus} q_{nu} v_{mnu} \left[\log(\xi_u \mu_{st} \pi_{us} M_{max_n}^{-1}) + \log\left(\binom{d_n}{b_n} \right) \right.$$

$$+ \log(\Gamma(b_n + \frac{\phi_u \eta_{nm}}{\tau})) + \log(\Gamma(\frac{1 - \phi_u \eta_{nm}}{\tau} + d_n - b_n)) + \log(\Gamma(\frac{1}{\tau}))$$

$$\left. - \log(\Gamma(\frac{1}{\tau} + d_n)) - \log(\Gamma(\frac{\phi_u \eta_{nm}}{\tau})) - \log(\Gamma(\frac{1 - \phi_u \eta_{nm}}{\tau})) \right]$$

Let's now compute derivatives. ψ_0 and ψ_1 denote the digamma and trigamma functions respectively.

$$\frac{\partial \mathcal{Q}(\boldsymbol{\theta}, \boldsymbol{\theta}^{i-1})}{\partial \tau} = \sum_{n=1}^{N} \sum_{u=1}^{J} \sum_{m=1}^{M_{max_n}} \frac{q_{nu} v_{mnu}}{\tau^2} \left[-\eta_{nm} \phi_u \psi_0(b_n + \frac{\phi_u \eta_{nm}}{\tau}) - (1 - \eta_{nm} \phi_u) \psi_0(\frac{1 - \phi_u \eta_{nm}}{\tau} + d_n - b_n) \right.$$

$$\left. -\psi_0(\frac{1}{\tau}) + \psi_0(\frac{1}{\tau} + d_n) + \eta_{nm} \phi_u \ \psi_0(\frac{\eta_{nm} \phi_u}{\tau}) + (1 - \eta_{nm} \phi_u) \ \psi_0(\frac{1 - \eta_{nm} \phi_u}{\tau}) \right]$$

$$\frac{\partial^2 \mathcal{Q}(\boldsymbol{\theta}, \boldsymbol{\theta}^{i-1})}{\partial \tau^2} = \sum_{n=1}^{N} \sum_{u=1}^{J} \sum_{m=1}^{M_{max_n}} q_{nu} v_{mnu} \left[\frac{2}{\tau^3} \left(\eta_{nm} \phi_u \psi_0(b_n + \frac{\phi_u \eta_{nm}}{\tau}) + (1 - \eta_{nm} \phi_u) \psi_0(\frac{1 - \phi_u \eta_{nm}}{\tau} + d_n - b_n) \right. \right.$$

$$+ \psi_0(\frac{1}{\tau}) - \psi_0(\frac{1}{\tau} + d_n) - \eta_{nm} \phi_u \ \psi_0(\frac{\eta_{nm} \phi_u}{\tau}) - (1 - \eta_{nm} \phi_u) \ \psi_0(\frac{1 - \eta_{nm} \phi_u}{\tau}) \bigg)$$

$$+ \frac{1}{\tau^4} \left(\eta_{nm} 2 \phi_u 2 \psi_1(b_n + \frac{\phi_u \eta_{nm}}{\tau}) + (1 - \eta_{nm} \phi_u) 2 \psi_1(\frac{1 - \phi_u \eta_{nm}}{\tau} + d_n - b_n) + \psi_1(\frac{1}{\tau}) \right.$$

$$\left. \left. - \psi_1(\frac{1}{\tau} + d_n) - \eta_{nm} 2 \phi_u 2 \ \psi_1(\frac{\eta_{nm} \phi_u}{\tau}) - (1 - \eta_{nm} \phi_u) 2 \ \psi_1(\frac{1 - \eta_{nm} \phi_u}{\tau}) \right) \right]$$

$$\frac{\partial \mathcal{Q}(\boldsymbol{\theta}, \boldsymbol{\theta}^{i-1})}{\partial \phi_u} = \sum_{n=1}^{N} \sum_{m=1}^{M_{max_n}} q_{nu} v_{mnu} \frac{\eta_{nm}}{\tau} \left[\psi_0(b_n + \frac{\phi_u \eta_{nm}}{\tau}) - \psi_0(\frac{1 - \phi_u \eta_{nm}}{\tau} + d_n - b_n) \right.$$

$$\left. - \psi_0(\frac{\eta_{nm} \phi_u}{\tau}) + \psi_0(\frac{1 - \eta_{nm} \phi_u}{\tau}) \right]$$

$$\frac{\partial^2 \mathcal{Q}(\boldsymbol{\theta}, \boldsymbol{\theta}^{i-1})}{\partial \phi_u^2} = \sum_{n=1}^{N} \sum_{m=1}^{M_{max_n}} q_{nu} v_{mnu} \frac{\eta_{nm}^2}{\tau^2} \left[\psi_1(b_n + \frac{\phi_u \eta_{nm}}{\tau}) + \psi_1(\frac{1 - \phi_u \eta_{nm}}{\tau} + d_n - b_n) \right.$$

$$\left. - \psi_1(\frac{\eta_{nm} \phi_u}{\tau}) - \psi_1(\frac{1 - \eta_{nm} \phi_u}{\tau}) \right]$$

$$\frac{\partial^2 \mathcal{Q}(\boldsymbol{\theta}, \boldsymbol{\theta}^{i-1})}{\partial \phi_u \partial \tau} = \sum_{n=1}^{N} \sum_{m=1}^{M_{max_n}} q_{nu} v_{mnu} \frac{\eta_{nm}}{\tau^2} \left[-\psi_0(b_n + \frac{\phi_u \eta_{nm}}{\tau}) - \frac{\eta_{nm} \ \phi_u}{\tau} \psi_1(b_n + \frac{\phi_u \eta_{nm}}{\tau}) + \psi_0(\frac{1 - \phi_u \eta_{nm}}{\tau} + d_n - b_n) \right.$$

$$+ \frac{(1 - \eta_{nm} \phi_u)}{\tau} \psi_1(\frac{1 - \phi_u \eta_{nm}}{\tau} + d_n - b_n) + \psi_0(\frac{\eta_{nm} \phi_u}{\tau}) + \frac{\phi_u \eta_{nm}}{\tau} \psi_1(\frac{\eta_{nm} \phi_u}{\tau})$$

$$\left. - \psi_0(\frac{1 - \eta_{nm} \phi_u}{\tau}) - \frac{1 - \phi_u \eta_{nm}}{\tau} \psi_1(\frac{1 - \eta_{nm} \phi_u}{\tau}) \right]$$

$$\frac{\partial^2 \mathcal{Q}(\boldsymbol{\theta}, \boldsymbol{\theta}^{i-1})}{\partial \phi_u \partial \phi_{u'}} = 0$$

For sake of completeness, we provide below a second, equivalent computation using another formulation following [59].

$$\mathcal{Q}(\boldsymbol{\theta},\boldsymbol{\theta^{i-1}}) = \sum_{n=1}^{N}\sum_{u=1}^{J}\sum_{s=1}^{L}\sum_{m=1}^{M_{max_n}} r_{nus}q_{nu}v_{mnu}$$

$$\log\left[\xi_u\mu_{st}\pi_{us}M_{max_n}^{-1}\binom{d_n}{b_n}\frac{\Gamma(b_n+\rho\phi_u\eta_{nm})\Gamma(\rho(1-\phi_u\eta_{nm})+d_n-b_n)}{\Gamma(\rho+d_n)}\frac{\Gamma(\rho)}{\Gamma(\rho\phi_u\eta_{nm})\Gamma(\rho(1-\phi_u\eta_{nm}))}\right]$$

$$=\sum_{n=1}^{N}\sum_{u=1}^{J}\sum_{s=1}^{L}\sum_{m=1}^{M_{max_n}} r_{nus}q_{nu}v_{mnu}$$

$$\log\left[\xi_u\mu_{st}\pi_{us}M_{max_n}^{-1}\binom{d_n}{b_n}\frac{\prod_{i=0}^{b_n-1}(\phi_u\eta_{nm}+\frac{i}{\rho})\prod_{i=0}^{d_n-b_n-1}(1-\phi_u\eta_{nm}+\frac{i}{\rho})}{\prod_{i=0}^{d_n-1}(1+\frac{i}{\rho})}\right]$$

$$=\sum_{n=1}^{N}\sum_{u=1}^{J}\sum_{s=1}^{L}\sum_{m=1}^{M_{max_n}} r_{nus}q_{nu}v_{mnu}\left[\log(\xi_u\mu_{st}\pi_{us}M_{max_n}^{-1})+\log(\binom{d_n}{b_n})+\sum_{i=0}^{b_n-1}\left[\log(\phi_u\eta_{nm}+\frac{i}{\rho})\right]\right.$$

$$\left.+\sum_{i=0}^{d_n-b_n-1}\left[\log(1-\phi_u\eta_{nm}+\frac{i}{\rho})\right]-\sum_{i=0}^{d_n-1}\left[\log(1+\frac{i}{\rho})\right]\right]$$

