

Bayesian Changepoint Analysis

I n a u g u r a l d i s s e r t a t i o n

zur

Erlangung des akademischen Grades eines
Doktors der Naturwissenschaften (Dr. rer. nat.)

der

Mathematisch-Naturwissenschaftlichen Fakultät

der

Universität Greifswald

Vorgelegt von

Tobias Siems

Dossenheim, der 25.10.2019

Dekan:

Prof. Dr. Werner Weitschies

1. Gutachter:

Prof. Dr. Volkmar Liebscher

2. Gutachter:

Prof. Dr. Paul Fearnhead

Tag des Promotionskolloquiums:

27.02.2020

Contents

1	The changepoint universe (a poetic view)	1
2	Introduction	2
2.1	Time series, segmentations and changepoints	5
2.2	Changepoint estimation via minimizing a penalized cost	6
2.3	Bayesian changepoint models	8
2.4	The purpose of this thesis and its structure	10
3	Markov chain Monte Carlo on finite state spaces	13
3.1	Convergence towards and existence of invariant distributions	14
4	A note on the Metropolis-Hastings acceptance probabilities for mixture spaces	19
4.1	Outline	19
4.2	Introduction	19
4.3	The Metropolis-Hastings algorithm in applied terms	24
4.3.1	Mixture state spaces and translations	26
4.3.2	Metropolis-within-Gibbs	33
4.4	A changepoint example	37
4.4.1	A plain implementation of birth and death	39
4.4.2	An ad-hoc implementation of birth and death	39
4.4.3	A post-hoc implementation of birth and death	40
4.4.4	Convergence	41
4.4.5	Results and conclusions	42
4.5	The Metropolis-Hastings algorithm in abstract terms	43
4.5.1	Mixture proposals in general and maximality in particular	44
4.6	Discussion and literature review	47

5	Exact inference in Bayesian changepoint models	51
5.1	Model assumption	53
5.2	Forward inference	55
5.2.1	Pruning	60
5.2.2	The exponential family	62
5.2.3	The Laplacian change in median model	64
5.2.4	Well-log example	68
5.3	Backward inference	70
5.3.1	Likelihood of changepoint locations and MAP estimator	72
5.3.2	Sampling	72
5.3.3	Entropy	73
5.4	Pointwise inference	75
5.4.1	Pointwise statements in the well-log example	78
5.5	Parameter estimation via likelihood and the EM algorithm	80
5.5.1	An EM algorithm for changepoint models	82
5.5.2	An EM algorithm for segment length	84
5.5.3	Parameter estimation in the well-log example	88
5.5.4	An EM algorithm for segment height and observation	93
5.5.5	EM algorithm for the scales in the well-log example	98
5.6	Discussion	102
6	Simultaneous credible regions for multiple changepoint locations	109
6.1	Outline	109
6.2	Introduction	109
6.3	The Sample Based Problem	112
6.3.1	Reformulation of the Sample Based Problem and its complexity	114
6.3.2	A greedy heuristic	115
6.4	Alternative approaches	116
6.4.1	Highest density regions	117

6.4.2	Marginal changepoint probabilities	118
6.5	Multiple changepoints and credible regions in more detail	119
6.5.1	Importance revisited	119
6.5.2	Exemplary comparison	120
6.5.3	Empirical proof of convergence and the accuracy of Greedy	123
6.6	An example of use for model selection	124
6.7	Simultaneous credible regions in the well-log example	125
6.8	Discussion	128
7	Final Remarks	131
8	Related own papers and significant contributions of other researchers to this thesis	132
A	Supplementary material	133
B	Appendix	133
B.1	How to draw a fixed number of changepoints uniformly	133
B.2	Hypergraphs and the complexity of the SBP	134
B.2.1	Hypergraphs	134
B.2.2	Computational complexity of the SBP	135
B.3	Proof of correctness of the ILP	139
	Symbols	141
	Figures	142
	Tables	143
	Pseudocodes	143
	References	143

1 The changepoint universe (a poetic view)

Changepoint problems unite three worlds: the world of counting, the world of placement and the world of change. These worlds constitute universes, which receive their rough shape from statistical models applied to data and are further refined through parameters.

