
Scaling up Continuous-Time Markov Chains Helps Resolve Underspecification

Alkis Gotovos
MIT
alkisg@mit.edu

Rebekka Burkholz
Harvard University
rburkholz@hsph.harvard.edu

John Quackenbush
Harvard University
johnq@hsph.harvard.edu

Stefanie Jegelka
MIT
stefje@mit.edu

Abstract

Modeling the time evolution of discrete sets of items (e.g., genetic mutations) is a fundamental problem in many biomedical applications. We approach this problem through the lens of continuous-time Markov chains, and show that the resulting learning task is generally underspecified in the usual setting of cross-sectional data. We explore a perhaps surprising remedy: including a number of additional independent items can help determine time order, and hence resolve underspecification. This is in sharp contrast to the common practice of limiting the analysis to a small subset of relevant items, which is followed largely due to poor scaling of existing methods. To put our theoretical insight into practice, we develop an approximate likelihood maximization method for learning continuous-time Markov chains, which can scale to hundreds of items and is orders of magnitude faster than previous methods. We demonstrate the effectiveness of our approach on synthetic and real cancer data.

1 Introduction

Modeling the time evolution of physical processes with discrete states is an important machine learning problem that spans a wide range of application domains, including molecular dynamics (Crommelin & Vanden-Eijnden, 2006), phylogenetics (Suchard et al., 2001), and computational medicine (Liu et al., 2015). Continuous-time Markov chains have found success in such modeling tasks, supported by extensive research into learning and inference in such models (Opper & Sanguinetti, 2008; Perkins, 2009; Archambeau & Opper, 2011; Rao & Teh, 2012).

While the majority of previous work has focused on modeling a relatively small number of discrete states, many interesting applications involve the interaction between sets of items, which results in exponentially large state spaces. For example, when modeling the time progression of cancer, we want to consider a number of genetic alterations (e.g., mutations, copy number variations, etc.) that exhibit complex time-related dependencies (Beerenwinkel et al., 2014). To model how sets of n such alterations evolve over time, we need to consider a state space of size 2^n . What makes the problem even more challenging is the fact that the available data sets are typically cross-sectional; that is, they consist of (unordered) sets of items observed at unknown time points without any further information about the history of the underlying process.

In this paper, we focus on the problem of learning a particular parametric family of continuous-time Markov chains from such cross-sectional data, and show that the resulting problem is in general underspecified. The issue of underspecification is many-faceted and has been shown to permeate a

large number of machine learning systems, often leading to poor generalization, lack of robustness, and spurious relationships when interpreting the resulting models (D’Amour et al., 2020). In our setting, we explore a perhaps surprising remedy: including a number of additional (approximately) independent items can help determine the time order of process events, and hence resolve underspecification. We theoretically show that these extra items act as a “background clock”, since counting the number of their occurrences in a data sample can help us estimate the time at which that sample was observed. In sharp contrast to common practice, which limits the analysis to a small subset of relevant, highly-interacting items, our insight suggests that scaling up the learning procedure can be crucial for the robustness of the inferred models. We thus show that items deemed a priori unimportant to the application at hand may in fact be particularly valuable for recovering the time properties of the underlying physical process.

Existing learning approaches are, unfortunately, not well-suited for practically applying this insight; for example, the state-of-the-art method by Schill et al. (2019) for learning continuous-time Markov chains to model cancer progression, scales exponentially in the number of items considered, thus limiting the analysis to around 20 items. To alleviate this issue, we propose an approximate likelihood maximization method that relies on a fast gradient approximation. On problems of 20 items our approach runs almost 1000 times faster than the state of the art, while it can also scale to problems involving hundreds of items. In experiments on real cancer data, we demonstrate how some previous results may have been artifacts of underspecification, and how scaling up the analysis can result in more robust models of cancer progression.

2 Problem setup

Given a ground set $V = \{1, \dots, n\}$, we consider a continuous-time Markov chain (CTMC) $\{X(t)\}_{t \geq 0}$ on state space 2^V (Grimmett & Stirzaker, 2001). Thus, the states of the Markov chain are subsets of V , and can equivalently be identified as binary vectors of size n . CTMCs are commonly represented by a generator matrix $Q \in \mathbb{R}^{2^n \times 2^n}$. Since the rows and columns of this matrix correspond to states in 2^V , we will use $q_{S \rightarrow R}$ to denote the entry of Q at row indexed by S and column indexed by R . For $S \neq R$, the entry $q_{S \rightarrow R}$ represents the infinitesimal rate of transitioning from state S to state R , that is,

$$q_{S \rightarrow R} = \lim_{\delta t \rightarrow 0^+} \frac{\mathbb{P}(X_{t+\delta t} = R | X_t = S)}{\delta t}, \text{ for } S \neq R,$$

with $q_{S \rightarrow R} \geq 0$. The quantity $\tilde{q}_S := -q_{S \rightarrow S}$ represents the total infinitesimal rate of leaving state S , with $\tilde{q}_S \geq 0$. Since the total rate of leaving S is equal to the sum of the rates of transitioning to any other state R , we have $\tilde{q}_S = \sum_{R \in 2^V} q_{S \rightarrow R}$, for all $S \in 2^V$.

We are interested in modeling physical processes that induce an increasing accumulation of items over time, for example, the accumulation of genetic alterations when modeling cancer progression. To this end, we impose the following two constraints on the CTMC. First, we assume that the chain always starts from the empty state, that is, $X(0) = \emptyset$. Second, we only allow transitions that add a single item to the current set, that is, for $S \neq R$, $q_{S \rightarrow R} \neq 0$ only if $R = S \cup \{i\}$ for some $i \in V$. Since items are never removed, this implies that $\lim_{t \rightarrow \infty} X_t = V$.

In this paper, we focus on learning the CTMC from a given data set. Since learning directly the exponentially sized generator Q is a hopeless endeavor, we impose a particular parametric form on this matrix, recently introduced by Schill et al. (2019). We define a parameter matrix $\Theta \in \mathbb{R}^{n \times n}$ (also denoted by θ as a vector), and parameterize the off-diagonal entries of the generator matrix as

$$q_{S \rightarrow S \cup \{j\}}(\theta) = \exp \left(\theta_{jj} + \sum_{i \in S} \theta_{ij} \right),$$

for all $S \subseteq V$, and $j \in V \setminus S$. The off-diagonal elements of Θ encode positive (attractive) or negative (repulsive) pairwise interactions between items in V . The presence of item i increases ($\theta_{ij} > 0$) or decreases ($\theta_{ij} < 0$) the rate of adding item j by a multiplicative factor of $w_{ij} := e^{\theta_{ij}}$. The diagonal elements θ_{jj} encode the intrinsic rate of adding element j when ignoring all other interactions.

Obtaining a sample from a CTMC under all aforementioned constraints can be done using the following sequential procedure. At each step, given the current state S , we draw a “holding time” h

Algorithm 1: Sampling a set from the marginal CTMC

Input: Parameters θ Draw $t_{\text{obs}} \sim \text{Exp}(1)$ $t \leftarrow 0, S \leftarrow \emptyset$ **while** $t < t_{\text{obs}}$ **do** Draw $h \sim \text{Exp}(\tilde{q}_S(\theta))$ Compute $p_i = q_{S \rightarrow S \cup \{i\}}(\theta) / \tilde{q}_S(\theta)$, for $i \in V \setminus S$ Draw $x \sim \text{Cat}(V \setminus S, (p_i)_{i \in V \setminus S})$ $t \leftarrow t + h, S \leftarrow S \cup \{x\}$ **return** S

from an exponential distribution with parameter \tilde{q}_S . This represents how much time will pass until we add the next item. Then, we draw the next item $x \in V \setminus S$ according to a categorical distribution with probabilities $\propto q_{S \rightarrow S \cup \{x\}}$. The result is a sequence of items $(\sigma_1, \dots, \sigma_n)$, $\sigma_i \in V$, together with a sequence of holding times (h_1, \dots, h_n) , $h_i \in \mathbb{R}$.

Ideally, our data set would consist of a number of such item sequences and holding times; we would then proceed to learn the parameters θ by maximizing the likelihood of our CTMC model. Unfortunately, such detailed data is typically not available in practice. In particular, it is rather common in biomedical applications to only have access to cross-sectional data, i.e., to only observe the current state (in the form of an unordered subset of V) at a particular point in time without knowing how the process reached that state. Furthermore, it is often the case that the time point of each observation can only be roughly estimated or, worse, is completely unknown.

Similarly to previous work, we assume that we are given a data set $\mathcal{D} = \{S^{(1)}, \dots, S^{(N)}\}$, $S^{(i)} \subseteq V$, with each $S^{(i)}$ representing a state observed at a potentially different unknown time T_{obs} . More specifically, we assume that T_{obs} is a random variable following an exponential distribution with density $p(t) = e^{-t}$ (Gerstung et al., 2009; Schill et al., 2019), and investigate the problem of maximizing the marginal likelihood,

$$p(\mathcal{D}; \theta) = \prod_{i=1}^N p(S^{(i)}; \theta) = \prod_{i=1}^N \int_0^\infty p(S^{(i)} | t; \theta) p(t) dt. \quad (1)$$

Algorithm 1 shows how to sample a set from $p(S; \theta)$. The difference from our previous description of sampling from a CTMC lies in the fact that the procedure is only run until we reach the observation time t_{obs} , and the result is an unordered set without any information about the holding times.

3 The value of “unimportant” items

Our ultimate goal is to retrieve the ordered interaction structure present in the data by learning the parameter vector θ . Before attempting to maximize the marginal likelihood of eq. (1), though, it is natural to first ask whether the limited information available in the data is sufficient to infer such ordered interactions in the first place. We start with a simple example to showcase this issue.

3.1 Warmup example

Consider a ground set $V = \{1, 2\}$ and true model parameters $\theta_{11}^* = \theta_{21}^* = 0$, and $\theta_{12}^* = -\theta_{22}^* = \alpha$, for some $\alpha > 0$. When α is sufficiently large, the probability of item 2 occurring before item 1 goes to zero, therefore the model encodes the fact that item 1 is a prerequisite (and therefore, always appears before) item 2. As a consequence, we will also have $p(\{2\}; \theta^*) \approx 0$ for the resulting marginal distribution of sets. One may intuitively think that this could be enough information to infer the true time order; the following proposition shows that this is not the case.