Let's set $\tau=\frac{1}{\rho}$. We are trying to compute maximum likelihood estimates for ϕ_u and τ.

$$\frac{\partial\mathcal{Q}(\boldsymbol{\theta},\boldsymbol{\theta^{i-1}})}{\partial\tau}=\sum_{n=1}^{N}\sum_{u=1}^{J}\sum_{m=1}^{M_{max_n}} q_{nu}v_{mnu}\left[\sum_{i=0}^{b_n-1}\left[\frac{i}{\phi_u\eta_{nm}+i\tau}\right]+\sum_{i=0}^{d_n-b_n-1}\left[\frac{i}{1-\phi_u\eta_{nm}+i\tau}\right]-\sum_{i=0}^{d_n-1}\left[\frac{i}{1+i\tau}\right]\right]$$

$$\frac{\partial^2\mathcal{Q}(\boldsymbol{\theta},\boldsymbol{\theta^{i-1}})}{\partial\tau^2}=\sum_{n=1}^{N}\sum_{u=1}^{J}\sum_{m=1}^{M_{max_n}} q_{nu}v_{mnu}\left[-\sum_{i=0}^{b_n-1}\left[\frac{i^2}{(\phi_u\eta_{nm}+i\tau)}\right]-\sum_{i=0}^{d_n-b_n-1}\left[\frac{i^2}{(1-\phi_u\eta_{nm}+\tau)^2}\right]+\sum_{i=0}^{d_n-1}\left[\frac{i^2}{(1+\tau)^2}\right]\right]$$

$$\frac{\partial^2\mathcal{Q}(\boldsymbol{\theta},\boldsymbol{\theta^{i-1}})}{\partial\phi_u\partial\tau}=\sum_{n=1}^{N}\sum_{m=1}^{M_{max_n}} q_{nu}v_{mnu}\left[-\sum_{i=0}^{b_n-1}\left[\frac{i\eta_{nm}}{(\phi_u\eta_{nm}+i\tau)^2}\right]+\sum_{i=0}^{d_n-b_n-1}\left[\frac{i\eta_{nm}}{(1-\phi_u\eta_{nm}+i\tau)^2}\right]\right]$$

$$\frac{\partial\mathcal{Q}(\boldsymbol{\theta},\boldsymbol{\theta^{i-1}})}{\partial\phi_u}=\sum_{n=1}^{N}\sum_{m=1}^{M_{max_n}} q_{nu}v_{mnu}\left[\sum_{i=0}^{b_n-1}\left[\frac{\eta_{nm}}{\phi_u\eta_{nm}+i\tau}\right]+\sum_{i=0}^{d_n-b_n-1}\left[\frac{-\eta_{nm}}{1-\phi_u\eta_{nm}+i\tau}\right]\right]$$

$$\frac{\partial^2\mathcal{Q}(\boldsymbol{\theta},\boldsymbol{\theta^{i-1}})}{\partial\phi_u^2}=\sum_{n=1}^{N}\sum_{m=1}^{M_{max_n}} q_{nu}v_{mnu}\left[-\sum_{i=0}^{b_n-1}\left[\frac{\eta_{nm}}{\phi_u\eta_{nm}+i\tau}\right]^2-\sum_{i=0}^{d_n-b_n-1}\left[\frac{\eta_{nm}}{1-\phi_u\eta_{nm}+i\tau}\right]^2\right]$$

$$\frac{\partial^2\mathcal{Q}(\boldsymbol{\theta},\boldsymbol{\theta^{i-1}})}{\partial\phi_{u'}\partial\phi_u}=0$$

We can then plug these formulas in the projected Newton algorithm to estimate ϕ^i and ρ^i. We repeat the E and M steps until $\|\boldsymbol{\theta^i}-\boldsymbol{\theta^{i-1}}\| < 10^{-5}\times J\times L$.

S1.2   Selecting the number of clones

As explained in Supplementary Section S1.1, the EM algorithm allows us to optimize all parameters of the CloneSig model for a given number of clones J. Here we explain how to estimate J. A first idea to automatize that choice is to rely on a model selection heuristics, such as the widely used Bayesian Information Criterion (BIC) [60], an asymptotic Bayesian criterion aiming at selecting the model best supported by the data. BIC is defined as

$$BIC(J)=\ell(\mathcal{D};\boldsymbol{\theta}_J)-\frac{D_J}{2}\log N,$$

where $\ell(\mathcal{D};\boldsymbol{\theta}_J)$ is the maximum log-likelihood as estimated by the EM procedure with J clones, and D_J is the degree of freedom of the model; by default, we take it equal to the number of free parameters, namely, $D_J = J*(L-1+2)$ for J clones, where L is the number of signatures. Indeed, for each clone, we have $L-1$ parameters for the signature proportions (π), the frequency of the clone (ϕ_u), and the proportion of the clone ξ_u. We have to remove 1 because $\sum_{u=1}^{J}\xi_u=1$, and add 1 for the overdispersion parameter τ.

On simulations, however, we found that while BIC correctly identifies the number of clones when the number of SNVs is large, it tends to performs poorly when the number of mutations is low (a few hundreds) in which case it

quasi systematically selects a single clone. On the other hand, when we observe the variation of the log-likelihood with the number of components J as for example in Supplementary Figure S1, we clearly see an "elbow" for some $J > 1$, suggesting that the information about J is properly captured by CloneSig's likelihood but not by BIC.

Figure S1: Evolution of the loglikelihood and BIC criterion for 2 simulated samples, with the same parameters and 200 mutations (up panels), and 2000 mutations (bottom panels). In both cases, the loglikelihood has an "elbow" at 3 clones indicating that the likelihood of the data increases much less at the addition of an additional mixture component beyond 3 components. The BIC criterion is maximal at 3 clones when 2000 mutations are observed, but at 1 clone in the case with 200 mutations.

We observed similar behaviors with other classical criteria such as the Akaike Information Criterion (AIC) [61], the Integrated Classification Likelihood (ICL) [62]), or the slope heuristics as described in [63]. This difficulty can be related to results from statistical theory of model selection and penalization suggesting that asymptotic results are known up to a factor when applied to smaller datasets [64], and therefore propose now as an alternative an empirical criterion that can be fit on data with known model, such as simulations. More precisely, we consider the following criterion:

$$BIC_\alpha(J) = \ell(\mathcal{D}; \boldsymbol{\theta}_J) - \alpha D_J \log N . \tag{8}$$

with $\alpha > 0$ is a free parameter to be user-defined or estimated, and D_J is a measure complexity of the model.

While we leave α as a user-defined parameter in the CloneSig software, we now propose a systematic approach to estimate it when we can simulate samples. For each simulated sample, we fit CloneSig for 1 to 10 clones. The objective is to estimate a parameter α such that $BIC_{\alpha,J}$ is maximal for the true number of clones J_{true} on all or most simulations. To achieve that, we formulate it as a standard supervised classification problem where for each simulation and each $J \neq J_{true}$, we want $BIC_\alpha(J_{true}) > BIC_\alpha(J)$; since $BIC_\alpha(J)$ is itself a linear function of α, we estimate α by minimizing a convex proxy to the number of errors, namely,

$$\min_\alpha \sum_{\mathcal{D}} \sum_{J \neq J_{true}} \phi\left(BIC_\alpha(J_{true}) - BIC_\alpha(J)\right) , \tag{9}$$

where $\phi(u) = \max(0, 1 - u)$ is the hinge loss that pushes its argument to be larger than one when minimized; solving (9) is a simple support vector machine (SVM) problem that we solve with a standard SVM solver.