Even though the defining elements of such universes are grasped so quickly, we can never fully understand them as a composite whole. Like a physicist, we need to rely on lenses revealing tiny pictures of the incredible truth, such as the astonishing glimpse into the counting world unveiled by Figure 1.

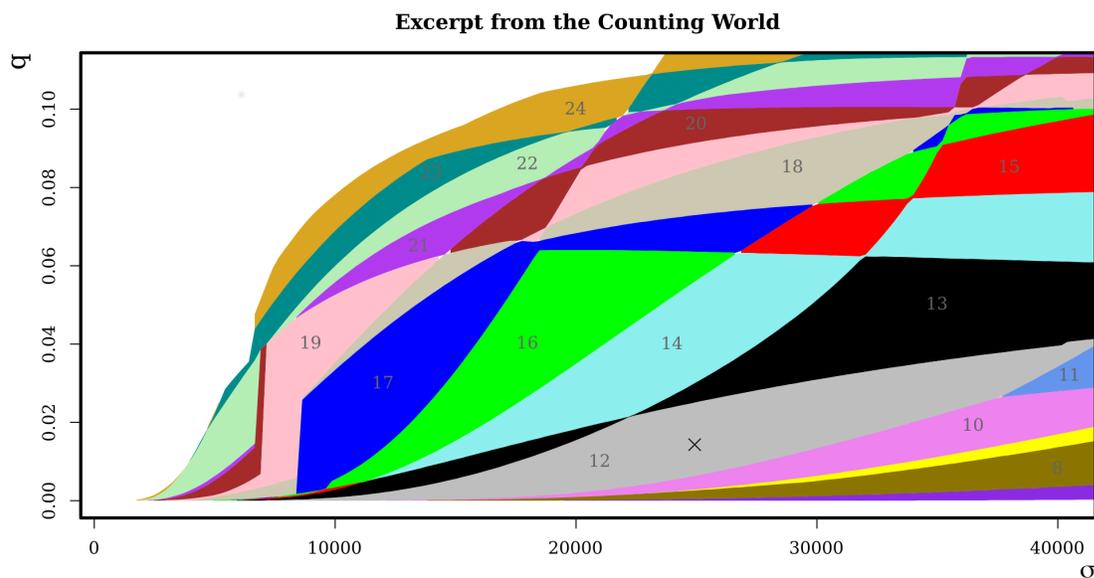


Figure 1: Data-driven partitioning of the parameter space of a changepoint model. Each of the colored regions stands for a certain changepoint count, obtained by computing the most likely changepoint configuration. q is the success probability of the length distribution of the segments and σ the scale of the data within those segments (see also Section 5.5.3).

This thesis gives you the unique opportunity to gain an understanding of the fantastic scale of Bayesian changepoint universes. To this end, we will look through a range of enlightening lenses, forged in the mighty fires of mathematics and operated on a highly sophisticated calculator, to discover places no one has known before.

2 Introduction

Changepoint analysis deals with time series where certain characteristics undergo occasional changes. Having observed such a time series, the latent changepoint process can be investigated in various ways. The following four examples provide an insight into the tasks that may arise in practice.

The examination of ion channels is a popular example for changepoint analysis. Ion channels are proteins and part of the cell membrane. Their purpose is to control the flow of ions, like potassium, into and out of the cell. New drugs must pass extensive tests for ion channel activity (Camerino et al., 2007). Depending on the type of the channel, they can be activated and deactivated through chemical reactions or electrical currents.