Proposition 1. *There is a one-dimensional family of parameters $\theta = \theta(s)$, and $s_1, s_2 \in \mathbb{R}$, such that $p(\cdot; \theta(s)) = p(\cdot; \theta^*)$, for all $s_1 < s < s_2$.*

The proof follows by simple algebra, and can be found in Appendix A together with an illustration of the parameter family. Interestingly, for $s \rightarrow s_1$, the model $\theta(s)$ encodes the fact that, when we

observe both items, item 2 always appears before item 1, which is the opposite of what is encoded in our true model.

3.2 Independent items as a background clock

While the above example paints a pessimistic picture for inferring the correct time order, we show that this is still possible given some additional side information. Suppose that we are given another ground set V_+ whose items have no interaction with the items in V , that is, the parameter matrix Θ_{full} of the resulting model over $V_{\text{full}} = V \cup V_+$ has a block diagonal structure with blocks $\Theta \in \mathbb{R}^{|V| \times |V|}$ and $\Theta_+ \in \mathbb{R}^{|V_+| \times |V_+|}$. Then, it is easy to see that the distributions over V and V_+ are conditionally independent given time t . That is, for any $S \subseteq V$, $S_+ \subseteq V_+$, we can write $p(S \cup S_+ | t; \Theta_{\text{full}}) = p(S | t; \Theta) p(S_+ | t; \Theta_+)$. Note that the same statement does not hold when considering the marginal distributions, i.e., $p(S \cup S_+; \Theta_{\text{full}}) \neq p(S; \Theta) p(S_+; \Theta_+)$. Using this conditional independence property, we can rewrite the marginal probability of $S \cup S_+$ as follows:

$$\begin{aligned} p(S \cup S_+; \Theta_{\text{full}}) &= \int_0^\infty p(S \cup S_+ | t; \Theta_{\text{full}}) p(t) dt \\ &= \int_0^\infty p(S | t; \Theta) p(S_+ | t; \Theta_+) p(t) dt && \text{(by cond. ind.)} \\ &= \int_0^\infty p(S | t; \Theta) p(t | S_+; \Theta_+) p(S_+; \Theta_+) dt && \text{(by Bayes' rule)} \\ \Rightarrow p(S | S_+; \Theta_{\text{full}}) &= \int_0^\infty p(S | t; \Theta) p(t | S_+; \Theta_+) dt && \text{(dividing with } p(S_+; \Theta_+)). \end{aligned} \quad (2)$$

Comparing equations (1) and (2), we see that the information about S gained by observing S_+ can be explained via a posterior observation time distribution $p(t | S_+; \Theta_+)$ that refines the prior $p(t)$ by taking into account S_+ .

To gain further insight into this time posterior, we analyze in more detail a simplified setup, in which V_+ consists of m identically distributed independent items parameterized by θ_+ , that is, $\Theta_+ = \theta_+ \mathbf{I}_m$. In this case, $|S_+| | t$ follows a binomial distribution over m variables with success parameter $1 - e^{-w_+ t}$, where $w_+ := e^{\theta_+}$. Intuitively, we expect these independent items to act as a ‘‘background clock’’: counting the number of items $|S_+|$ should give us an indication about the observation time t . Moreover, as we increase the number of items m we should be able to get increasingly accurate estimates of the observation time. The following theorem formalizes these ideas.

Theorem 1. *Let $S_+ \subseteq V$ be randomly drawn according to the CTMC with parameter matrix Θ_+ , and let t^* be the true observation time of S_+ . Then, for any $\delta \in (0, 1)$, there exists an m_0 , such that for all $m \geq m_0$, the mean and variance of the posterior time distribution $p(t | S_+; \Theta)$ can be bounded as follows with probability at least $1 - \delta$,*

$$\begin{aligned} |M_{\text{post}} - t^*| &\leq C_1(\theta_+, t^*) \sqrt{\frac{\log m}{m}} + \mathcal{O}\left(\frac{\log m}{m}\right) \\ V_{\text{post}} &\leq C_2(\theta_+, t^*) \frac{1}{m} + \mathcal{O}\left(\frac{1}{m^2}\right). \end{aligned}$$

The quantities C_1 and C_2 are constant w.r.t. m and encode how suitable a particular rate (quantified through θ_+) of an item is for estimating the specific observation time t^* . Intuitively, larger rates are better suited for estimating smaller times and vice versa. We further illustrate this effect in [Appendix B](#) together with the detailed proof of the theorem. Note that for $m \rightarrow \infty$ we get $p(t | S_+; \Theta_+) \rightarrow \delta(t - t^*)$ in distribution, and therefore $p(S | S_+; \Theta_{\text{full}}) \rightarrow p(S | t^*; \Theta)$. This means that observing an arbitrarily large amount of independent items is equivalent to knowing the true observation time. In [Appendix C](#) we show that a similar asymptotic convergence behavior can be observed even when the items in V_+ are not identically distributed or exhibit some dependencies.

3.3 Discussion

To illustrate how this result applies to our two-item example discussed before, we define $\Theta_+ = \mathbf{I}_m[\theta_+^{(1)}, \dots, \theta_+^{(m)}]^T$, where each $\theta_+^{(i)}$ is drawn uniformly in $[-4, -2]$. This corresponds to individual




Figure 1: Posterior time densities for $N = 500$ samples. Posteriors of samples with $S = \{1\}$ are colored red, while those of samples with $S = \{1, 2\}$ are colored blue. The dashed line highlights the time posterior of a specific sample, while the black square denotes its true observation time.

frequencies of items in S_+ that roughly range from 0.01 to 0.1. We then draw $N = 500$ samples from the CTMC defined by θ_{full} . Figure 1 shows the posterior time density corresponding to each sample with $S = \{1\}$ (red) and $S = \{1, 2\}$ (blue), for different numbers m of independent items. To avoid clutter, we do not show the densities for samples with $S = \emptyset$. As expected, when $m = 0$ all samples follow the prior time density, and we have no way to distinguish their time order. However, as we keep increasing m we notice that the two colors keep separating from each other in time, and it becomes clear that the red posteriors tend to concentrate around earlier observation times than the blue ones. This is a clear indication that item 1 comes before item 2, and the additional time information obtained by the posterior time estimates should help pinpoint the true model within the family discussed in Proposition 1.

Due to computational considerations, it is common practice to constrain the ground set to a small subset of items that interact strongly and are deemed interesting to the application at hand; for example, genetic alterations that are known to have a role in cancer development (Raphael & Vandin, 2015; Schill et al., 2019). Our results in this section challenge this practice. We show that items believed to be independent of the “interesting items”, and thus considered unimportant to the analysis, can in fact be particularly valuable in estimating the observation time of data samples and recovering the time order properties of the underlying physical process.

For the remainder of this paper, we will not assume any prior knowledge about the block-diagonal structure of Θ_{full} or the values of the parameters θ_+ . Furthermore, we will make no distinction between “items of interest” and independent items. If there is indeed a block-diagonal parameter structure in the data, the addition of a regularization term that promotes sparsity can help recover this structure in the presence of noise. For this reason, similarly to what was used by Schill et al. (2019), the final objective we maximize in practice is the L_1 -regularized marginal log-likelihood,

$$F(\mathcal{D}; \theta) = \frac{1}{N} \sum_{d=1}^N \log p(S^{(d)}; \theta) - \lambda \sum_{i \neq j} |\theta_{ij}|. \quad (3)$$

4 Efficient approximate likelihood maximization

To put the previously discussed theoretical insight into practice, we need to be able to efficiently maximize the above objective for ground sets V containing potentially hundreds of items. When using a first-order method for optimization, the bulk of the required computation is devoted to obtaining the gradient of the marginal log-likelihood $\log p(S; \theta)$ with respect to the parameters θ . Computing the exact gradient requires $\mathcal{O}(2^n)$ computation; we propose here a method to compute a gradient approximation in a much more efficient manner. We start by deriving expressions for the likelihood in some simplified setups, and then use these results as building blocks for our setup of interest.

4.1 Full sequences

First, we consider a setup in which we observe the arrival order of the items in the CTMC, which corresponds to building up an ordered sequence σ instead of a set S in [Algorithm 1](#). Furthermore, we assume that $t_{\text{obs}} = +\infty$, which means that the returned sequence will be a permutation of the ground set V . By definition of the CTMC, the probability of observing a sequence $\sigma = (\sigma_1, \dots, \sigma_n) \in \mathcal{S}_V$, where \mathcal{S}_V denotes the permutation group over V , can be written as

$$p(\sigma; \theta) = \prod_{i=1}^n \mathbb{P}(X_i = \sigma_{[i]} \mid X_{i-1} = \sigma_{[i-1]}) = \prod_{i=1}^n \frac{q_{\sigma_{[i-1]} \rightarrow \sigma_{[i]}}(\theta)}{\tilde{q}_{\sigma_{[i-1]}}(\theta)}. \quad (4)$$

We denote subsequences using the bracket notation, $\sigma_{[i]} := (\sigma_1, \dots, \sigma_i)$, with $\sigma_{[0]} = ()$. Also, with a slight abuse of notation we directly use sequences in some places that require sets; therefore, the notation \tilde{q}_σ is to be understood as $\tilde{q}_{\text{set}(\sigma)}$, where $\text{set}(\sigma) := \{\sigma_i \mid i \in \{1, \dots, |\sigma|\}\}$.

4.2 Partial sequences given time

As a next step, instead of $t_{\text{obs}} = +\infty$, we assume that we are given the observation time t_{obs} , and would like to derive the probability of observing a partial sequence $\sigma = (\sigma_1, \dots, \sigma_k)$, with $k \leq n$, at time t_{obs} . We can express this as the probability of the intersection of two events in the space of outcomes of the continuous-time Markov chain. Event \mathcal{A} identifies the outcomes in which the first k ordered elements agree with our partial sequence σ . Event \mathcal{B} identifies the outcomes for which the observation time t_{obs} falls between the addition of the k -th and $k+1$ -th elements; that is, $T_k < t_{\text{obs}} < T_{k+1}$. We can then express the desired probability as $p(\sigma \mid t_{\text{obs}}; \theta) = \mathbb{P}(\mathcal{A}) \mathbb{P}(\mathcal{B} \mid \mathcal{A})$.