The second important aspect of (8) is D_J, that measures the complexity of the model with J clones. The original BIC penalizes the "dimension of the model" [60], that can be interpreted as the degree of freedom of the model, and

we now discuss different possible definitions for it. The parameters ϕ, ξ and ρ determining the CCFs and proportions of the different clones in the mixture must clearly be counted as in BIC. Regarding the signatures however, one can notice that the signatures are neither orthogonal (some signatures are very similar), nor independent (some signatures are associated with the same underlying biological process). Instead of just counting the number of signatures, we therefore propose to estimate the degree of freedom dof_L of the matrix with L signatures by the number of eigenvalues of the cosine similarity matrix greater than 0.5 in absolute value. As shown in Figure S2, dof_L is roughly proportional to L, at least for L up to 20. Another source of degree of freedom is the copy number. Indeed, for each observed mutation, several values of the number of mutated copies are considered, so if the maximal average multiplicity for mutations in the sample is $M_{max_{avg}}$, a unique clone CCF corresponds in average to $M_{max_{avg}}$ possible VAFs, adding some freedom to the model. We therefore consider four possible definitions for D_J, indexed with letters A to D.

$$D_J^A = J \times (L+1) \times M_{max_{avg}} \,, \tag{10}$$

$$D_J^B = J \times (L+1) \,, \tag{11}$$

$$D_J^C = J \times (\text{dof}_L + 1) \times M_{max_{avg}} \,, \tag{12}$$

$$D_J^D = J \times (\text{dof}_L + 1) \,. \tag{13}$$

Figure S2: Variation of the degree of freedom of a subset of cancer type-specific signatures (31 distinct types illustrated in Supplementary Table S4) or for all available signatures depending on the number of signatures (65, or 47 if the potential artefactual signatures are excluded). The left panel shows the dependence for type-specific subsets only, and the middle and right panel for the 47 and 65 signatures additionally. We see that the dependence with the number of signature (slope) is different in the two cases, but that the removal of the 12 artefactual signatures has a moderate impact on the slope compared to the use of only the subsets.

Moreover, if we consider the variations of the degree of freedom associated with L signatures, dof_L, as a function of L for the 31 available cancer types, and for the all 47 signatures, we note that there is a gap, as the maximal number of signatures in one cancer type is 23, and that the slope seems different for a subset or for all the signatures (see Supplementary Figure S2). The dependency being quite different, this raises the question of whether we should estimate a single α for all situations (i.e., a unique BIC model), or whether we should fit two BIC models: one for the cases where CloneSig is run with only cancer type-specific signatures, and one for the case where CloneSig is run with all the 47 signatures.

For each possible definition of D_J (10)-(13), and for each setting (estimating a unique or two separate BIC models), we ran simulations to estimate the value of α such that $BIC_{\alpha,j}, j \in \{1, \ldots, 10\}$ is maximal for the true value of J, by solving (9). To evaluate the results, we split the dataset into a train (80% of data) and a test set (20%), and assess the accuracy of J estimation on the test set. To evaluate the stability of the learnt parameter α, we compute the 95% confidence interval over 10 independent train-test splits. The values for learnt coefficients, averaged over 10 independent train/test splits for each case are presented in Table S1. The test accuracies for different criteria and different learning settings are presented in Figure S3. We first see that, as mentioned earlier, standard model selection criteria (BIC, AIC, ICL) perform overall poorly. Second, we notice that the "separate" strategy is usually slightly better than the "full" strategy, i.e., learning a single α for CloneSig with all 47 signatures or only a subset is not as good as learning two different α's. As for the definition of D_J, we see in both cases that using the degree of freedom of the signature matrix is better than counting the number of columns, and that taking into account the variations in copy numbers through $M_{max_{avg}}$ does not bring any benefit. A complete overview of the number of clones found

	D_J^A	D_J^B	D_J^C	D_J^D
separate model (subset)	-0.026 ± 0.000171	-0.043 ± 0.000200	-0.040 ± 0.000294	-0.065 ± 0.000357
unique model	-0.013 ± 0.000228	-0.021 ± 0.000101	-0.026 ± 0.000160	-0.042 ± 0.000199
separate model (47 signatures)	-0.0098 ± 0.000048	-0.016 ± 0.000068	-0.021 ± 0.000109	-0.034 ± 0.000149

Table S1: Values for the coefficients α for different penalty shapes and training subset. We see that the coefficients for the whole dataset and for the 47 signatures examples are close. Overall, the confidence interval for the coefficients are small.

Figure S3: Test accuracy of various model selection criteria. BIC, AIC and ICL are standard model selection. The others are attempts to learn a valid criterion on simulated data.

over the test set for each penalization strategy is given in Figure S4. In conclusion, we use in all our experiments an adaptive BIC criterion based on D_j^D as a measure of degree of freedom, and α estimated separately when CloneSig is fitted with 47 signatures or with a cancer-specific subset.

S1.3   Statistical test for signature change

To assess whether a signature change between clones is statistically significant, we design and calibrate a statistical test. To that end, we compare the likelihood of a CloneSig model with J clones as determined by the model selection criterion, and the likelihood of a model with the same clones but a single mixture of signatures common to all the clones (and found by fitting all observed mutations together). The objective of the test is to determine whether the difference between the two likelihoods is significant. To that end, we implement a likelihood-ratio test based on the statistics:

$$\lambda = \frac{\ell_{sigCst}}{\ell_{sigChange}} .$$

Following Neyman-Pearson lemma [65] one can set a threshold c to reject the null hypothesis that there is no signature change if λ is lower or equal to c with a certain level of significance α determined by the distributions of the likelihood of the model. As this distribution is unknown, we apply the Wilks theorem stating that asymptotically, $-2\log(\lambda)$ follows a chi-squared distribution of parameter the difference in dimensionality between the two alternative models [66].

As previously illustrated for the model selection criterion, the number of parameters is different from the degree of freedom in the case of CloneSig, so we resort to simulations to fit the degree of freedom of the test. We simulate a dataset with a similar mixture of signatures for all clones of each sample, and focused on samples with at least 2 clones, as described in Material and Methods. For the purpose of calibration, we use the true number of clones to fit the two alternative models. The objective of this approach is to fit a chi-squared distribution on the empirical distribution of $-2\log(\lambda)$ obtained in simulations. This is achieved again in two settings: fitting with all 47 signatures or with a cancer type-specific subset of signatures. In both cases, the distribution for each number of clones J evokes indeed a chi-squared distribution (Figure S5)

To fit the degree of freedom to use in the implementation of the test, as the degree of freedom of a chi-squared-distributed variable is its mean, we train a linear ridge regression model to fit $-2\log(\lambda)$ to relevant covariates. Four covariates were initially considered: the number of clones, the degree of freedom of the input signature matrix, the

Figure S4: Number of clones found with different model selection criteria on the test set (not used to fit the model selection criteria). This illustrates the improved accuracy of the adapted BIC criterion compared to classical criteria

Figure S5: Empirical distribution of $-2\log(\lambda)$, with $\lambda = \frac{\ell_{sigCst}}{\ell_{sigChange}}$ obtained by fitting CloneSig with the true number of clones on simulated data, either with all 47 signatures (left), or with a subset of cancer type-specific signatures. The distribution is estimated separately for each number of clones.

number of mutations, and the diploid proportion of the genome. We found that the last variable has no visible correlation with the target variable (see Supplementary Figure S6). Additionally, when added to the model, with

Figure S6: Correlation of $-2\log(\lambda)$, with $\lambda = \frac{\ell_{sigCst}}{\ell_{sigChange}}$ with potentially relevant covariates.

standard scaling of input variables, it has coefficients more than ten times smaller than the ones of the number of clones, the signature degree of freedom, and the number of mutations. We therefore compute the final model on the three retained (unscaled) variables, and we average the values of the coefficients over 10-fold cross-validation. The resulting coefficients are reported in Table S2.

To finally ascertain the validity of the test, we now check the uniform distribution of the p-values for negative samples in Figure S7. There is a slight deviation from the uniform distribution, probably due to the fact that CloneSig does not necessarily converge to the true model likelihood (and instead to a local maxima), and thus does not respect the conditions of application of Wilks theorem.

We finally explore the sensitivity of the test on the maximum cosine distance between signatures. The dataset

	Intercept	Number of Clones coefficient	Degree of freedom coefficient	Number of mutations coefficient
separate model (subset)	-14.235 ± 0.0901	5.038 ± 0.0125	1.233 ± 0.00916	0.000992 ± 0.000014
unique model	-17.620 ± 0.0577	6.436 ± 0.0142	0.909 ± 0.00220	0.00134 ± 0.000012
separate model (47 signatures)	-4.606 ± 0.06111	5.038 ± 0.0183	0 ± 0	0.00168 ± 0.000017

Table S2: values for the coefficient α for different penalty shapes and training subset. We see that the coefficients for the whole dataset and for the 47 signatures examples are close. Overall, the confidence interval for the coefficients are small.

Figure S7: Empirical distribution of the p-values of the calibrated test of significance of signature change for negative simulated samples.

used for that purpose consists of 3,600 samples with the number of clones varying between 2 and 6. For each number of clones, we drew 30 distinct π matrices with distinct maximal cosine distances between the mutation type profiles. For each number of clones and π matrix, we generated a sample with varying number of observed mutations, diploid percent of the genome, and sequencing depth. Figure S8 illustrates the proportion of samples where the test p-value is below 0.05 depending on the maximal distance between two subclones. We observe that detection is more efficient as the distance between clones becomes larger. Dependence on other variables is explored in Supplementary Figure S9. We see in particular the importance of the number of observed mutations, and of the choice of input signatures.