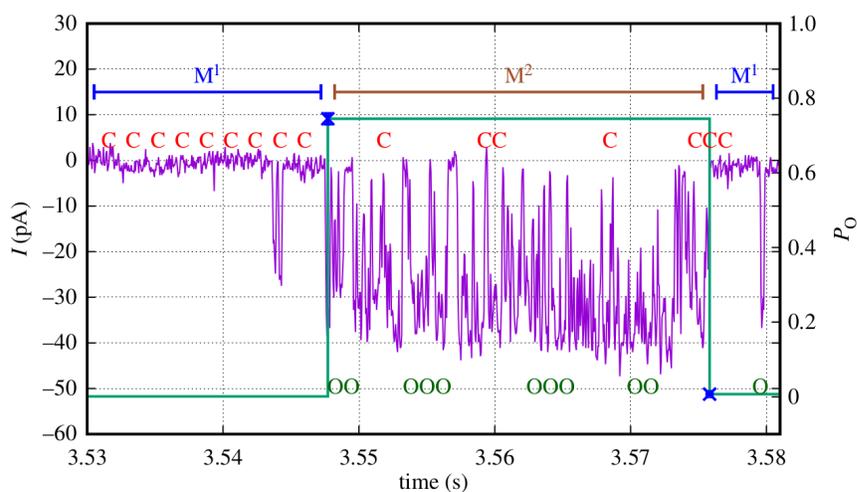


Figure 2: Ion channel data (Siekmann et al., 2016).

A Nobel prize winning technique to gather data from single ion channels is the patch clamp technique (Sakmann and Neher, 1984). It allows to record the electrical current that is generated by the flow of ions. Figure 2 shows a data set gathered from an inositol trisphosphate receptor, which controls the flow of calcium ions (Siekmann et al., 2016). The purple lines represent the measured currents, the C's and O's mark the times where the channel was closed or opened, respectively. This is called gating and can be understood

as a changepoint process. However, there is a second superordinate changepoint process indicated by M^1 and M^2 . It models the modes of the channel’s activity.

These modes reflect the stochastic nature of how the cell controls the flow of ions. Thus, it is of high interest to find and understand the possible modes an ion channel takes. Each mode can be represented through a time continuous Markov chain that describes its gating behavior (Siekmann et al., 2016). The transitions between the modes may again be modeled as a Markov chain leading to a hierarchical model. A challenging problem hereby is the amount of quantities which are to be estimated in order to determine the model and the associated non-identifiability. Inference in these models may be drawn by means of approximate sampling through Markov chain Monte Carlo (MCMC) (Siekmann et al., 2011, 2012, 2014, 2016).

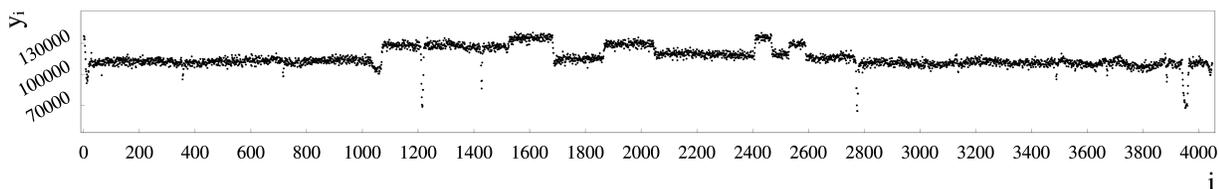


Figure 3: Well-Log data (Fearnhead, 2006).

A different example for changepoint analysis, which arises in geology, is the well-log drilling data in Figure 3. It shows the nuclear-magnetic response of underground rocks (Fearnhead, 2006).

In order to build a changepoint model here, we may regard coherent parts of the data as noisy versions of one and the same state and allow these states to change occasionally. This can be understood as a change in location problem whereby the states are interpreted as the location parameter of the datapoints, e.g. median or mean. Sequences of related datapoints are then considered to stem from the same material.

Interestingly, there are challenging perturbations in form of heavy outliers and smaller abrupt, but also gradual location changes. These irregularities are prone to be confused with the desired location changes, which gives rise to robust changepoint inference (Fearnhead,

2006; Fearnhead and Rigaiil, 2019; Weinmann et al., 2015). The task here is to maintain a good balance between sensitivity and specificity.