The probability of event \mathcal{A} can be directly written using (4), except that, in this case the product is taken from $i = 1$ to k , instead of n . Regarding event \mathcal{B} , first note that the sequence of events $(\{T_i < t_{\text{obs}} \mid \mathcal{A}\})_{i=1}^n$ is decreasing. As a result, we can write

$$\mathbb{P}(\mathcal{B} \mid \mathcal{A}) = \mathbb{P}(T_k < t_{\text{obs}} < T_{k+1} \mid \mathcal{A}) = \mathbb{P}(T_k < t_{\text{obs}} \mid \mathcal{A}) - \mathbb{P}(T_{k+1} < t_{\text{obs}} \mid \mathcal{A}). \quad (5)$$

By definition of the CTMC, each holding time H_i is an exponentially distributed random variable whose parameter depends on the $i-1$ items that have been added up to that point. We can also define the jump time $T_k := \sum_{i=1}^k H_i$, which represents the time at which the k -th change of state occurs. When we condition on knowing the first k items of the sequence, i.e., event \mathcal{A} , the jump times T_k and T_{k+1} are distributed as sums of k and $k+1$ independent exponential random variables respectively. The CDF of the sum of r independent exponential variables with rates $\lambda_1, \dots, \lambda_r$, also known as the hypoexponential distribution ([Bibinger, 2013](#)), is given by

$$F_{\text{HEXP}}(y; \lambda_1, \dots, \lambda_r) = \left(\prod_{i=1}^r \lambda_i \right) \sum_{i=1}^r \frac{1 - e^{-\lambda_i y}}{\lambda_i \prod_{j \neq i} (\lambda_j - \lambda_i)}.$$

In our case, the rates λ_i of each exponential distribution are $\tilde{q}_{[\sigma_i]}(\theta)$, thus we can compute the terms of (5) as $\mathbb{P}(\mathcal{B} \mid \mathcal{A}) = F_{\text{HEXP}}(t_{\text{obs}}; \tilde{q}_{[\sigma_1]}, \dots, \tilde{q}_{[\sigma_k]}) - F_{\text{HEXP}}(t_{\text{obs}}; \tilde{q}_{[\sigma_1]}, \dots, \tilde{q}_{[\sigma_{k+1}]})$.

4.3 Marginal partial sequences

Next we consider the probability of observing a partial sequence $\sigma = (\sigma_1, \dots, \sigma_k)$, with $k \leq n$, without knowledge of the observation time t_{obs} . We define events \mathcal{A} and \mathcal{B} as before, except that in

this case the observation time T_{obs} is a random variable, and event \mathcal{B} is defined by $T_k < T_{\text{obs}} < T_{k+1}$. In particular, we assume that T_{obs} is an exponential random variable with rate $\lambda_{\text{obs}} = 1$. Rather than naively integrating over time the probability derived before, we can obtain the following greatly simplified expression by making use of the memoryless property of exponential random variables (Bertsekas & Tsitsiklis, 2008) when computing $\mathbb{P}(\mathcal{B} \mid \mathcal{A})$.

Proposition 2. *The marginal probability of a partial sequence $\sigma = (\sigma_1, \dots, \sigma_k)$ can be written as*

$$p(\sigma; \theta) = \left(\prod_{i=1}^k \frac{q_{\sigma_{[i-1]} \rightarrow \sigma_{[i]}}(\theta)}{1 + \tilde{q}_{\sigma_{[i-1]}}(\theta)} \right) \frac{1}{1 + \tilde{q}_{\sigma_{[k]}}(\theta)}. \quad (6)$$

The proof can be found in Appendix D, but the form of this expression allows for an intuitive interpretation. We can modify the original continuous-time Markov chain by adding an extra dummy state X_{obs} , such that for all $X \subseteq V$, $q_{X \rightarrow X_{\text{obs}}}(\theta) = 1$, and $q_{X_{\text{obs}} \rightarrow X}(\theta) = 0$. The first property implies that there is a fixed rate 1 to transition from any state to X_{obs} , while the second property implies that X_{obs} is a terminal state. It is easy to see then, that eq. (6) expresses exactly the probability of observing a “full” sequence σ in this modified chain. Note that a “full” sequence in this chain is no longer a permutation of V , but rather any sequence that ends with X_{obs} .

4.4 Marginal sets

Finally, we consider the probability of observing a set $S \subseteq V$, without knowledge of the order in which the individual items arrived, and still under the assumption that the observation time T_{obs} is an exponential random variable with rate $\lambda_{\text{obs}} = 1$. Making use of the result in the previous section, we can compute this probability by summing over all partial sequences that are permutations of S , that is, $p(S; \theta) = \sum_{\sigma \in \mathcal{S}_S} p(\sigma; \theta)$. However, this computation is infeasible for all but very small set sizes, since it requires summing over $|S|!$ terms. To alleviate this problem, we propose a method to approximate the gradient of $\log p(S; \theta)$. This gradient can be written as

$$\nabla_{\theta} \log p(S; \theta) = \frac{1}{p(S; \theta)} \nabla_{\theta} p(S; \theta) = \frac{1}{p(S; \theta)} \sum_{\sigma \in \mathcal{S}_S} \nabla_{\theta} p(\sigma; \theta) = \sum_{\sigma \in \mathcal{S}_S} \frac{p(\sigma; \theta)}{p(S; \theta)} \nabla_{\theta} \log p(\sigma; \theta).$$

In the last expression, we observe that the gradient corresponding to each permutation σ is weighed by the probability of observing that permutation of the given set S . This suggests a stochastic approximation of the desired gradient by first sampling m permutations $\sigma^{(1)}, \dots, \sigma^{(M)}$ according to $p(\cdot \mid S; \theta) := p(\cdot; \theta)/p(S; \theta)$, and then computing the average of the resulting gradients,

$$\nabla_{\theta} \log p(S; \theta) \approx \frac{1}{M} \sum_{i=1}^M \nabla_{\theta} \log p(\sigma^{(i)}; \theta). \quad (7)$$

Conceptually similar approaches have been used to approximate the gradient of the log-normalizer when maximizing the likelihood of energy-based probabilistic models (Song & Kingma, 2021).

It remains to show how to obtain samples from $p(\cdot \mid S; \theta)$, which we do by using a Markov chain Monte Carlo method on \mathcal{S}_S . More concretely, we employ a Metropolis-Hastings chain (Levin et al., 2009); at each time step, given the current permutation σ , the chain proposes a new permutation σ_{new} according to proposal distribution $Q(\sigma_{\text{new}} \mid \sigma)$, and transitions to σ_{new} with probability

$$p_{\text{accept}} = \min \left(1, \frac{p(\sigma_{\text{new}} \mid S; \theta) Q(\sigma \mid \sigma_{\text{new}}; \theta)}{p(\sigma \mid S; \theta) Q(\sigma_{\text{new}} \mid \sigma; \theta)} \right).$$

A simple choice for Q is the uniform distribution over all permutations regardless of the current state. We instead employ a more sophisticated proposal that is based on the form of the marginal probability (6), and leads to faster and more stable learning in practice. While a detailed discussion and theoretical analysis of mixing is beyond the scope of this paper, in Appendix E we present our proposal in detail, and compare it to the uniform.

5 Further related work

There has been a long line of research focused on understanding and reconstructing the evolutionary history of tumors (Nik-Zainal et al., 2012; Welch et al., 2012; Turajlic et al., 2015; Mitchell




Figure 2: The KL divergence of recovered vs. true model as a function of the number m of extra independent items, evaluated on two synthetic examples with known true parameters. The dashed lines represent the learned models given the true observations times. Error bars are over 96 repetitions.

et al., 2018; Jolly & Van Loo, 2018). However, these papers do not attempt to systematically model the timed interaction of genetic alterations across a data set. More recently, Gerstung et al. (2020) showed that modern data resulting from deep sequencing procedures may contain partial order information about specific genetic events. Incorporating this additional information into an analysis like ours, could potentially help with further reducing underspecification. Finally, there is evidence that specific classes of mutational processes generate mutations at a constant rate, and are largely unaffected by other factors (Alexandrov et al., 2015; Campbell et al., 2020; Alexandrov et al., 2020). The resulting so-called mutational signatures have been likened to a biological clock, and could be candidates for explicitly defining the set V_+ discussed earlier (cf. Section 3.2).

6 Experiments

The following experimental setup is common to all our experiments. To optimize the objective (3), we use a proximal AdaGrad method (Duchi et al., 2011). We fix the initial step size to $\eta = 1$, and the regularization weight to $\lambda = 0.01$. To initialize the parameter matrix Θ , we train a diagonal model for 50 epochs, and then draw each off-diagonal entry from $\text{Unif}([-0.2, 0.2])$. For the gradient approximation (7), we use $M = 50$ samples in addition to 10 burn-in samples that are discarded. The code used to run our experiments can be found at <https://github.com/3lectrologos/time/tree/clean>.

6.1 Synthetic data

We start by evaluating our approach on two synthetic data sets drawn from known true models. The first data set is identical to the one discussed in Figure 1, and contains observations of two items that tend to follow a specific time order (item 1 comes before item 2). The second data set contains observations of five items following the graphical structure shown on the upper right of Figure 2. Edges with arrows \leftarrow represent positive interactions ($\theta_{ij} > 0$), while edges with vertical lines \vdash represent negative interactions ($\theta_{ij} < 0$). This type of structure, containing groups of items with bidirectional negative interactions, is common in genomics data; in fact, this exact structure can be found in the model learned from real cancer data in our following experiment (see Appendix F).

Our goal is to evaluate how well we can recover the true model by optimizing the regularized marginal likelihood; in particular, we want to examine the effect of adding a number m of independent items to the data. To quantify the quality of recovery, we compare the true distribution p_{true} of marginal sequences, computed using eq. (6), to the distribution of marginal sequences p^* of the learned model, approximated by drawing 10^6 marginal sequences from that model. In Figure 2 we depict the KL divergence $d_{\text{KL}}(p^* \| p_{\text{true}})$ for increasing numbers m of extra items. To compare against the ideal case, we also train a model by maximizing the regularized marginal likelihood given the true observation times (cf. Section 4.2). The KL divergence between this model and the true one is shown as a dashed line in the figure. We repeated each experiment 96 times; sources of randomness include the choice of independent items, the parameter initialization, and the gradient approximation. We see that the recovery quality improves for both data sets with increasing m , and rapidly approaches the ideal. This shows that observing a large enough number of additional items is practically equivalent to knowing the true observation times, and thus validates Theorem 1.