Figure S8: Percentage of significant tests depending on the max distance between 2 clones, quantized in 30 bins.

Figure S9: Percentage of significant tests depending several variables: number of mutations, and percentage of diploid genome, and sequencing depth, and input signature choice setting.

S1.4   Several "modes" to run CloneSig

A crucial difficulty in performing mutational signature activity deconvolution is the identifiability of the problem. Indeed, several mixtures of signatures may provide satisfying results. The most common approach to address this issue is to reduce the number of candidate signatures, in particular by using only signatures known to be active in the cancer type of the considered tumor sample [23] (approach `cancer_type`). An alternative approach is to perform two successive fits, the first one on all mutations in the sample in order to select potentially active signatures by keeping those with a contribution greater than a threshold, and the second one to refit those selected signatures with varying number of clones. This avoids the situation where a lot of signatures have very small contributions to the final mixture [18] (approach `prefit`). Those two alternatives are implemented in CloneSig (see Figure S10) and also tested for all methods tested (see supplementary Figures S13-S26). For the subclonal reconstruction problem, we see that the two approaches that limit the number of signatures have similar performance and improve the accuracy of CloneSig, especially in cases with few mutations. However, for the signature activity deconvolution problem, even though the `prefit` approach exhibits improved performance compared to taking all signatures, the `cancer_type` approach shows significantly better results. The results were similar for the other signature activity deconvolution methods, so for the rest of the analysis, we retain the `cancer_type` approach, and report only one result per method, to simplify the interpretation.

Figure S10: CloneSig's performance for 3 different input signature strategies: use all available signature (all), a subset of cancer type-specific signatures (cancer_type), or proceed in two steps by first fitting all mutations together to select potential signatures, and then actally run CloneSig with the selected subset (prefit). Additionally, the contribution of CloneSig'a approach for accounting for copy number was evaluated, by implementing the simpler approach from Palimpsest [51] (all_nuclonal).

Note S2   Full benchmarking results

We considered four datasets for the evaluation of CloneSig, and comparison of its performances to other relevant methods. Representative results obtained on the PhylogicSim500 dataset are presented in the main text, and results on other datasets in this note, for brevity.

S2.1   CloneSigSim dataset

The CloneSigSim dataset was generated using CloneSig's statistical model, and provides a very thorough analysis of CloneSig and other relevant methods in different situations, representing both WES and WGS data. A total of 10,584 samples were generated, with varying number of mutations, number of clones, percentage of diploid genome, active signatures, and variation of signature activity between clones or not.

Figure S11 summarizes the performance of the different methods according to the different metrics, and under different scenarios, where we vary respectively the number of clones in the simulation (more clones should be more challenging), the number of mutations available (more mutations should help), and the percentage of diploid genome (a higher percentage should be easier).

Regarding the estimation of the number of clones (score1B), CloneSig is the best method in all settings, except in the presence of 5 clones or more, and except for 2 clones where it is outperformed by Palimpsest, which by design systematically predicts two clones. It is in particular the only method achieving a perfect accuracy in identifying samples with one or two clones, and exhibits the best performance for score1B up to 4 clones (with the exception of Palimpsest for two clones). CloneSig, TrackSig and TrackSigFreq see their performance decrease with the number of clones, as expected, while surprisingly Ccube, DPClust and PhylogicNDT have the opposite behavior and achieve better results when the number of clones is large. During the experiments we noticed that PyClone tends to find large numbers of clones with only one mutation, so we ignore these clones when we compute score1B in order not to excessively penalize PyClone for this problematic behavior. PyClone, SciClone and Palimpsest have overall a stable performance with varying numbers of clones. Regarding the impact of the number of mutations on score1B, we see that CloneSig outperforms all other methods when enough mutations are observed. As expected, CloneSig, TrackSig and TrackSigFreq improve when the number of SNV increases, and we confirm that TrackSig requires at least 1,000 SNVs to be competitive with other methods in this experiment, while CloneSig reaches the best performance of TrackSig with as few as 100 SNVs. A surprising result is that for PyClone, SciClone, DPClust, PhylogicNDT and Ccube, score1B decreases with the number of observed mutations, which may suggest a bad calibration of the clone number estimate for large numbers of SNVs; for CloneSig we designed a specific, adaptive estimator for the number of clones since we observed that standard statistical approaches for model selection perform poorly in this setting (see Material and Methods and Supplementary Section S1.2). The percentage of diploid genome has no visible impact on the performance of any method. Regarding score1C, which focuses not on the number of clones estimated but on their ability to correctly recapitulate the distribution of CCF values, we also see that all methods except SciClone, TrackSig and TracksigFreq have almost perfect performance in all settings. TrackSig and TrackSigFreq perform slightly worse, especially as the number of clones increases, but this may be explained by its poor performance when the number of mutations is too low, as performance matches the other methods for 5,000 mutations. Finally, SciClone is clearly the worse method for score1C, particularly with 1 to 3 clones.

Besides the ability of different methods to reconstruct the correct number of subclones and their CCF, as assessed by score1B and score1C, we measure with score2A their ability to correctly assign individual mutations to their clones, an important step for downstream analysis of mutations in each subclone. According to score2A, CloneSig outperforms all other methods in almost all scenarios, illustrating the improved accuracy of accounting for both CCF and mutational signatures when achieving ITH reconstruction. For all methods, score2A decreases when the number of clones increases and when the percentage of diploid genomes decreases, as expected, but the relative order of methods does not change, with CloneSig followed by a group of three methods with similar performances: PyClone, Ccube and Palimpsest. DPClust is slightly behind, and SciClone performs poorly except when the genome is fully diploid, in which case it gets competitive with Palimpsest but still below CloneSig, PyClone and Ccube. The number of mutations has a limited impact on the performance of all methods except for TrackSig and TrackSigFreq, which only becomes competitive after 1,000 mutations. CloneSig with 100 mutations still outperforms TrackSig with 1,000 mutations, though. Finally, when we assess the capacity of each method to simply discriminate clonal from subclonal mutations using score2C, a measure meant not to penalize Palimpsest which only performs that task, we see again that CloneSig is among the best in all scenarios, except under 100 mutations, and very close to by Ccube, PyClone Palimpsest, as well as TrackSig and TrackSigFreq with 5,000 mutations. Palimpsest is a bit below these methods, while SciClone, DPClust, PhylogicNDT TrackSig and TrackSigFreq with 1,000 mutations or less are clearly not competitive for this metric.

Overall, these experiments show that CloneSig performs as well as or better than the state-of-the-art according to all metrics considered and in most simulated scenarios, confirming that accounting for the mutation type for each

Figure S11: **Comparison of CloneSig, TrackSig, TrackSigFreq, Palimpsest, PyClone, SciClone, DPClust, PhylogicNDT and Ccube for subclonal reconstruction on the CloneSigSim dataset.** Each row corresponds to one score, as detailed in the main text. In short, score1B evaluates the number of clones found by the method, score1C the resulting mutation CCF distribution, score2A the co-clustering of mutations in the defined clones, and score2C the classification of subclonal versus clonal mutations. All scores are normalized between 0 and 1, with 1 being the best and 0 the worst. Each column corresponds to a setting where one parameter in the simulation varies: the true number of clones (left), the observed number of mutations (middle), and the diploid proportion of the genome (right). Each point represents the average of the score over all available simulated samples. Bootstrap sampling of the scores was used to compute 95% confidence intervals, which are not visible if smaller than the dot indicating each value in each plot. We have ensured that all scores are comparable by removing for all methods the corresponding sample are removed when one method fails on the sample, except for systematic issues (such as SciClone and entirely non-diploid genomes, or TrackSig and TrackSigFreq for under 100 mutations).

mutation, in addition to its CCF, improves the accuracy of subclonal reconstruction. We also confirm that TrackSig and TrackSigFreq, the only existing methods that combine CCF and mutational signature information to detect subclones, require at least 1,000 mutations to obtain results competitive with other methods in our benchmark, while CloneSig reaches good accuracy in all scores with as few as 100 mutations. For smaller number of mutations, state-of-the art subclonal reconstruction methods outperform CloneSig, indicating that the signature signal may bring more noise than help at such a small resolution.