This exciting data set serves as the main demonstration example throughout this thesis.

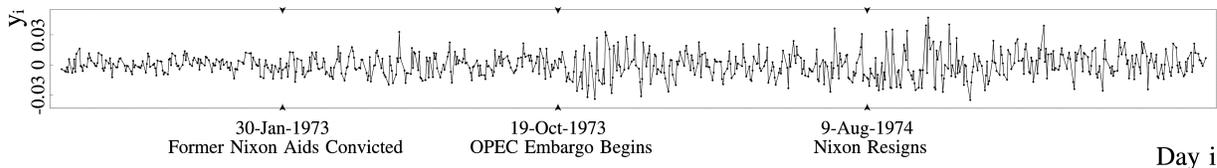


Figure 4: Daily Dow Jones returns between 1972 and 1975 (Adams and MacKay, 2007).

Stock markets pose another important changepoint application. Consider for example Figure 4. It depicts daily Dow Jones returns observed between 1972 and 1975 (Adams and MacKay, 2007) and highlights three events on January 1973, October 1973 and August 1974.

This can be understood as a change in variation problem, whereby the data is split into regions of similar spread. Changepoint analysis may be used for a real time detection of high risk or to understand the effect of political decisions on the economy.

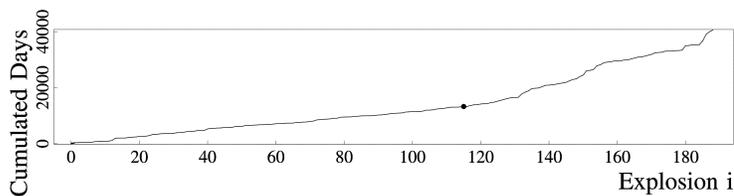


Figure 5: Coal mine explosions killing ten or more men between March 15, 1851 and March 22, 1962 (Jarrett, 1979; Adams and MacKay, 2007). The abscissa shows the explosion count and the ordinate the number of days until the explosions took place. The black dot marks the Coal Mines Regulations Act in 1887.

Our last example concerns coal mine explosions between March 15, 1851 and March 22, 1962 that killed ten or more people. Figure 5 counts up the days from one explosion to

the next. The black dot points to the Coal Mines Regulations Act in 1887, which also roughly marks a change in the period of time it took from one explosion to the next.

In a changepoint approach, the days between consecutive explosions could be modeled as geometrically distributed and changes in success rate would be examined. The task here is to elaborate the effect of the regulations act by detecting or rejecting one single change. This kind of single changepoint approach can be applied to a large variety of applications where we expose if something has generally changed.

For the time being, this should be enough to motivate changepoint analysis. However, for further elaborations of different changepoint approaches see, for example, Eckley et al. (2011).

2.1 Time series, segmentations and changepoints

Time series are ordered sequences of data. This order can, for example, account for successive loci on a DNA sequence or indeed the times the observations were made and allows us to speak of segmentations.

Let \mathbb{Y} be a set and $y = (y_1, \dots, y_n) \in \mathbb{Y}^n$ be a time series with indexes $1, \dots, n$. Any sequence, (y_j, \dots, y_i) with $1 \leq j \leq i \leq n$, of successive datapoints is considered as a *segment* with *segment length* $i - j + 1$. A *segmentation* of y is a collection of non-overlapping segments that covers y as a whole.

A segmentation is uniquely determined by a number $k \in \{0, \dots, n - 1\}$ and a sequence of successive indexes $1 < \tau_1 < \dots < \tau_k \leq n$ as follows

$$(y_1, \dots, y_{\tau_1-1}), (y_{\tau_1}, \dots, y_{\tau_2-1}), \dots, (y_{\tau_{k-1}}, \dots, y_{\tau_k-1}), (y_{\tau_k}, \dots, y_n)$$

whereby $k = 0$ indicates that the whole time series corresponds to one segment. In the context of changepoints, k represents the number of changepoints and the τ_1, \dots, τ_k are the *changepoint locations* or simply *changepoints*. For now, changepoints at timepoint one are excluded. However, this will be relaxed later on.