Table 1: Learning run times (Intel Core i9 CPU)

Method	$n = 10$	$n = 15$	$n = 20$	$n = 50$	$n = 100$
Schill et al. (2019)	2 s	43 s	121 m	–	–
Ours	3 s	4 s	8 s	1 m 5 s	33 m 43 s

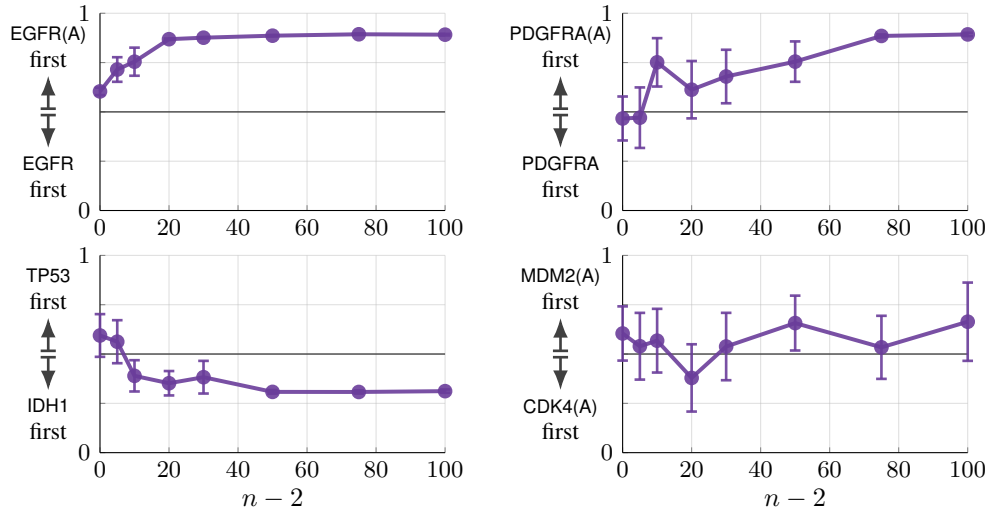


Figure 3: Results on the TCGA glioblastoma data set indicating the tendencies in the inferred time order of four selected pairs of genetic alterations. Error bars are over 24 repetitions.

6.2 Real cancer data

We next evaluate our approach on a data set that contains $N = 378$ tumor samples of glioblastoma multiforme, an aggressive type of brain cancer. The data is part of the TCGA PanCancer Atlas project (<https://www.cancer.gov/tcga>), and we obtained a preprocessed version via cBioPortal (Cerami et al., 2012). We followed some further filtering procedures from previous work (Leiserson et al., 2013) and ended up with a ground set of $n = 410$ items, each of which represents a point mutation, amplification, or deletion of a specific gene. Analyzing the interaction structure between such genetic alterations is of fundamental importance to cancer research, as it can help illuminate the processes that are involved in cancer initiation and progression.

For the remainder of this section, we assume that the items in V are ordered by decreasing frequency, and when selecting a subset we keep the most frequent items. In Table 1, we show that our learning approach runs about 1000 times faster on a ground set of size $n = 20$ compared to the exponentially scaling exact approach proposed by Schill et al. (2019).

In the absence of ground truth, we use the following procedure to evaluate how some of the learned parameters behave as we increase the size of the ground set. Given a pair of items $a, b \in V$, and $n-2$ other (most frequent) items, we first learn a CTMC model on the n total items using our approach. Next we approximate the sequence distribution of just a and b (marginally over time and over all other items) by sampling 10^6 marginal sequences from the learned model. Finally, we compute the proportion of sequences in which a occurred before b , given that both were observed; when this proportion is greater than 0.5, we infer that a is more likely to occur before b , and vice versa. In Figure 3 we present these computed proportions for four chosen pairs of genetic alterations that are interesting both from a biological standpoint, as well as in terms of learning behavior.

First, we discover two significant ordered interactions, shown in the top row of the figure, which were not reported in previous work analyzing glioblastoma data (Raphael & Vandin, 2015; Cristea et al., 2017; Schill et al., 2019). Our results suggest that amplification events for genes EGFR and PDGFRA tend to occur before their respective point mutation events. Interestingly, the co-occurrence of amplifications and mutations for EGFR has been observed before (Leiserson et al.,

2013; Sanchez-Vega et al., 2018), but there is little known about their time order. Note that the time order of the PDGFRA alterations can be robustly inferred only after including more than 70 additional items in the analysis.

Second, we confirm some previously reported interactions, for example, the tendency of IDH1 mutations to occur before TP53 ones, as shown in the bottom left of the figure. This was observed before by Schill et al. (2019), and is also supported by some biological evidence (Watanabe et al., 2009). Again, we can see that our result becomes robust only after $n = 50$. This indicates that the previous observation of this time order may have been in part due to a fortuitous choice of the optimization setup, e.g., fixed initialization (see Appendix G). The final result in the bottom right of the figure seems to support this claim: the amplifications of MDM2 and CDK4 do not seem to have a clear time order, even though previous work has inferred both that CDK4(A) occurs before MDM2(A) (Cristea et al., 2017), as well as the opposite order (Schill et al., 2019).

Finally, the parameter matrix Θ learned from this data set (see Appendix G) has indeed an approximately block-diagonal structure, which confirms the practical relevance of Theorem 1.

7 Conclusion

Our experiments demonstrated that, in the presence of underspecification, special care is required when interpreting results that may be significantly affected by particular choices in the optimization setup. In this paper, we proposed an effective way to mitigate such effects, namely scaling up the problem and taking into account additional items that are commonly discarded as insignificant. We hope that our approach will inspire further work into combating underspecification and learning robust time evolution models.

Acknowledgments and Disclosure of Funding

This work was partially supported by a Swiss National Science Foundation Early Postdoc.Mobility fellowship, by NSF CAREER award 1553284, NSF BIGDATA IIS-1741341, and NSF award 1717610, as well as a grant from the US National Cancer Institute (1R35CA220523).

References

- Alexandrov, L. B., Jones, P. H., Wedge, D. C., Sale, J. E., Campbell, P. J., Nik-Zainal, S., and Stratton, M. R. Clock-like mutational processes in human somatic cells. *Nature Genetics*, 2015.
- Alexandrov, L. B., Kim, J., Haradhvala, N. J., Huang, M. N., Tian Ng, A. W., Wu, Y., Boot, A., Covington, K. R., Gordenin, D. A., Bergstrom, E. N., Islam, S. M. A., Lopez-Bigas, N., Klimczak, L. J., McPherson, J. R., Morganella, S., Sabarinathan, R., Wheeler, D. A., Mustonen, V., Getz, G., Rozen, S. G., and Stratton, M. R. The repertoire of mutational signatures in human cancer. *Nature*, 2020.
- Archambeau, C. and Opper, M. Approximate inference for continuous-time markov processes. In Barber, D., Cemgil, A. T., and Chiappa, S. (eds.), *Bayesian Time Series Models*. 2011.
- Beerenwinkel, N., Schwarz, R. F., Gerstung, M., and Markowitz, F. Cancer evolution: Mathematical models and computational inference. *Systematic Biology*, 2014.
- Bertsekas, D. P. and Tsitsiklis, J. N. *Introduction to probability*. Athena Scientific, 2008.
- Bibinger, M. Notes on the sum and maximum of independent exponentially distributed random variables with different scale parameters. *arXiv*, art. 1307.3945, 2013.
- Campbell, P. J., Getz, G., Korbel, J. O., Stuart, J. M., Jennings, J. L., Stein, L. D., Perry, M. D., Nahal-Bose, H. K., Ouellette, B. F. F., Li, C. H., Rheinbay, E., Nielsen, G. P., Sgroi, D. C., Wu, C.-L., Faquin, W. C., Deshpande, V., Boutros, P. C., Lazar, A. J., Hoadley, K. A., Louis, D. N., Dursi, L. J., Yung, C. K., Bailey, M. H., Saksena, G., Raine, K. M., Buchhalter, I., Kleinheinz, K., Schlesner, M., Zhang, J., Wang, W., Wheeler, D. A., Ding, L., Simpson, J. T., OConnor, B. D., Yakneen, S., Ellrott, K., Miyoshi, N., Butler, A. P., Royo, R., Shorser, S. I., Vazquez, M.,