In addition to ITH inference in terms of subclones, CloneSig estimates the mutational processes involved in the tumor and in the different subclones. We now assess the accuracy of this estimation on simulated data, using six performance scores detailed in the Material and Methods section. In short, score_sig_1A is the Euclidean distance between the normalized mutation type counts and the reconstructed profile (activity-weighted sum of all signatures); score_sig_1B is the Euclidean distance between the true and the reconstructed profile; score_sig_1C measures the identification of the true signatures; score_sig_1D is the proportion of signatures for which the true causal signature is correctly identified; and score_sig_1E reports the median of the distribution of the cosine distance between the true and the predicted mutation type profile that generated each mutation. We compare CloneSig to the two other methods that perform both ITH and mutational process estimation, namely, TrackSig and Palimpsest, and add also deconstructSigs [13] in the benchmark, a method that optimizes the mixture of mutational signature of a sample through multiple linear regressions without performing subclonal reconstruction.

Figure S12 shows the performance of the different methods according to the different metrics. For Score_sig_1A and Score_sig_1B, all methods exhibit overall similar performances, with a small advantage for CloneSig, TrackSig and TrackSigFreq over Palimpsest and deconstructSigs in several scenarios. For Score_sig_1C, CloneSig, TrackSig and TrackSigFreq exhibit the best AUC to detect present signatures. It may be related to a better sensitivity as CloneSig and TrackSig perform signature activity deconvolution in smaller subsets of mutations. All methods perform similarly with respect to Score_sig_1D. The median cosine distance (Score_sig_1E) is also slightly better for CloneSig than for other methods in all settings.

Overall, as for ITH inference, we conclude that CloneSig is as good as or better than all other methods in all scenarios tested. Further results where we vary other parameters in each methods, notably the set of mutations used as inputs or the set of signatures used as prior knowledge, can be found in Supplementary Figures S13 to S27; they confirm the good performance of CloneSig in all settings tested.

To fully assess CloneSig's performance on the CloneSigSim data, in comparison with other state-of-the-art approaches for subclonal reconstruction and signature activity deconvolution, we report here the full results with all tested "modes" (all signatures, a subset of cancer-type-specific signatures, or a pre-fit step where only the most prominent signatures found on the whole set of mutations are then retained for the true signature activity deconvolution for CloneSig, TrackSig and Palimpsest). In this extensive version of the results, we report all metrics used to create score2C (AUC, specificity, sensitivity), score_sig_1C and score_sig_1E (max_diff_distrib_mut, median_diff_distrib_mut, perc_dist_5 and perc_dist_10).

Regarding the subclonal reconstruction problem, for all metrics, there is little difference between the different modes of each signature-aware method, except for score2C_sensitivity for CloneSig, where the use of the cancer-type-specific subset exhibits better results. For signature activity deconvolution, there is a higher variability of results with respect to the run mode. CloneSig is the best performing method, except for one metric: max_diff_distrib_mut. For Score_sig_1C, the mode cancer-type-specific subset for CloneSig achieves a very good specificity, but the other modes have a high proportion of false positive signatures.

Additionally, we conduct a similar benchmark in the case where there is no signature change between subclones, and present results in Supplementary Figures S13 to S27, panels b, c, f. The improvement of CloneSig over other methods in subclonal reconstruction is partially lost in this setting, but CloneSig remains competitive, and the best performing method for score 1B up to 3 clones. A similar trend is visible for all scores for the subclonal reconstruction problem, with slightly worse scores, and higher inter-quartile space when there is no signature variation between clones. For the signature activity deconvolution problem, most metrics are unaffected, except for score_sig_1E, where all methods perform better and close the gap with CloneSig. Overall, CloneSig performs better than other methods when there are differences of signature activities between subclones, and remains competitive with other approaches in the absence of signature change.

The runtimes of all methods for those simulations are presented in Figure S27. The main determinant of runtime is the number of input mutations for all methods. CloneSig is slower than methods involving variational inference for the subclonal reconstruction problem, but is significantly faster than PyClone, especially for high numbers of mutations, thus illustrating its scalability to both WES and WGS data.

Figure S12: **Comparison of CloneSig, TrackSig, Palimpsest, and deconstructSigs for signature activity deconvolution on the CloneSigSim dataset.** Several metrics have been implemented, and are detailed in the main text. Scores_sig_1A and 1B are distances between the estimated mutation type profile (as defined by the signature activity proportions) to the true mutation profile (defined using the parameters used for simulations, 1A), and the empirical observed mutation profile (defined using available observed mutations, 1B), and is better when close to 0. Score_sig_1C is the area under the ROC curve for the classification of signatures as active or inactive in the sample, and is better when close to 1. Score_sig_1D is the proportion of mutations for which the correct signature was attributed, and is better when close to 1. Finally, Score_sig_1E is the median distance to the true mutation type profile of the clone to which a mutation was attributed from the true distribution of its original clone in the simulation, and is better when close to 0. The results are presented depending on several relevant covariates: the true number of clones (left), the number of mutations (middle), and the diploid proportion of the genome (right). Each point represents the average of the score over all available simulated samples. Bootstrap sampling of the scores was used to compute 95% confidence intervals.

Figure S13: Score_1B for ITH methods on the CloneSigSim dataset, with varying number of clones (a,b), number of observed mutations (c,d) and diploid percent of the genome (e,f). Panels a, c and e correspond to simulations with varying signature between clones, and b, d, f to simulations with constant signatures.

Figure S14: Score_2A for ITH methods on the CloneSigSim dataset, with varying number of clones (a, b), number of observed mutations (c, d) and diploid percent of the genome (e, f). Panels a, c and e correspond to simulations with varying signature between clones, and b, d, f to simulations with constant signatures.

Figure S15: Score_2C (area under the curve) for ITH methods on the CloneSigSim dataset, with varying number of clones (a, b), number of observed mutations (c, d) and diploid percent of the genome (e, f). Panels a, c and e correspond to simulations with varying signature between clones, and b, d, f to simulations with constant signatures.

Figure S16: Score_2C (sensitivity) for ITH methods on the CloneSigSim dataset, with varying number of clones (a, b), number of observed mutations (c, d) and diploid percent of the genome (e, f). Panels a, c and e correspond to simulations with varying signature between clones, and b, d, f to simulations with constant signatures.

Figure S17: Score_2C (specificity) for ITH methods on the CloneSigSim dataset, with varying number of clones (a, b), number of observed mutations (c, d) and diploid percent of the genome (e, f). Panels a, c and e correspond to simulations with varying signature between clones, and b, d, f to simulations with constant signatures.

Figure S18: Score_sig_1B for signature activity deconvolution methods on the CloneSigSim dataset, with varying number of clones (a, b), number of observed mutations (c, d) and diploid percent of the genome (e, f). Panels a, c and e correspond to simulations with varying signature between clones, and b, d, f to simulations with constant signatures.

Figure S19: Score_sig_1C (accuracy) for signature activity deconvolution methods on the CloneSigSim dataset, with varying number of clones (a, b), number of observed mutations (c, d) and diploid percent of the genome (e, f). Panels a, c and e correspond to simulations with varying signature between clones, and b, d, f to simulations with constant signatures.

Figure S20: Score_sig_1C (sensitivity) for signature activity deconvolution methods on the CloneSigSim dataset, with varying number of clones (a, b), number of observed mutations (c, d) and diploid percent of the genome (e, f). Panels a, c and e correspond to simulations with varying signature between clones, and b, d, f to simulations with constant signatures.

Figure S21: Score_sig_1C (specificity) for signature activity deconvolution methods on the CloneSigSim dataset, with varying number of clones (a, b), number of observed mutations (c, d) and diploid percent of the genome (e, f). Panels a, c and e correspond to simulations with varying signature between clones, and b, d, f to simulations with constant signatures.

Figure S22: Score_sig_1D for signature activity deconvolution methods on the CloneSigSim dataset, with varying number of clones (a, b), number of observed mutations (c, d) and diploid percent of the genome (e, f). Panels a, c and e correspond to simulations with varying signature between clones, and b, d, f to simulations with constant signatures.

Figure S23: Maximal cosine distance between the true and estimated mutation type profile for signature activity deconvolution methods on the CloneSigSim dataset, with varying number of clones (a, b), number of observed mutations (c, d) and diploid percent of the genome (e, f). Panels a, c and e correspond to simulations with varying signature between clones, and b, d, f to simulations with constant signatures.

Figure S24: Median cosine distance between the true and estimated mutation type profile for signature activity deconvolution methods on the CloneSigSim dataset, with varying number of clones (a, b), number of observed mutations (c, d) and diploid percent of the genome (e, f). Panels a, c and e correspond to simulations with varying signature between clones, and b, d, f to simulations with constant signatures.

Figure S25: Proportion of SNVs with cosine distance between the true and estimated mutation type profile under 0.05 for signature activity deconvolution methods on the CloneSigSim dataset, with varying number of clones (a, b), number of observed mutations (c, d) and diploid percent of the genome (e, f). Panels a, c and e correspond to simulations with varying signature between clones, and b, d, f to simulations with constant signatures.