As mentioned earlier, the segmentation features changes in certain characteristics of the time series, e.g. location or scale. These characteristics are expressed through parameters,

the so-called *segment heights* that take values within a set, say \mathbb{X} . Consequently, we assign each segment of a segmentation a value within \mathbb{X} . In practice, a change in segment height typically triggers a changepoint.

In non-trivial cases, the segment heights are latent and algorithms for finding changepoints in a time series trace possible timepoints where they change their value. In doing so, they strive for persistency but also selectivity or as the statistician would put it, they govern a trade-off between specificity and sensitivity. While a too high sensitivity yields too many changepoints overfitting the segmentation to the data, a too high specificity disregards valuable changepoints and impairs the changepoint detection.

The existing literature captures approaches to detect at most one changepoint as well as multiple changepoints. Even though single changepoint methods can be applied recursively to detect multiple changepoints (Fryzlewicz et al., 2014), in this thesis, we focus solely on techniques that surmise multiple changepoints in the first place.

2.2 Changepoint estimation via minimizing a penalized cost

A natural approach for finding multiple changepoints in a time series is to assign each segmentation a rating and then select one of the best ones. This rating should express the goodness of fit, but also regularize the number of changepoints to avoid overfitting. However, instead of speaking about goodness of fit and regularization, it is common to use opposite terms like *cost* and *penalization*. Thus, changepoints are inferred by *minimizing a penalized cost*.

For the sake of computational convenience, the costs and penalizations are often defined individually for the segments instead of the whole segmentation at once. This can be achieved through a *segmental cost* $\mathbf{f} : \bigcup_{i=1}^n \mathbb{Y}^i \rightarrow \mathbb{R}_{\geq 0}$ and a penalization constant $\gamma \geq 0$. For the time being, we may ignore the segment heights.

Minimizing the penalized cost then yields the optimization problem which finds one

element of

$$\operatorname{argmin}_{1 < \tau_1 < \dots < \tau_k \leq n \text{ and } k \in \mathbb{N}} \left\{ \mathbf{f}(y_{1:\tau_1-1}) + \left(\sum_{i=1}^{k-1} \mathbf{f}(y_{\tau_i:\tau_{i+1}-1}) \right) + \mathbf{f}(y_{\tau_k:n}) + \gamma \cdot k \right\} \quad (1)$$

This can be solved through a Viterbi algorithm (Viterbi, 1967), also known as Bellman recursion (Bellman, 1966) or dynamic programming. In oversimplified terms, the algorithm performs a forward search from 1 to n , whereby in each iteration i it performs a backward search in order to evaluate a best candidate for the most recent changepoint prior to i . Subject to the condition that \mathbf{f} can be augmented by new observations in constant time, i.e. $\mathbf{f}(v_{1:i})$ can be derived from $\mathbf{f}(v_{1:i-1})$ in constant time for all $v_{1:i} \in \mathbb{Y}^i$, this yields a space and time complexity of $\mathcal{O}(n^2)$. A long series of papers explains this algorithm in detail (see for example Jackson et al. (2005); Friedrich et al. (2008) and Killick et al. (2012)).

Furthermore, an approach developed by Friedrich et al. (2008) computes solutions to (1) for all possible γ at once. It does so, by exploiting the linearity of the penalized cost w.r.t. γ . Almost identical ideas are provided by Haynes et al. (2017).

A discovery that earned considerable popularity is described in Killick et al. (2012). It allows under some circumstances to reduce the expected space and time complexity of solving (1) to $\mathcal{O}(n)$. We briefly discuss the idea behind it.