- Rausch, T., Tiao, G., Waszak, S. M., Rodriguez-Martin, B., Shringarpure, S., Wu, D.-Y., Demidov, G. M., Delaneau, O., Hayashi, S., Imoto, S., Habermann, N., Segre, A. V., Garrison, E., Cafferkey, A., Alvarez, E. G., Heredia-Genestar, J. M., Muyas, F., Drechsel, O., Bruzos, A. L., Temes, J., Zamora, J., Baez-Ortega, A., Kim, H.-L., Mashl, R. J., Ye, K., DiBiase, A., Huang, K.-I., Letunic, I., McLellan, M. D., Newhouse, S. J., Shmaya, T., Kumar, S., Wedge, D. C., Wright, M. H., Yellapantula, V. D., Gerstein, M., Khurana, E., Marques-Bonet, T., Navarro, A., Bustamante, C. D., Siebert, R., Nakagawa, H., Easton, D. F., Ossowski, S., Tubio, J. M. C., De La Vega, F. M., Estivill, X., Yuen, D., Mihaiescu, G. L., Omberg, L., Ferretti, V., Sabarinathan, R., Pich, O., Gonzalez-Perez, A., Taylor-Weiner, A., Fittall, M. W., Demeulemeester, J., Tarabichi, M., Roberts, N. D., Van Loo, P., Cortés-Ciriano, I., Urban, L., Park, P., Zhu, B., Pitkänen, E., Li, Y., Saini, N., Klimczak, L. J., Weischenfeldt, J., Sidiropoulos, N., Alexandrov, L. B., Rabionet, R., Escaramis, G., Bosio, M., Holik, A. Z., Susak, H., Prasad, A., Erkek, S., Calabrese, C., Raeder, B., Harrington, E., Mayes, S., Turner, D., Juul, S., Roberts, S. A., Song, L., Koster, R., Mirabello, L., Hua, X., Tanskanen, T. J., Tojo, M., Chen, J., Aaltonen, L. A., Rättsch, G., Schwarz, R. F., Butte, A. J., Brazma, A., Chanock, S. J., Chatterjee, N., Stegle, O., Harismendy, O., Bova, G. S., Gordenin, D. A., Haan, D., Sieverling, L., Feuerbach, L., Chalmers, D., Joly, Y., Knoppers, B., Molnár-Gábor, F., Phillips, M., Thorogood, A., Townend, D., Goldman, M., Fonseca, N. A., Xiang, Q., Craft, B., Piñeiro-Yañez, E., Muñoz, A., Petryszak, R., Füllgrabe, A., Al-Shahrour, F., Keays, M., Haussler, D., Weinstein, J., Huber, W., Valencia, A., Papatheodorou, I., Zhu, J., Fan, Y., Torrents, D., Bieg, M., Chen, K., Chong, Z., Cibulskis, K., Eils, R., Fulton, R. S., Gelpi, J. L., Gonzalez, S., Gut, I. G., Hach, F., Heinold, M., Hu, T., Huang, V., Hutter, B., Jäger, N., Jung, J., Kumar, Y., Lalasingh, C., Leshchiner, I., Livitz, D., Ma, E. Z., Maruvka, Y. E., Milovanovic, A., Nielsen, M. M., Paramasivam, N., Pedersen, J. S., Puiggròs, M., Sahinalp, S. C., Sarrafi, I., Stewart, C., Stobbe, M. D., Wala, J. A., Wang, J., Wendl, M., Werner, J., Wu, Z., Xue, H., Yamaguchi, T. N., Yellapantula, V., Davis-Dusenbery, B. N., Grossman, R. L., Kim, Y., Heinold, M. C., Hinton, J., Jones, D. R., Menzies, A., Stebbings, L., Hess, J. M., Rosenberg, M., Dunford, A. J., Gupta, M., Imielinski, M., Meyerson, M., Beroukhi, R., Reimand, J., Dhingra, P., Favero, F., Dentre, S., Wintersinger, J., Rudneva, V., Park, J. W., Hong, E. P., Heo, S. G., Kahles, A., Lehmann, K.-V., Soulette, C. M., Shiraishi, Y., Liu, F., He, Y., Demirciolu, D., Davidson, N. R., Greger, L., Li, S., Liu, D., Stark, S. G., Zhang, F., Amin, S. B., Bailey, P., Chateigner, A., Frenkel-Morgenstern, M., Hou, Y., Huska, M. R., Kilpinen, H., Lamaze, F. C., Li, C., Li, X., Li, X., Liu, X., Marin, M. G., Markowski, J., Nandi, T., Ojesina, A. I., Pan-Hammarström, Q., Park, P. J., Pedamallu, C. S., Su, H., Tan, P., Teh, B. T., Wang, J., Xiong, H., Ye, C., Yung, C., Zhang, X., Zheng, L., Zhu, S., Awadalla, P., Creighton, C. J., Wu, K., Yang, H., Göke, J., Zhang, Z., Brooks, A. N., Fittall, M. W., Martincorena, I., Rubio-Perez, C., Juul, M., Schumacher, S., Shapira, O., Tamborero, D., Mularoni, L., Hornshøj, H., Deu-Pons, J., Muiños, F., Bertl, J., Guo, Q., Gonzalez-Perez, A., Xiang, Q., and The ICGC/TCGA Pan-Cancer Analysis of Whole Genomes Consortium. Pan-cancer analysis of whole genomes. *Nature*, 2020.
- Cerami, E., Gao, J., Dogrusoz, U., Gross, B. E., Sumer, S. O., Aksoy, B. A., Jacobsen, A., Byrne, C. J., Heuer, M. L., Larsson, E., Antipin, Y., Reva, B., Goldberg, A. P., Sander, C., and Schultz, N. The cBio cancer genomics portal: An open platform for exploring multidimensional cancer genomics data. *Cancer Discovery*, 2012.
- Cristea, S., Kuipers, J., and Beerenwinkel, N. pathTiMEX: Joint inference of mutually exclusive cancer pathways and their progression dynamics. *Journal of Computational Biology*, 2017.
- Crommelin, D. T. and Vanden-Eijnden, E. Fitting timeseries by continuous-time markov chains: A quadratic programming approach. *Journal of Computational Physics*, 2006.
- D’Amour, A., Heller, K., Moldovan, D., Adlam, B., Alipanahi, B., Beutel, A., Chen, C., Deaton, J., Eisenstein, J., Hoffman, M. D., Hormozdiari, F., Houlsby, N., Hou, S., Jerfel, G., Karthikesalingam, A., Lucic, M., Ma, Y., McLean, C., Mincu, D., Mitani, A., Montanari, A., Nado, Z., Natarajan, V., Nielson, C., Osborne, T. F., Raman, R., Ramasamy, K., Sayres, R., Schrouff, J., Seneviratne, M., Sequeira, S., Suresh, H., Veitch, V., Vladymyrov, M., Wang, X., Webster, K., Yadlowsky, S., Yun, T., Zhai, X., and Sculley, D. Underspecification presents challenges for credibility in modern machine learning. *arXiv*, art. 2011.03395, 2020.
- Duchi, J., Hazan, E., and Singer, Y. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research (JMLR)*, 2011.

- Gerstung, M., Baudis, M., Moch, H., and Beerenwinkel, N. Quantifying cancer progression with conjunctive bayesian networks. *Bioinformatics*, 2009.
- Gerstung, M., Jolly, C., Leshchiner, I., Dentre, S. C., Gonzalez, S., Rosebrock, D., Mitchell, T. J., Rubanova, Y., Anur, P., Yu, K., Tarabichi, M., Deshwar, A., Wintersinger, J., Kleinheinz, K., Vázquez-García, I., Haase, K., Jerman, L., Sengupta, S., Macintyre, G., Malikić, S., Donmez, N., Livitz, D. G., Cmero, M., Demeulemeester, J., Schumacher, S., Fan, Y., Yao, X., Lee, J., Schlesner, M., Boutros, P. C., Bowtell, D. D., Zhu, H., Getz, G., Imielinski, M., Beroukhim, R., Sahinalp, S. C., Ji, Y., Peifer, M., Markowitz, F., Mustonen, V., Yuan, K., Wang, W., Morris, Q. D., Spellman, P. T., Wedge, D. C., and Loo, P. V. The evolutionary history of 2,658 cancers. *Nature*, 2020.
- Grimmett, G. and Stirzaker, D. *Probability and random processes*. Oxford University Press, 2001.
- Jolly, C. and Van Loo, P. Timing somatic events in the evolution of cancer. *Genome Biology*, 2018.
- Leiserson, M. D. M., Blokh, D., Sharan, R., and Raphael, B. J. Simultaneous identification of multiple driver pathways in cancer. *PLoS Computational Biology*, 2013.
- Levin, D. A., Peres, Y., and Wilmer, E. L. *Markov chains and mixing times*. American Mathematical Society, 2009.
- Liu, Y.-Y., Li, S., Li, F., Song, L., and Rehg, J. M. Efficient learning of continuous-time hidden markov models for disease progression. In *Neural Information Processing Systems (NeurIPS)*, 2015.
- Mengersen, K. L. and Tweedie, R. L. Rates of convergence of the Hastings and Metropolis algorithms. *The Annals of Statistics*, 1996.
- Mitchell, T. J., Turajlic, S., Rowan, A., Nicol, D., Farmery, J. H. R., O’Brien, T., Martincorena, I., Tarpey, P., Angelopoulos, N., Yates, L. R., Butler, A. P., Raine, K., Stewart, G. D., Challacombe, B., Fernando, A., Lopez, J. I., Hazell, S., Chandra, A., Chowdhury, S., Rudman, S., Soultati, A., Stamp, G., Fotiadis, N., Pickering, L., Au, L., Spain, L., Lynch, J., Stares, M., Teague, J., Maura, F., Wedge, D. C., Horswell, S., Chambers, T., Litchfield, K., Xu, H., Stewart, A., Elaidi, R., Oudard, S., McGranahan, N., Csabai, I., Gore, M., Futreal, P. A., Larkin, J., Lynch, A. G., Szallasi, Z., Swanton, C., Campbell, P. J., and TRACERx Renal Consortium. Timing the landmark in the evolution of clear cell renal cell cancer: TRACERx renal. *Cell*, 2018.
- Nik-Zainal, S., VanĀLoo, 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., Hinton, J., Menzies, A., Stebbings, L. A., Leroy, C., Jia, M., Rance, R., Mudie, L. J., Gamble, S. J., Stephens, P. J., McLaren, S., Tarpey, P. S., Papaemmanuil, E., Davies, H. R., Varela, I., McBride, D. J., Bignell, G. R., Leung, K., Butler, A. P., Teague, J. W., Martin, S., Jönsson, G., Mariani, O., Boyault, S., Miron, P., Fatima, A., Langerød, A., Aparicio, S. A. J. R., Tutt, A., Sieuwerts, A. M., Borg, k., Thomas, G., Salomon, A. V., Richardson, A. L., Børresen-Dale, A.-L., Futreal, P. A., Stratton, M. R., and Campbell, P. J. The life history of 21 breast cancers. *Cell*, 2012.
- Opper, M. and Sanguinetti, G. Variational inference for markov jump processes. In *Neural Information Processing Systems (NeurIPS)*, 2008.
- Perkins, T. Maximum likelihood trajectories for continuous-time markov chains. In *Neural Information Processing Systems (NeurIPS)*, 2009.
- Rao, V. and Teh, Y. Mcmc for continuous-time discrete-state systems. In *Neural Information Processing Systems (NeurIPS)*, 2012.
- Raphael, B. J. and Vandin, F. Simultaneous inference of cancer pathways and tumor progression from cross-sectional mutation data. *Journal of Computational Biology*, 2015.
- Sanchez-Vega, F., Mina, M., and more. Oncogenic signaling pathways in the cancer genome atlas. *Cell*, 2018.
- Schill, R., Solbrig, S., Wettig, T., and Spang, R. Modelling cancer progression using mutual hazard networks. *Bioinformatics*, 2019.

- Song, Y. and Kingma, D. P. How to train your energy-based models. *arXiv*, art. 2101.03288, 2021.
- Suchard, M. A., Weiss, R. E., and Sinsheimer, J. S. Bayesian selection of continuous-time markov chain evolutionary models. *Molecular Biology and Evolution*, 2001.
- Turajlic, S., McGranahan, N., and Swanton, C. Inferring mutational timing and reconstructing tumour evolutionary histories. *Biochimica et Biophysica Acta (BBA) - Reviews on Cancer*, 2015.
- Watanabe, T., Nobusawa, S., Kleihues, P., and Ohgaki, H. IDH1 mutations are early events in the development of astrocytomas and oligodendrogliomas. *The American Journal of Pathology*, 2009.
- Welch, J. S., Ley, T. J., Link, D. C., Miller, C. A., Larson, D. E., Koboldt, D. C., Wartman, L. D., Lamprecht, T. L., Liu, F., Xia, J., Kandoth, C., Fulton, R. S., McLellan, M. D., Dooling, D. J., Wallis, J. W., Chen, K., Harris, C. C., Schmidt, H. K., Kalicki-Veizer, J. M., Lu, C., Zhang, Q., Lin, L., O’Laughlin, M. D., McMichael, J. F., Delehaunty, K. D., Fulton, L. A., Magrini, V. J., McGrath, S. D., Demeter, R. T., Vickery, T. L., Hundal, J., Cook, L. L., Swift, G. W., Reed, J. P., Alldredge, P. A., Wylie, T. N., Walker, J. R., Watson, M. A., Heath, S. E., Shannon, W. D., Varghese, N., Nagarajan, R., Payton, J. E., Baty, J. D., Kulkarni, S., Klco, J. M., Tomasson, M. H., Westervelt, P., Walter, M. J., Graubert, T. A., DiPersio, J. F., Ding, L., Mardis, E. R., and Wilson, R. K. The origin and evolution of mutations in acute myeloid leukemia. *Cell*, 2012.