Figure S26: Proportion of SNVs with cosine distance between the true and estimated mutation type profile under 0.10 for signature activity deconvolution methods on the CloneSigSim dataset, with varying number of clones (a, b), number of observed mutations (c, d) and diploid percent of the genome (e, f). Panels a, c and e correspond to simulations with varying signature between clones, and b, d, f to simulations with constant signatures.

Figure S27: Runtime for ITH reconstruction and signature activity deconvolution methods on the Clone-SigSim dataset, with varying number of clones (a), number of observed mutations (b) and diploid percent of the genome (c). Results with varying signature between clones only are shown but similar results were obtained on simulations with constant signatures.

S2.2   DREAM challenge dataset

We tested all ITH inference methods on the DREAM Challenge dataset [20], with a time limit of 48 hours. Figure S28 shows the subclonal reconstruction performance metrics. We observe overall very similar performances for CloneSig, Ccube, Palimpsest, DPClust and PhylogicNDT in most cases, except for score2A where Ccube, CloneSig, Phylogic-NDT and Palimpsest clearly outperform other methods. PyClone failed to run in less than 48 hours for most samples and was not represented on this Figure. Complete results per sample are available in Supplementary Figures S29 and S30. Results for signature reconstruction metrics are shows in Supplementary Figures S31-S33.

Figure S28: **Comparison of CloneSig, TrackSig, TrackSigFreq, Palimpsest, SciClone, DPClust, PhylogicNDT and Ccube for subclonal reconstruction on the DREAM dataset.** Each row corresponds to one score, as detailed in the main text. All scores are normalized between 0 and 1, with 1 being the best and 0 the worst. Each column corresponds to a setting where one parameter in the simulation varies: the true number of clones (left), the sequencing depth (middle), and the simulated tumor (right). Each point represents the average of the score over all available simulated samples. Bootstrap sampling of the scores was used to compute 95% confidence intervals. We have ensured that all scores are comparable by removing for all methods the corresponding sample are removed when one method fails on the sample, except for systematic issues (such as SciClone and entirely non-diploid genomes, or TrackSig and TrackSigFreq for under 100 mutations). Score1B evaluates the number of clones found by the method, score1C the resulting mutation CCF distribution, score2A the co-clustering of mutations in the defined clones, and score2C the classification of subclonal versus clonal mutations.

Figure S29: **Complete results for the DREAM dataset, score1B and score1C**

Figure S30: **Complete results for the DREAM dataset, score2A and score2C AUC**

Figure S31: **Results for the DREAM dataset, signatures-related scores**

Figure S32: **Complete results for the DREAM dataset, score_sig_1A and score_sig_1B**

Figure S33: **Complete results for the DREAM dataset, score_sig_1C and score_sig_1E**

S2.3 SimClone1000 dataset

We tested all evaluated methods on the SimClone1000 simulated dataset generated in [2], for samples with less than 10,000 mutations. Results for samples in which all methods ran successfully are presented in Figure S34 for the subclonal reconstruction task, and in Figure S35 for the signature activity deconvolution task. This illustrates CloneSig's competitive performances for evaluated tasks on an independent dataset.

Figure S34: **Comparison of CloneSig, TrackSig, TrackSigFreq, Palimpsest, SciClone, DPClust, PhylogicNDT and Ccube for subclonal reconstruction on the SimClone1000 dataset.** Each row corresponds to one score, as detailed in the main text. All scores are normalized between 0 and 1, with 1 being the best and 0 the worst. Each column corresponds to a setting where one parameter in the simulation varies: the true number of clones (left), the sequencing depth (middle), and the simulated tumor (right). Each point represents the average of the score over all available simulated samples. Bootstrap sampling of the scores was used to compute 95% confidence intervals. We have ensured that all scores are comparable by removing for all methods the corresponding sample are removed when one method fails on the sample. Score1B evaluates the number of clones found by the method, score1C the resulting mutation CCF distribution, score2A the co-clustering of mutations in the defined clones, and score2C the classification of subclonal versus clonal mutations.

Figure S35: **Comparison of CloneSig, TrackSig, Palimpsest, and deconstructSigs for signature activity deconvolution on the SimClone1000 dataset.** Several metrics have been implemented, and are detailed in the main text. Scores_sig_1A and 1B are distances between the estimated mutation type profile (as defined by the signature activity proportions) to the true mutation profile (defined using the parameters used for simulations, 1A), and the empirical observed mutation profile (defined using available observed mutations, 1B), and is better when close to 0. Score_sig_1C is the area under the ROC curve for the classification of signatures as active or inactive in the sample, and is better when close to 1. Score_sig_1D is the proportion of mutations for which the correct signature was attributed, and is better when close to 1. Finally, Score_sig_1E is the median distance to the true mutation type profile of the clone to which a mutation was attributed from the true distribution of its original clone in the simulation, and is better when close to 0. The results are presented depending on several relevant covariates: the true number of clones (left), the number of mutations (middle), and the diploid proportion of the genome (right). Each point represents the average of the score over all available simulated samples. Bootstrap sampling of the scores was used to compute 95% confidence intervals.

Note S3   Complete overview of TCGA results

We now use CloneSig on real data, to analyze ITH and mutational process changes in a large cohort of 8,951 tumor WES samples from the TCGA spanning 31 cancer types. An overview of the main characteristics of the cohort is presented in Table S3.

For each sample in the cohort, we estimate with CloneSig the number of subclones present in the tumor, the signatures active in each subclone, and test for the presence of a signature change between clones. Figure S36 shows a global summary of the signature changes found in the cohort. For each cancer type, it shows the proportion of samples where a signature change is found, and a visual summary of the proportion of samples where each individual signature is found to increase or to decrease in the largest subclone, compared to the clonal mutations. The thickness of each bar, in addition, indicates the median change of each signature.

A multivariate Cox model was fitted to the TCGA data to compare three groups of tumors from the TCGA: homogeneous tumors with one clone, tumors with multiple clones but no significant signature change, and tumors with multiple clones and a significant signature change, with a stratification with respect to the tumor type. It indicates for 2 clones a hazard ratio (HR) of 1.02 (95% confidence interval (CI): $[0.91, 1.13]$, $p = 0.74$), and for 3 clones a HR of 0.97 (CI= $[0.87, 1.09]$, $p = 0.65$), hence no significant difference in survival between the three groups.

Finally, detailed results on each cancer type are presented in Supplementary Figures S38 to S68.

We additionally generated similar plots using a log-fold change in Figure S37, a relative variation of signature activity measure.

Figure S36: **Mutational signature changes in the TCGA cohort.** Each plot corresponds to one cancer type, indicates the number of samples with a significant signature change compared to the total number of samples, and shows on the right panel an increase of a signature in the largest subclone, compared to clonal mutations, and on the left panel a decrease. The length of each bar corresponds to the number of patients with such changes, and the thickness to the median absolute observed change.

Figure S37: **Mutational signature changes in the TCGA cohort.** Each plot corresponds to one cancer type, indicates the number of samples with a significant signature change compared to the total number of samples, and shows log fold change of signature activity in the largest subclone, compared to clonal mutations. The small dots correspond to individual samples, and the larger dot to the median change. Bootstrap sampling of the scores was used to compute 95% confidence intervals. The vertical dotted line in each plot represent no change.