Assume that $\mathbf{f}(v_{1:j}) + \mathbf{f}(v_{j+1:i}) \leq \mathbf{f}(v_{1:i})$ for any sequence $v_{1:i} \in \mathbb{Y}^i$ and $1 \leq j < i$. This makes sense, because splitting a segment usually results in a tighter fit and thus a lower cost. Let further \mathcal{F}_i be a solution to (1) w.r.t. $y_{1:i}$, i.e.

$$\mathcal{F}_i \in \operatorname{argmin}_{1 < \tau_1 < \dots < \tau_k \leq i \text{ and } k \in \mathbb{N}} \left\{ \mathbf{f}(y_{1:\tau_1-1}) + \left(\sum_{\ell=1}^{k-1} \mathbf{f}(y_{\tau_\ell:\tau_{\ell+1}-1}) \right) + \mathbf{f}(y_{\tau_k:i}) + \gamma \cdot k \right\}$$

If for $1 \leq j < i \leq n$ we get $\mathcal{F}_j + \mathbf{f}(y_{j+1:i}) \geq \mathcal{F}_i$, we can take an arbitrary $\ell > i$ and get

$$\mathcal{F}_j + \mathbf{f}(y_{j+1:\ell}) + \gamma \geq \mathcal{F}_j + \mathbf{f}(y_{j+1:i}) + \mathbf{f}(y_{i+1:\ell}) + \gamma \geq \mathcal{F}_i + \mathbf{f}(y_{i+1:\ell}) + \gamma$$

This states that placing the most recent changepoint prior to ℓ at i , is at least as good as placing it at j . Thus, we may skip j in any backward search past i . An implementation of the corresponding algorithm called PELT, can be found in the R-Package *changepoint* (Killick et al., 2012).

The starting point for setting up the segmental cost is usually another function $\mathbf{g} : \bigcup_i \mathbb{Y}^i \times \mathbb{X} \rightarrow \mathbb{R}_{\geq 0}$, which takes the data within a segment and the segment height. \mathbf{g} can, for example, be derived as the negative loglikelihood of a statistical model. However, M-estimators are also viable (Friedrich et al., 2008; Fearnhead and Rigall, 2019).

We may put \mathbf{g} directly into (1) and additionally minimize over the individual segment heights or equivalently, use $\mathbf{f}(\cdot) = \min_{x \in \mathbb{X}} \mathbf{g}(\cdot, x)$. The solution can then be interpreted as a piecewise constant function with values corresponding to the optimal segment heights. Consequently, we can speak of a regression method here.

In contrast, the Bayesian statistician may integrate over the segment heights in order to get rid of them completely. This involves a more complex approach since we need to put a prior on the segment heights and perform an integration for each considered segment.

Li and Lund (2012) make use of priors and run a genetic algorithm, which detects multiple changepoints by minimizing a functional similar to the penalized cost. Another non-Bayesian method to infer multiple changepoints through a dynamic programming algorithm can be found in Frick et al. (2014).

2.3 Bayesian changepoint models

By applying priors, we pave the way for an interpretation of segmentations as random. This allows for the treatment of the data, changepoints and segment heights within the same methodological realm of stochastics. As a result, a more complex modeling approach is required here.

To begin with, the changepoints are represented by continuous or discrete random times, $(\mathbf{T}_i)_{i \in \mathbb{N}_+}$ over $\mathbb{R}_{\geq 0}$. The changepoint process starts at a certain timepoint, say $t_0 \in \mathbb{R}$, whereby the first changepoint is $t_0 + \mathbf{T}_1$, the second $t_0 + \mathbf{T}_1 + \mathbf{T}_2$ and so forth. Consequently, \mathbf{T}_i represents the length of the i -th segment.

Furthermore, the segment heights are modeled accordingly through a family of random variables $(\boldsymbol{\theta}_i)_{i \in \mathbb{N}_+}$ over a measurable space $(\mathbb{X}, \mathcal{X})$, whereby $\boldsymbol{\theta}_i$ represents the i -th segment height. Finally, the observations constitute a family $(\mathbf{Y}_i)_{i=1, \dots, n}$ over a measurable space

$(\mathbb{Y}, \mathcal{Y})$. Each observation \mathbf{Y}_i was measured at a timepoint t_i , whereby $t_0 \leq t_1 < \dots < t_n$.