A Proof of Proposition 1

Proposition 1. *There is a one-dimensional family of parameters $\boldsymbol{\theta} = \boldsymbol{\theta}(s)$, and $s_1, s_2 \in \mathbb{R}$, such that $p(\cdot; \boldsymbol{\theta}(s)) = p(\cdot; \boldsymbol{\theta}^*)$, for all $s \in (s_1, s_2)$.*

Proof. We use (6) to write down the following probabilities,

$$\begin{aligned} p_0(\boldsymbol{\theta}) &:= p(\emptyset; \boldsymbol{\theta}) = p((); \boldsymbol{\theta}) = \frac{1}{1 + w_{11} + w_{22}} \\ p_1(\boldsymbol{\theta}) &:= p(\{1\}; \boldsymbol{\theta}) = p((1); \boldsymbol{\theta}) = \frac{w_{11}}{1 + w_{11} + w_{22}} \frac{1}{1 + w_{22}w_{12}} = \frac{p_0 w_{11}}{1 + w_{22}w_{12}} \\ p_2(\boldsymbol{\theta}) &:= p(\{2\}; \boldsymbol{\theta}) = p((2); \boldsymbol{\theta}) = \frac{w_{22}}{1 + w_{11} + w_{22}} \frac{1}{1 + w_{11}w_{21}} = \frac{p_0 w_{22}}{1 + w_{11}w_{21}}. \end{aligned}$$

We define the probabilities of the true model $p_i^* := p_i(\boldsymbol{\theta}^*)$, for $i = 0, 1, 2$, and solve the system

$$\begin{aligned} p_i(\boldsymbol{\theta}) &= p_i^*, \quad i \in \{0, 1, 2\} \\ w_{ij} &> 0, \quad i, j \in \{1, 2\}. \end{aligned}$$

The solution is the following parametric family,

$$\begin{aligned} w_{11} &= s \\ w_{22} &= 1/p_0^* - 1 - s \\ w_{21} &= \frac{1 - p_0^* - p_2^* - p_0^* s}{p_2^* s} \\ w_{12} &= \frac{p_0^* s - p_1^*}{p_1^* (1/p_0^* - 1 - s)}, \end{aligned}$$

for $s \in (s_1, s_2)$, where $s_1 = p_1^*/p_0^*$ and $s_2 = (1 - p_0^* - p_2^*)/p_0^*$. Then, the parameter family $\boldsymbol{\theta} = \boldsymbol{\theta}(s)$ can be computed as $\theta_{ij}(s) = \log w_{ij}(s)$, for $i, j \in \{1, 2\}$. \square

Illustration of the parameter family




Figure 4: The family of parameters that result in the same marginal distribution of sets $p(S; \boldsymbol{\theta}^*)$ for the two-item example discussed in Section 3 and for $\alpha = 4$. The square marks highlight the true model parameters $\theta_{11}^* = 0, \theta_{22}^* = -\alpha, \theta_{12}^* = \alpha, \theta_{21}^* = 0$.

B Proof of Theorem 1

We consider a CTMC with parameter matrix $\Theta_+ = \theta_+ \mathbf{I}_m$. As a reminder, for simplicity of notation we define $w_+ = e^{\theta_+}$. The following lemma derives the mean and variance of the posterior time distribution, for any fixed observation S_+ .

Lemma 1. *For any $m > 0$, and any observation $S_+ \subseteq V$ with $|S_+| = k$, if we define $r = k/m$, then the posterior time distribution $p(t | S_+; \theta_+)$ has mean and variance given by*

$$\begin{aligned} M_{\text{post}} &= \frac{1}{w_+} (\psi(\alpha + \beta) - \psi(\alpha)) \\ V_{\text{post}} &= \frac{1}{w_+^2} (\psi_1(\alpha) - \psi_1(\alpha + \beta)), \end{aligned}$$

where $\alpha = 1/w_+ + (1-r)m + 1$, $\beta = rm + 1$, and ψ, ψ_1 are the digamma and trigamma functions respectively.

Proof. At any time $t \geq 0$, and for any $x \in V$, if we denote by X the randomly observed set, then the indicator random variable $\mathbb{1}[x \in X]$ is Bernoulli with success probability $1 - e^{-tw_+}$. Therefore, the variable $|X|$ follows a binomial distribution with parameters m and $1 - e^{-tw_+}$, and the posterior time density can be written as

$$\begin{aligned} p(t | S_+; \theta_+) &= p(t | |X| = k; \theta_+) \\ &\propto p(|X| = k | t; \theta_+) p(t) \\ &= (1 - e^{-tw_+})^k (e^{-tw_+})^{m-k} e^{-t} \\ &= (1 - e^{-tw_+})^{rm} e^{-t(1+(1-r)w_+m)}. \end{aligned}$$

Consider now the variable transformation $y = e^{-tw_+}$ with derivative $dy/dt = -w_+ e^{-tw_+}$. If T is a random variable with density $p(\cdot | S_+; \theta_+)$, then the random variable $Y = e^{-Tw_+}$ has density given by

$$p(y | S_+; \theta) \propto y^{1/w_+ + (1-r)m} (1 - y)^{rm},$$

that is, $Y \sim \text{Beta}(\alpha, \beta)$, where $\alpha = 1/w_+ + (1-r)m + 1$, and $\beta = rm + 1$. We can then write the posterior time mean and variance as

$$\begin{aligned} M_{\text{post}} &= \mathbb{E}[T] = -\frac{1}{w_+} \mathbb{E}[\log Y] \\ V_{\text{post}} &= \text{Var}[T] = \frac{1}{w_+^2} \text{Var}[\log Y]. \end{aligned}$$

The mean and variance of $\log Y$ can be computed with the help of the digamma function ψ , and trigamma function ψ_1 , which are defined as

$$\begin{aligned} \psi(z) &= \frac{d}{dz} \log \Gamma(z) \\ \psi_1(z) &= \frac{d^2}{dz^2} \log \Gamma(z), \end{aligned}$$

where Γ denotes the gamma function. We then have

$$\begin{aligned} M_{\text{post}} &= -\frac{1}{w_+} \mathbb{E}[\log Y] = \frac{1}{w_+} (\psi(\alpha + \beta) - \psi(\alpha)) \\ V_{\text{post}} &= \frac{1}{w_+^2} \text{Var}[\log Y] = \frac{1}{w_+^2} (\psi_1(\alpha) - \psi_1(\alpha + \beta)). \end{aligned}$$

□

Since the size k of the observed set is randomly distributed given the observation time, the next lemma shows how this size concentrates around the mean.

Lemma 2. Let $S_+ \subseteq V$ be randomly drawn according to the CTMC with parameter matrix Θ_+ , let t^* be the true observation time, and define random variable $r = |S_+|/m$. Then, for any $\delta \in (0, 1)$, and any $m \geq \sqrt{2/\delta}$, the following holds with probability at least $1 - \delta$,

$$\left| r - (1 - e^{-w_+ t^*}) \right| \leq \sqrt{\frac{\log m}{m}}.$$

Proof. In the preceding lemma we saw that $|S_+|$ follows a binomial distribution with parameters m and $1 - e^{-w_+ t^*}$. Applying Hoeffding's inequality gives us

$$\begin{aligned} & \mathbb{P} \left[\left| |S_+| - m(1 - e^{-w_+ t^*}) \right| \leq \sqrt{m \log m} \right] \geq 1 - \frac{2}{m^2} \\ \Rightarrow & \mathbb{P} \left[\left| r - (1 - e^{-w_+ t^*}) \right| \leq \sqrt{\frac{\log m}{m}} \right] \geq 1 - \frac{2}{m^2}. \end{aligned}$$

□

Theorem 1. Let $S_+ \subseteq V$ be randomly drawn according to the CTMC with parameter matrix Θ_+ , and let t^* be the true observation time of S_+ . Then, for any $\delta \in (0, 1)$, there exists a m_0 , such that for all $m \geq m_0$, the mean and variance of the posterior time distribution $p(t | S_+; \theta)$ can be bounded as follows with probability at least $1 - \delta$,

$$\begin{aligned} |M_{\text{post}} - t^*| & \leq C_1(\theta_+, t^*) \sqrt{\frac{\log m}{m}} + \mathcal{O}\left(\frac{\log m}{m}\right) \\ V_{\text{post}} & \leq C_2(\theta_+, t^*) \frac{1}{m} + \mathcal{O}\left(\frac{1}{m^2}\right). \end{aligned}$$

Proof. We start with the posterior mean. By Lemma 1 we have

$$\begin{aligned} M_{\text{post}} & = \frac{1}{w_+} (\psi(\alpha + \beta) - \psi(\alpha)) \\ & = \frac{1}{w_+} \left(\log(\alpha + \beta) - \frac{1}{2(\alpha + \beta)} - \log(\alpha) + \frac{1}{2\alpha} + \mathcal{O}\left(\frac{1}{(\alpha + \beta)^2}\right) + \mathcal{O}\left(\frac{1}{\alpha^2}\right) \right) \\ & \hspace{15em} \text{(by } \psi \text{ series expansion)} \\ & = \frac{1}{w_+} \left(\log\left(\frac{1}{w_+ m} + \frac{2}{m} + 1\right) - \log\left(\frac{1}{w_+ m} + (1 - r) + \frac{1}{m}\right) + \mathcal{O}\left(\frac{1}{m}\right) \right) \\ & \hspace{15em} \text{(by Lemma 2)} \\ & = -\frac{1}{w_+} \log(1 - r) + \mathcal{O}\left(\frac{1}{m}\right), \hspace{10em} \text{(by Taylor expansion of log)} \end{aligned}$$

which holds with probability at least $1 - \delta$, for all $m \geq \max(\sqrt{2/\delta}, m_1)$, where m_1 is the smallest positive integer that satisfies $\left| \frac{1+2w_+}{w_+ m_1} \right| \leq 1$, and $\left| \frac{1+w_+}{w_+ m_1} \right| \leq 1 - r$. These conditions are necessary for the Taylor expansions of the logarithms to be converging. Now we can bound the distance of the posterior mean from the true observation time as follows,

$$\begin{aligned} |M_{\text{post}} - t^*| & = \left| -\frac{1}{w_+} \log(1 - r) + \frac{1}{w_+} \log\left(e^{-w_+ t^*}\right) + \mathcal{O}\left(\frac{1}{m}\right) \right| \\ & = \frac{1}{w_+} \left| \log\left((1 - r)e^{-w_+ t^*}\right) \right| + \mathcal{O}\left(\frac{1}{m}\right) \\ & = \frac{1}{w_+} \left| \log\left(1 + (1 - r - e^{-w_+ t^*})e^{w_+ t^*}\right) \right| + \mathcal{O}\left(\frac{1}{m}\right) \\ & = \frac{1}{w_+} \left| (1 - r - e^{-w_+ t^*})e^{w_+ t^*} \right| + \mathcal{O}\left(\frac{1}{m}\right) + \mathcal{O}\left(\left(1 - r - e^{-w_+ t^*}\right)^2\right) \\ & \hspace{15em} \text{(by Taylor expansion of } \log(1 + \cdot) \text{)} \\ & \leq \frac{e^{w_+ t^*}}{w_+} \sqrt{\frac{\log m}{m}} + \mathcal{O}\left(\frac{\log m}{m}\right), \hspace{5em} \text{(by Lemma 2)} \end{aligned}$$

which holds with probability at least $1 - \delta$, for all $m \geq \max(\sqrt{2/\delta}, m_1, m_2)$, where m_2 is the smallest positive integer that satisfies $|(1 - r - e^{-w_+ t^*})e^{w_+ t^*}| \leq 1$. Again, this condition is necessary for the Taylor expansion of the logarithm to be converging. The bound of the theorem follows by defining $C_1(w_+, t^*) := \frac{e^{w_+ t^*}}{w_+}$.