Figure S38: Panel a: Stratification of patients depending on their pattern of signature change for ACC patients (77 patients, including 21 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S39: Panel a: Stratification of patients depending on their pattern of signature change for BLCA patients (353 patients, including 192 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S40: Panel a: Stratification of patients depending on their pattern of signature change for BRCA patients (930 patients, including 338 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S41: Panel a: Stratification of patients depending on their pattern of signature change for CESC patients (275 patients, including 186 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S42: Panel a: Stratification of patients depending on their pattern of signature change for CHOL patients (35 patients, including 7 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S43: Panel a: Stratification of patients depending on their pattern of signature change for COAD-READ patients (457 patients, including 330 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S44: Panel a: Stratification of patients depending on their pattern of signature change for DLBC patients (37 patients, including 16 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S45: Panel a: Stratification of patients depending on their pattern of signature change for ESCA patients (180 patients, including 91 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S46: Panel a: Stratification of patients depending on their pattern of signature change for GBM patients (327 patients, including 140 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF). Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S47: Panel a: Stratification of patients depending on their pattern of signature change for HNSC patients (445 patients, including 134 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S48: Panel a: Stratification of patients depending on their pattern of signature change for KICH patients (60 patients, including 43 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S49: Panel a: Stratification of patients depending on their pattern of signature change for KIRC patients (271 patients, including 47 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S50: Panel a: Stratification of patients depending on their pattern of signature change for KIRP patients (242 patients, including 90 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S51: Panel a: Stratification of patients depending on their pattern of signature change for LGG patients (455 patients, including 77 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S52: Panel a: Stratification of patients depending on their pattern of signature change for LIHC patients (347 patients, including 127 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S53: Panel a: Stratification of patients depending on their pattern of signature change for LUAD patients (433 patients, including 214 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S54: Panel a: Stratification of patients depending on their pattern of signature change for LUSC patients (423 patients, including 235 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S55: Panel a: Stratification of patients depending on their pattern of signature change for MESO patients (78 patients, including 14 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S56: Panel a: Stratification of patients depending on their pattern of signature change for OV patients (390 patients, including 220 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S57: Panel a: Stratification of patients depending on their pattern of signature change for PAAD patients (150 patients, including 23 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S58: Panel a: Stratification of patients depending on their pattern of signature change for PCPG patients (141 patients, including 5 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S59: Panel a: Stratification of patients depending on their pattern of signature change for PRAD patients (458 patients, including 66 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S60: Panel a: Stratification of patients depending on their pattern of signature change for SARC patients (210 patients, including 84 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S61: Panel a: Stratification of patients depending on their pattern of signature change for SKCM patients (424 patients, including 203 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S62: Panel a: Stratification of patients depending on their pattern of signature change for STAD patients (418 patients, including 179 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S63: Panel a: Stratification of patients depending on their pattern of signature change for TGCT patients (128 patients, including 10 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S64: Panel a: Stratification of patients depending on their pattern of signature change for THCA patients (467 patients, including 29 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF). Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S65: Panel a: Stratification of patients depending on their pattern of signature change for THYM patients (121 patients, including 32 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S66: Panel a: Stratification of patients depending on their pattern of signature change for UCEC patients (486 patients, including 236 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S67: Panel a: Stratification of patients depending on their pattern of signature change for UCS patients (53 patients, including 40 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S68: Panel a: Stratification of patients depending on their pattern of signature change for UVM patients (80 patients, including 24 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Note S4 Complete overview of PCAWG results

We applied CloneSig to the pan-cancer PCAWG cohort of 2,632 patients, spanning 32 cancer types.

For each sample in the cohort, we estimate with CloneSig the number of subclones present in the tumor, the signatures active in each subclone, and test for the presence of a significant signature change between clones. Based on samples exhibiting a significant signature change, we have attempted to identify the signatures that are the most variant for each cancer type. To that end, we have computed the absolute difference in signature activity between the largest subclone and the set of clonal mutations. We neglected cases where the absolute difference was below 0.05. Figure S69 shows a global summary of the signature changes found in the cohort. For each cancer type, it shows the proportion of samples where a signature change is found, and a visual summary of the proportion of samples where each individual signature is found to increase or to decrease in the largest subclone, compared to the clonal mutations. The thickness of each bar, in addition, indicates the median change of each signature. We retained only signatures found variant in more than 10% of each cohort samples.

We additionally generated similar plots using a log-fold change, a relative variation of signature activity measure.

Overall, CloneSig detects a significant change in signature activity from the protected set of mutations in 61% of the PCAWG cohort, in line with previous reports of 40% [21] or 76% [2], depending on the used method, although these proportions vary between cancer types. In terms of signature changes, we recover patterns already observed: for example, lifestyle-associated signatures associated with tobacco-smoking (signature 4) and UV light exposure (signature 7) decrease systematically in lung tumors and oral cancers and skin melanoma respectively, and signature 12 decreases strongly in liver tumors.

More precisely, patterns of change detected by CloneSig on the PCAWG cohort are similar to what was described previously for Biliary adenocarcinoma (Biliary-AdenoCA), Osteosarcoma (Bone-Osteosarc), Medulloblastoma (CNS-Medullo), Pilocytic astrocytoma (CNS-PiloAstro), Colorectal adenocarcinoma (ColoRect-AdenoCA), Chromophobe renal cell carcinoma (Kidney-ChRCC), Papillary renal cell carcinoma (Kidney-PRCC), Lung adenocarcinoma (Lung-AdenoCA), Lung squamous cell lung carcinoma (Lung-SCC), Melanoma (Skin-Melanoma), Gastric adenocarcinoma (Stomach-AdenoCA) and Thyroid adenocarcinoma (Thy-AdenoCA) [21, 2].

In addition, CloneSig detects several new patterns of variations. In Breast tumors, CloneSig finds signatures 2, 13 (activity of the AID/APOBEC family of cytidine deaminases) and 18 (possibly damage by reactive oxygen species) to be rather increasing in activity in subclones than varying in both directions [2, 21], and additionally finds signature 8 to be increasing, as well as signatures 30 (deficiency in base excision repair), 37, 41, that were not reported, but no variation in signature 17, reported in [21] as specifically subclonal, though in a very small number of samples, not reaching 10% of samples. In Oesophageal adenocarcinoma (Eso-AdenoCA), CloneSig finds increase in signature 3 activity and a decrease for signature 17, as reported in [2], but not in [21], where both signatures are found to vary in both directions; similarly, previous reports for signatures 2 and 13 differ, and CloneSig finds an increase, closer to what was found in [21]. In Clear cell renal cell carcinoma (Kidney-RCC), CloneSig uncovers an increase in signature 3 previously outlined in [28], but not in [21], as well as an increase in signature 29, also not previously reported, and similarly to [21] a decrease in signature 40, and an increase in signature 1 (deamination of 5-methylcytosine to thymine). In B-cell non-Hodgkin lymphoma (Lymph-BNHL), CloneSig and Tracksig [2] both identify an indrease in signature 17 and strong variations in signature 9 (mutations induced during replication by polymerase η), CloneSig additionally finds an un-reported increase in signature 37. In Ovarian adenocarcinoma (Ovary-AdenoCA), reports are again discordant between the two previous analyses of the cohort, with a strong increase in signature 39, found also in [2], but not in [21], but variations for signatures 3, 5, 8, and 40, similarly to [21]. Small differences are found in Pancreatic adenocarcinoma (Panc-AdenoCA), with all analyses reporting a strong decrease in signatures 1 and 5, an increase in signatures 18, 2 and 13, and 3, the two other studies report a similar trend for signature 17, not found by CloneSig. For Pancreatic neuroendocrine tumours (Panc-Endocrine), both CloneSig and [21] identify increases in signatures 8, 30 and 36, and a decrease in signature 5, while TrackSig [2] misses signature 30 and 36, but additionally identifies signature 18, and finds signature 3 increasing while varying in both directions in the other two studies. For Prostate adenocarcinoma (Prost-AdenoCA), both CloneSig and TrackSig [2] found strong increase of signatures 37 and 3, absent in [21]; all studies find a decreasing tendency for signatures 1 and 5, and variations in signatures 40 and 18 in both directions, except for CloneSig in which signature 18 is increasing only. CloneSig additionally reports increase in variations in signatures 8 and 41, already reported in [22] for signature 8. In Endometrial adenocarcinoma (Uterus-AdenoCA), all studies find a decrease in signatures 1 and 5, CloneSig and [21] find an increase in signatures 2 and 13, found varying in both directions in TrackSig [2].

Overall, we note a good agreement between the different analyses of the PCAWG cohort, with two lines of explanation of the few observed discrepancies: the unidentifiability of signature activity deconvolution, discussed in more details in the general discussion, and the point of view of the different analysis, with our analysis and [2] focusing on most varying signatures, with changes in signature activity quantified in absolute difference, while [21] focuses more on a general description of the cohort and report changes in log-fold change, and goes beyond SNVs by integrating Doublet-based substitutions (DBS) and Small insertions and deletions (ID) signature changes.

Figure S69: **Mutational signature changes in the PCAWG cohort.** Each plot corresponds to one cancer type, indicates the number of samples with a significant signature change compared to the total number of samples, and shows on the right panel an increase of a signature in the largest subclone, compared to clonal mutations, and on the left panel a decrease. The length of each bar corresponds to the number of patients with such changes, and the thickness to the median absolute observed change.

Figure S70: **Mutational signature changes in the PCAWG cohort.** Each plot corresponds to one cancer type, indicates the number of samples with a significant signature change compared to the total number of samples, and shows log fold change of signature activity in the largest subclone, compared to clonal mutations. The small dots correspond to individual samples, and the larger dot to the median change. Bootstrap sampling of the scores was used to compute 95% confidence intervals. The vertical dotted line in each plot represent no change.