An in some ways minimal requirement for a statistical model to be expedient, is the feasibility of its likelihood. This enables the use of MCMC techniques like Metropolis-Hastings and thus, pioneers an abundance of inference strategies. However, before we can set up the likelihood, we need to consider an issue that typically results from the data collection.

In practice, t_0 often remains unknown. It is thus reasonable to build the changepoint model solely within $[t_1, t_n]$. Consequently, we interpret the first segment as a clipped random length with a yet to be determined distribution. Given the first changepoint within $[t_1, t_n]$, the second segment length will follow the same distribution as the \mathbf{T}_i 's and so forth. As soon as a segment exceeds t_n the process ends irrespective of its concrete length.

We may account for the lack of knowledge about t_0 by considering the limit $t_0 \rightarrow -\infty$. To this end, assume that $(\mathbf{T}_i)_{i \in \mathbb{N}_+}$ is a renewal process (Doob, 1948) and independent of all other random variables within the changepoint model. Renewal theory (Lawrance, 1973) teaches us that the distribution of the random length from t_1 to the first changepoint exhibits the density $\mathbb{P}(t_1 + \mathbf{T}_1 > t) / \mathbb{E}(\mathbf{T}_1)$ w.r.t. the Lebesgue measure. These partial lengths are referred to as residual times or forward recurrence times.

The case where $-\infty < t_0 < t_1$ turns out to be more challenging since the distribution of the residual times now depends on t_0 , and can be very difficult to compute especially in the continuous case.

There are two important exceptions here. The residual times of an exponential renewal process are independent of t_0 and correspond to the same exponential distribution again. Similarly, for a geometric renewal process that exhibits changepoints only on a known discrete grid, we can again take the same geometric distribution for the residual times.

A similar issue arises with the segment heights. A not too restrictive assumption is that $(\boldsymbol{\theta}_i)_{i \in \mathbb{N}_+}$ is described by a Markov chain. If $t_0 < t_1$, related theory may be applicable to derive the distribution of the segment height of the segment that contains t_1 .

Having considered this, we can now set up the joint likelihood of the data, changepoint locations and segment heights in order to use MCMC methods like Metropolis-Hastings to

perform approximate inference.

2.4 The purpose of this thesis and its structure

In this thesis, we elaborate upon Bayesian changepoint analysis, whereby our focus is on three big topics: approximate sampling via MCMC, exact inference and uncertainty quantification. Besides, modeling matters are discussed in an ongoing fashion. Our findings are underpinned through several changepoint examples with a focus on the aforementioned well-log drilling data.

In the following, you find four strongly interrelated, but also widely self-contained sections. While Section 3 provides an understanding of the mathematical foundations of MCMC sampling, Section 4 considers the Metropolis-Hastings algorithm particularly for changepoint related scenarios. Both sections do not intend to develop completely new methods. Instead, Section 3 explains important convergence theorems and gives very basic and simple proofs. Furthermore, our elaboration of the Metropolis-Hastings algorithm aims at clarifying the apparently difficult topic of sampling within mixtures of spaces.

Section 5 talks extensively about exact inference in Bayesian changepoint models with independence induced by changepoints. It develops the bulk of existing algorithms, like exact sampling, and builds novel, very efficient algorithms for pointwise inference and an EM algorithm. At the same time, new scopes for applying approximate changepoint samples are opened up.

Finally, Section 6 develops a notion of credible regions for multiple changepoint locations. They are built from either exact or approximate changepoint samples and illustrated by a new type of plot. This plot greatly facilitates uncertainty quantification, model selection and changepoint analysis. The implementation of our credible regions gives rise to a novel NP-complete optimization problem, which is approached either exactly by solving an ILP or through a fast and accurate greedy heuristic.