The argument for the posterior variance follows a similar structure. By [Lemma 1](#) we have

$$\begin{aligned} V_{\text{post}} &= \frac{1}{w_+^2} (\psi_1(\alpha) - \psi_1(\alpha + \beta)) \\ &= \frac{1}{w_+^2} \left(\frac{1}{\alpha} - \frac{1}{\alpha + \beta} + \mathcal{O}\left(\frac{1}{(\alpha + \beta)^2}\right) + \mathcal{O}\left(\frac{1}{\alpha^2}\right) \right) \quad (\text{by } \psi_1 \text{ series expansion}) \\ &= \frac{r}{w_+^2 (1-r)} \frac{1}{m} + \mathcal{O}\left(\frac{1}{m^2}\right). \end{aligned}$$

Using [Lemma 2](#), we can show that $r/(1-r)$ is upper bounded by a constant in m with probability at least $1 - \delta$. In particular, for all $m \geq \max(\sqrt{2/\delta}, m_3)$, where $m_3 := 1/(1 - e^{-2w_+ t^*})$, we get $1 - r \geq e^{-w_+ t^*}/2 > 0$. It follows that

$$\frac{r}{1-r} \leq \frac{2 - e^{-w_+ t^*}}{e^{-w_+ t^*}}.$$

The posterior variance bound follows by defining $C_2(w_+, t^*) := \frac{2 - e^{-w_+ t^*}}{w_+^2 e^{-w_+ t^*}}$.

To conclude the formulation of the theorem, we define $m_0 := \max(\sqrt{2/\delta}, m_1, m_2, m_3)$. \square

Illustration of C_1 and C_2

The following figure shows the constants $C_1(w_+, t^*)$ and $C_2(w_+, t^*)$ for different values of w_+ and t^* . Note that as we keep increasing the true observation time t^* , the values of w_+ that minimize C_1 and C_2 keep decreasing. This verifies the intuition provided in the main text that larger rates for the independent elements in V_+ , that is, larger values of w_+ will be more suitable (i.e., will have faster convergence in m) for determining smaller observation times t^* and vice versa. We also see that larger times are in general harder to estimate as evidenced by the increase of the minimum values of $C_1(\cdot, t^*)$ and $C_2(\cdot, t^*)$ with increasing t^* .




Figure 5: The terms C_1 and C_2 that appear in the bounds of [Theorem 1](#), plotted for different values of w_+ and t^* .

C Violating the assumptions of Theorem 1

We show that the time posterior can behave in a similar way to what is described in Theorem 1 even when the assumption about Θ_+ having the form $\theta_+ \mathbf{I}_m$ is relaxed.

In a similar setup to Figure 1, given a matrix Θ_+ , we draw $N = 1000$ samples and compute the variance of the posterior time distribution for each sample. We then plot the average variance over the samples for each choice of Θ_+ .

Three of the lines in Figure 6 satisfy the assumptions of Theorem 1 each with a different choice of θ_+ . The fourth line corresponds to independently drawing each diagonal entry of Θ_+ uniformly from $[-3, -1]$. Both axes are plotted in logarithmic scale, and the error bars denote two standard errors of the mean. The dashed line denotes the function c/m for some constant c , and is provided for reference. Note that the setting where the diagonal parameters are drawn uniformly exhibits similar behavior to the constant diagonal setting. In particular, we can see that in all cases the variances asymptotically decrease with rate $1/m$, as shown in our theorem.

Next, we assume that Θ_+ is block-diagonal with blocks of size 2, and introduce positive off-diagonal parameters γ within each block. In other words, we violate the independence assumption of Theorem 1, and introduce increasingly attractive behavior between pairs of elements. To make sure that we are able to compare the different choices of γ on equal grounds, we adjust the diagonal parameters of Θ_+ , so that the marginal frequency of each item in S_+ stays constant as we vary γ . In Figure 7, we see that the posterior variance increases up to some point with increasing interaction strength. Intuitively, we expect that for high enough γ , the two elements of each block are practically co-occurring (i.e., either neither or both are observed), and therefore the effect of each block on the time posterior is reduced to the effect of a single element. Figure 7 confirms this intuition; for example, the average variance for $m = 100$ and $\gamma = 12$ can be seen to be roughly the same as the average variance for $m = 50$ and $\gamma = 0$ (independent case).

In Figure 8, we show analogous results for $\gamma < 0$, that is, repulsive interaction between pairs of items in S_+ . In this case, the effect on the posterior time estimates seems to be less noticeable compared to the attractive case.




Figure 6: The average variance of the time posteriors for different choices of diagonal Θ_+ .




Figure 7: The average variance of the time posteriors for different strengths of attractive interaction. The independent case $\gamma = 0$ is equivalent to the $\theta_+ = -2$ case in Figure 6.




Figure 8: The average variance of the time posteriors for different strengths of repulsive interaction. The independent case $\gamma = 0$ is equivalent to the $\theta_+ = -2$ case in Figure 6.

D Proof of Proposition 2

Proposition 2. *The marginal probability of a partial sequence $\sigma = (\sigma_1, \dots, \sigma_k)$ can be written as*

$$p(\sigma; \boldsymbol{\theta}) = \left(\prod_{i=1}^k \frac{q_{\sigma_{[i-1]} \rightarrow \sigma_{[i]}}(\boldsymbol{\theta})}{1 + \tilde{q}_{\sigma_{[i-1]}}(\boldsymbol{\theta})} \right) \frac{1}{1 + \tilde{q}_{\sigma_{[k]}}(\boldsymbol{\theta})}.$$

Proof. Event \mathcal{B} is defined by $T_k < T_{\text{obs}} < T_{k+1}$. To compute $\mathbb{P}(\mathcal{B} \mid \mathcal{A})$ we focus first on the event defined by one of the above inequalities, and rewrite its probability using the definition of jump times and the chain rule as follows,

$$\begin{aligned} \mathbb{P}(T_{\text{obs}} > T_k \mid \mathcal{A}) &= \mathbb{P}\left(T_{\text{obs}} > \sum_{i=1}^k H_i \mid \mathcal{A}\right) \\ &= \prod_{j=1}^k \mathbb{P}\left(T_{\text{obs}} > \sum_{i=1}^j H_i \mid T_{\text{obs}} > \sum_{i=1}^{j-1} H_i, \mathcal{A}\right). \end{aligned}$$

The crucial observation here is that we can greatly simplify the above expression by making use of the memoryless property of exponential random variables (Bertsekas & Tsitsiklis, 2008). According to this property, if Z is an exponential random variable, then $\mathbb{P}(Z > x + y \mid Z > x) = P(Z > y)$, for all $x, y \in \mathbb{R}$. The above expression then becomes

$$\begin{aligned} \mathbb{P}(T_{\text{obs}} > T_k \mid \mathcal{A}) &= \prod_{j=1}^k \mathbb{P}(T_{\text{obs}} > H_j \mid \mathcal{A}) \\ &= \prod_{j=1}^k \mathbb{P}(H_j - T_{\text{obs}} < 0 \mid \mathcal{A}) \\ &= \prod_{j=1}^k \frac{\tilde{q}_{\sigma_{[j-1]}}(\boldsymbol{\theta})}{1 + \tilde{q}_{\sigma_{[j-1]}}(\boldsymbol{\theta})}. \end{aligned} \tag{8}$$

The last equality comes from the fact that the distribution of the difference of two independent exponential random variables with rates λ_1, λ_2 is an asymmetric Laplace distribution with parameters λ_2 and λ_1 on the negative and positive half-lines respectively.