Figure S71: Panel a: Stratification of patients depending on their pattern of signature change for Biliary-AdenoCA patients (30 patients, including 13 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S72: Panel a: Stratification of patients depending on their pattern of signature change for Bladder-TCC patients (22 patients, including 15 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S73: Panel a: Stratification of patients depending on their pattern of signature change for Bone-Leiomyo patients (27 patients, including 23 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S74: Panel a: Stratification of patients depending on their pattern of signature change for Bone-Osteosarc patients (37 patients, including 27 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S75: Panel a: Stratification of patients depending on their pattern of signature change for Bone-Other patients (26 patients, including 7 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S76: Panel a: Stratification of patients depending on their pattern of signature change for Breast patients (188 patients, including 131 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF). Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S77: Panel a: Stratification of patients depending on their pattern of signature change for CNS-GBM patients (29 patients, including 18 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S78: Panel a: Stratification of patients depending on their pattern of signature change for CNS-Medullo patients (145 patients, including 79 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S79: Panel a: Stratification of patients depending on their pattern of signature change for CNS-Oligo patients (18 patients, including 6 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S80: Panel a: Stratification of patients depending on their pattern of signature change for CNS-PiloAstro patients (89 patients, including 16 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S81: Panel a: Stratification of patients depending on their pattern of signature change for Cervix-SCC patients (18 patients, including 15 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S82: Panel a: Stratification of patients depending on their pattern of signature change for ColoRect-AdenoCA patients (56 patients, including 40 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S83: Panel a: Stratification of patients depending on their pattern of signature change for Eso-AdenoCA patients (97 patients, including 58 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S84: Panel a: Stratification of patients depending on their pattern of signature change for Head-SCC patients (52 patients, including 32 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S85: Panel a: Stratification of patients depending on their pattern of signature change for Kidney-ChRCC patients (45 patients, including 32 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S86: Panel a: Stratification of patients depending on their pattern of signature change for Kidney-PRCC patients (33 patients, including 25 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S87: Panel a: Stratification of patients depending on their pattern of signature change for Kidney-RCC patients (101 patients, including 68 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S88: Panel a: Stratification of patients depending on their pattern of signature change for Liver-HCC patients (319 patients, including 172 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S89: Panel a: Stratification of patients depending on their pattern of signature change for Lung-AdenoCA patients (31 patients, including 13 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S90: Panel a: Stratification of patients depending on their pattern of signature change for Lung-SCC patients (48 patients, including 19 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S91: Panel a: Stratification of patients depending on their pattern of signature change for Lymph-BNHL patients (101 patients, including 53 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S92: Panel a: Stratification of patients depending on their pattern of signature change for Lymph-CLL patients (92 patients, including 57 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S93: Panel a: Stratification of patients depending on their pattern of signature change for Myeloid-AML patients (16 patients, including 10 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S94: Panel a: Stratification of patients depending on their pattern of signature change for Myeloid-MDS/MPN patients (54 patients, including 31 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S95: Panel a: Stratification of patients depending on their pattern of signature change for Ovary-AdenoCA patients (103 patients, including 83 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S96: Panel a: Stratification of patients depending on their pattern of signature change for Panc-AdenoCA patients (233 patients, including 182 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S97: Panel a: Stratification of patients depending on their pattern of signature change for Panc-Endocrine patients (82 patients, including 68 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S98: Panel a: Stratification of patients depending on their pattern of signature change for Prost-AdenoCA patients (279 patients, including 178 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S99: Panel a: Stratification of patients depending on their pattern of signature change for Skin-Melanoma patients (99 patients, including 48 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S100: Panel a: Stratification of patients depending on their pattern of signature change for Stomach-AdenoCA patients (70 patients, including 31 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S101: Panel a: Stratification of patients depending on their pattern of signature change for Thy-AdenoCA patients (48 patients, including 22 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Figure S102: Panel a: Stratification of patients depending on their pattern of signature change for Uterus-AdenoCA patients (44 patients, including 35 with a significant signature change). The heatmap represents the difference between the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).Panel b: Stratification of patients depending on their complete pattern of signature exposure. The heatmap represents the signature activity in the largest subclone (in terms of number of mutations) and the clonal mutations (defined as belonging to the clone of highest CCF).

Supplementary figures

Figure S103: **CloneSig analysis of 1002 SNVs obtained by WES of a SARC sample (patient TCGA-3B-A9HT)**

Figure S104: **CloneSig analysis of 221 SNVs obtained by WES of a OV sample (patient TCGA-04-1341)**

Figure S105: **CloneSig analysis of 526 SNVs obtained by WES of a LUSC sample (patient TCGA-18-4721)**

Figure S106: **CloneSig analysis of 475 SNVs obtained by WES of a GBM sample (patient TCGA-19-2629)**

Figure S107: **CloneSig analysis of 407 SNVs obtained by WES of a OV sample (patient TCGA-20-0990)**

Figure S108: **CloneSig analysis of 460 SNVs obtained by WES of a OV sample (patient TCGA-20-1685)**

Figure S109: An example of empirical distribution of the total copy number for samples with 2000 mutations. In the first panel, labeled "default behavior", the user does not specify the percentage of genome that is diploid, and the total copy number values are drawn as specified in the Methods section. On the other panels, the user specifies a desired percentage of genome that is diploid (0, 0.2, 0.4, 0.6, 0.8, 1) respectively for the cases shown. The distribution is slightly different from the default behavior.

Figure S110: CloneSig's ability to distinguish 2 clones depending on the CCF distance between the two clones, and other relevant variables: number of mutations, percentage of diploid genome, sequencing depth, and the mode of input signature selection

Figure S111: Evaluation of the power of separation of two clones with varying CCF distance and mutation profile distance, for the different methods considered that account both for CCF mutations, and mutational context. Only simulations with 1000 observed signatures were considered, as TrackSig and TrackSigFreq hardly capture any signal with inferior numbers of signatures.

Figure S112: Influce of the number of observed signatures for each method's ability to distinguish clonal and subclonal mutations for simulations with 2 clones with a varying CCF difference.

Supplementary tables

Cancer type	Mean nb mutations (protected)	Mean nb mutations (public)	Standard deviation nb mutations (protected)	Standard deviation nb mutations (public)	Number of samples	Median followup (months)	Number of events
ACC	324.34	110.58	467.43	240.88	77	39.22	27
BLCA	732.18	350.38	886.65	424.23	354	17.21	153
BRCA	472.91	121.88	981.68	372.65	931	27.00	123
CESC	921.07	366.51	2869.13	1323.24	275	21.42	67
CHOL	358.97	100.97	505.33	225.34	35	12.65	15
COADREAD	2085.67	621.14	6247.55	1708.77	458	21.42	93
DLBC	568.30	203.68	276.72	123.15	37	24.67	5
ESCA	707.39	247.19	560.34	251.02	180	13.02	75
GBM	790.05	245.66	2583.80	1191.70	327	11.27	246
HNSC	454.21	201.73	543.69	271.16	445	20.96	185
KICH	209.32	50.12	264.68	142.42	60	85.66	8
KIRC	330.32	73.08	338.69	52.35	271	36.33	64
KIRP	280.86	82.57	130.50	37.90	242	25.13	37
LGG	212.73	76.89	1462.32	770.54	455	20.04	115
LIHC	511.53	157.00	445.93	174.27	347	19.25	117
LUAD	892.89	381.00	985.35	393.91	433	22.01	146
LUSC	909.66	382.83	686.09	289.58	423	22.27	169
MESO	203.23	47.08	114.19	48.74	78	NA	0
OV	593.79	160.15	562.30	152.24	390	31.34	227
PAAD	560.23	203.05	3780.28	1828.93	150	15.14	83
PCPG	78.70	14.09	16.71	7.43	141	NA	0
PRAD	171.30	61.78	824.76	467.50	458	30.80	8
SARC	424.27	130.72	723.66	309.04	210	30.96	81
SKCM	1876.97	886.84	2621.06	1204.90	423	35.28	184
STAD	989.96	455.07	1822.53	909.56	418	14.01	163
TGCT	148.88	22.16	33.90	11.64	128	43.05	3
THCA	120.98	16.22	95.88	14.59	467	31.01	12
THYM	253.69	36.40	175.18	83.19	121	39.17	8
UCEC	4647.22	1791.01	12341.73	4832.10	487	30.12	80
UCS	680.43	198.00	1743.77	738.28	53	NA	0
UVM	88.54	24.60	96.34	62.57	80	27.52	11

Table S3: Characteristics of the TCGA cohort used in this study.

Table S4: Table of presence of signatures in the different cancer types. A 1 indicates the presence of the signature, and a 0 an absence. The background color indicates whether the presence of the signature comes from SigProfiler detection in the TCGA (lavender) or from the literature (green), in which case, the matching paper is referenced in the table cell. Signatures suspected to be artefacts were removed and are indicated in dark pink.