An analogous derivation gives us

$$\begin{aligned} \mathbb{P}(T_{\text{obs}} < T_{k+1} \mid T_{\text{obs}} > T_k, \mathcal{A}) &= 1 - \mathbb{P}(T_{\text{obs}} > T_{k+1} \mid T_{\text{obs}} > T_k, \mathcal{A}) \\ &= 1 - \mathbb{P}(T_{\text{obs}} > H_{k+1} \mid \mathcal{A}) && \text{(by memorylessness)} \\ &= 1 - \frac{\tilde{q}_{\sigma_{[k]}}(\boldsymbol{\theta})}{1 + \tilde{q}_{\sigma_{[k]}}(\boldsymbol{\theta})} && \text{(by difference of exp.)} \\ &= \frac{1}{1 + \tilde{q}_{\sigma_{[k]}}(\boldsymbol{\theta})}. \end{aligned} \tag{9}$$

Finally, combining the results of (8) and (9), we get the probability of a partial sequence,

$$\begin{aligned} p(\sigma; \boldsymbol{\theta}) &= \mathbb{P}(\mathcal{A}) \mathbb{P}(\mathcal{B} \mid \mathcal{A}) \\ &= \mathbb{P}(\mathcal{A}) \mathbb{P}(T_{\text{obs}} > T_k \mid \mathcal{A}) \mathbb{P}(T_{\text{obs}} < T_{k+1} \mid T_{\text{obs}} > T_k, \mathcal{A}) \\ &= \left(\prod_{i=1}^k \frac{q_{\sigma_{[i-1]} \rightarrow \sigma_{[i]}}(\boldsymbol{\theta})}{1 + \tilde{q}_{\sigma_{[i-1]}}(\boldsymbol{\theta})} \right) \frac{1}{1 + \tilde{q}_{\sigma_{[k]}}(\boldsymbol{\theta})}. \end{aligned}$$

□

E MCMC proposal

As a reminder, we would like to draw samples from $p(\cdot | S; \theta)$, which we do by using a Metropolis-Hastings chain over state space \mathcal{S}_S . At each time step, given the current permutation σ , the chain proposes a new permutation σ_{new} according to proposal distribution $Q(\sigma_{\text{new}} | \sigma)$, and transitions to σ_{new} with probability

$$p_{\text{accept}} = \min \left(1, \frac{p(\sigma_{\text{new}} | S; \theta) Q(\sigma | \sigma_{\text{new}}; \theta)}{p(\sigma | S; \theta) Q(\sigma_{\text{new}} | \sigma; \theta)} \right). \quad (10)$$

We focus on proposal distributions $Q(\sigma_{\text{new}} | \sigma; \theta) = Q(\sigma_{\text{new}}; \theta)$ that do not depend on the current state σ . A simple such choice is the uniform proposal $Q(\sigma_{\text{new}}) = 1/|S|!$. In general, the convergence rate of such proposals depends crucially on the minimum ratio $Q(\cdot)/p(\cdot | S; \theta)$ over all states (Mengersen & Tweedie, 1996); intuitively, we want the proposal to match as well as possible the true distribution, and, in particular, to have a high chance to propose states that have high probability according to $p(\cdot | S; \theta)$.

Algorithm 2 shows how to draw a permutation σ from our proposal Q , and at the same time compute its unnormalized probability $Q_{\text{val}} \propto Q(\sigma; \theta)$. This is enough to use our proposal in the Metropolis-Hastings chain, since the acceptance probability in (10) only requires computing ratios $Q(\sigma; \theta)/Q(\sigma_{\text{new}}; \theta)$. As shown in the algorithm, given a set S , we iteratively add items to σ by considering at each step a weight u_ν for each candidate item $\nu \in S \setminus \sigma$. The form of u_ν bears a strong similarity to the form of the marginal sequence probability derived in (6). In fact, it is easy to show that for $|V| = 2$ this algorithm samples exactly from the correct distribution, that is, $Q(\cdot) = p(\cdot | S; \theta)$. More generally, the factor c_ν approximates the denominator of eq. (6), and intuitively takes into account the interactions between the elements that are already in σ and all other elements in V . The factor $\prod w_{\nu j}$ computes the numerator of eq. (6), and intuitively takes into account the interactions that would occur between item ν (were we to add this item at this position) and all remaining items in S that would be added subsequently.

To compare this proposal to the uniform one, we use the same synthetic data set as in Figure 1, but constrain the size of the ground set to be at most $n = 20$, so that we can compare to the exact gradients. We first run 100 gradient ascent epochs using the exact gradients, and then approximate the gradient of the marginal likelihood at that point using the Metropolis-Hastings sampler with the two different proposals. If we denote by \mathbf{g} the true gradient, and by $\hat{\mathbf{g}}$ the gradient approximation, Figure 9 plots the error $\|\hat{\mathbf{g}} - \mathbf{g}\|_2$ as a function of the number of samples used. Note that for $n = 10$ the two proposals seem to perform similarly, but as n gets larger the advantage of our proposal becomes increasingly more pronounced. In practice we have observed that for larger ground sets our proposal tends to provide considerable improvements to the convergence speed of the optimization algorithm, and also be much less sensitive to the choice of step size compared to the uniform proposal.

Algorithm 2: Drawing from proposal Q

Input: Parameters θ , set S

$Q_{\text{val}} \leftarrow 0$

$\sigma \leftarrow ()$

for $k = 1, \dots, |S|$ **do**

for $\nu \in S \setminus \sigma$ **do**

$\sigma' \leftarrow \sigma \cup \{\nu\}$

$d_\nu \leftarrow 1 + \sum_{j \in V \setminus \sigma'} w_{jj} \prod_{i \in \sigma'} w_{ij}$

$u_\nu \leftarrow \frac{1}{d_\nu} \prod_{j \in S \setminus \sigma'} w_{\nu j}$

 Draw $x \sim \text{Cat}(S \setminus \sigma, (u_\nu / \sum_j u_j)_{\nu \in S \setminus \sigma})$

$\sigma \leftarrow \sigma \parallel x$

$Q_{\text{val}} \leftarrow Q_{\text{val}} u_x / \sum_j v_j$

return σ, Q_{val}




Figure 9: The norm of the difference between the true gradient and the gradient approximation using sampling. Dashed lines correspond to the uniform proposal, while solid lines correspond to our proposal described in Algorithm 2. Each color denotes a different ground set size.


F Synthetic data sets

Two-item data set



$$\Theta^* = \begin{bmatrix} 0 & 4 \\ 0 & -4 \end{bmatrix}$$

Five-item data set



$$\Theta^* = \begin{bmatrix} -1 & 4 & 0 & 0 & 0 \\ 0 & -1 & -2 & -2 & 0 \\ 0 & -2 & -1 & -2 & 0 \\ 0 & -2 & -2 & -0.5 & 4 \\ 0 & 0 & 0 & 0 & -4 \end{bmatrix}$$

G Further experimental results

Figure 10 shows results on the TCGA glioblastoma data set used by Schill et al. (2019), and originally preprocessed by Leiserson et al. (2013). (This is an older version of the data used in our experiments.) We run our method on the same subset of $n = 20$ genetic alterations chosen by Schill et al. (2019). The two plots in the top row show the learned parameter matrices Θ for two different random initializations. The left matrix is practically identical to the result reported by Schill et al. (2019), while the right matrix has some notable differences. To further highlight these differences, the bottom plot shows the range (max - min value) of each learned parameter across 20 random initializations.

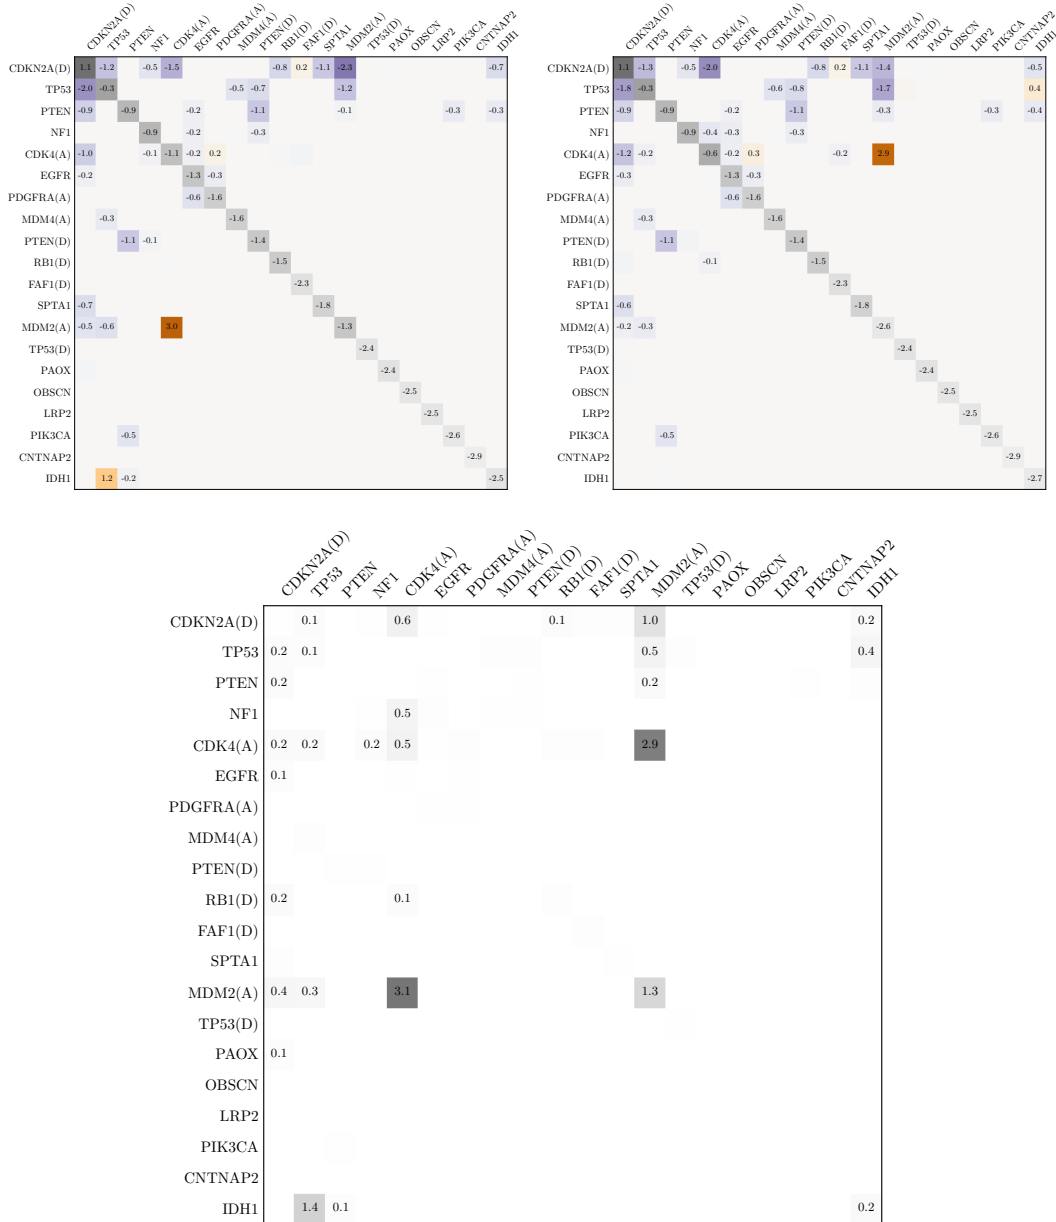


Figure 10: Top: Learned parameter matrices Θ on the reduced data set ($n = 20$) used by Schill et al. (2019) for two different random initializations. Orange shades denote $\theta_{ij} > 0$, while purple shades denote $\theta_{ij} < 0$. The zero entries are left blank. Bottom: Learned parameter ranges across 20 random initializations.

Approximately block-diagonal interaction structure in real data




Figure 11: The learned parameter matrix Θ for the $n = 150$ most frequent genetic alterations in the TCGA glioblastoma data set discussed in Section 6. Note that the matrix consists of a smaller block of complex dependencies followed by a larger block of approximately independent items.