The Sections 4 and 6 begin with an outline to summarize their purposes. Besides, each of the following four sections provides an introductory part, which puts its topic into

context with existing research and further clarifies the actual subject. Sections 4, 5 and 6 conclude with a discussion. Section 7 concludes this whole thesis with final remarks.

Section 8 lists related papers that I co-authored and considerable contributions of other researchers to this thesis.

References to the supplementary material can be found in Section A. Section B contains the appendix. At the very end, we provide lists of symbols, figures, tables, pseudocodes and references.

3 Markov chain Monte Carlo on finite state spaces

In this section, we elaborate the idea behind Markov chain Monte Carlo (MCMC) methods in a mathematically coherent, yet simple and understandable way. Therefore, we prove a pivotal convergence theorem for finite Markov chains and a minimal version of the Perron-Frobenius theorem. Only very basic knowledge about matrices, convergence of real sequences and probability theory is required here. A convenient summary of analogous results for general state spaces can be found in Tierney (1994).

MCMC techniques aim at drawing samples from prespecified distributions. They do so in an indirect, approximate fashion through Markov chains. This is important since the distributions deployed in practice are often too complex to be dealt with directly or even unavailable in closed form.

There exists a tremendous number of scientific articles and books about MCMC. See, for example, Bishop and Mitchell (2014) for a vivid and more comprehensive introduction without mathematical proofs.

Let π be a probability distribution over a finite state space \mathcal{S} and $\pi_s \in [0, 1]$ the probability of state $s \in \mathcal{S}$ under π . By virtue of the strong law of large numbers, independent samples from π (so-called *i.i.d. samples*) can be used universally to approximate expectations w.r.t. π . Thus, for a set of such samples, say $x_1, \dots, x_m \in \mathcal{S}$, and an arbitrary function $\mathbf{f} : \mathcal{S} \rightarrow \mathbb{R}^\ell$, we get that $\frac{1}{m} \sum_{i=1}^m \mathbf{f}(x_i) \approx \sum_{s \in \mathcal{S}} \pi_s \mathbf{f}(s)$. This simple recipe poses one of the most powerful tools in statistics.

An example for \mathbf{f} is the indicator function $\mathbb{1}\{s \in A\}$ for events $A \subseteq \mathcal{S}$. It is one if the condition in the brackets is true and zero otherwise. Its expectation yields the probability of A .

A *Markov chain* over \mathcal{S} is defined through an (arbitrary) initial state $s_0 \in \mathcal{S}$ and a *transition kernel* μ . μ is a non-negative function over $\mathcal{S} \times \mathcal{S}$ such that $\sum_{s \in \mathcal{S}} \mu_{zs} = 1$ for all $z \in \mathcal{S}$. μ_z can be interpreted as a conditional distribution.

The Markov chain starts in state s_0 and evolves according to μ in an iterative fashion: the distribution of the first link in the chain is μ_{s_0} and given the first link, say s_1 , the

distribution of the second link is μ_{s_1} and so forth. This results in a potentially infinite sequence of random variables $\mathbf{X} = (\mathbf{X}_0, \mathbf{X}_1, \mathbf{X}_2, \dots)$, whereby \mathbf{X}_n represents the n -th link in the chain. Consequently, we get that $\mathbb{P}(\mathbf{X}_n = s \mid \mathbf{X}_{n-1} = z) = \mu_{zs}$ for $n > 0$ and $\mathbb{P}(\mathbf{X}_0 = s_0) = 1$.

Later on, we will deal with the (unconditional) distributions of the n -th links. To this end, w.l.o.g. we assume that $\mathcal{S} = \{1, \dots, k\}$. Therewith, we describe μ as a matrix in $[0, 1]^{k \times k}$ and π as a column vector in $[0, 1]^k$. Any such quadratic matrix with non-negative entries and rows that sum to one is called *stochastic matrix*. Therewith, $\mathbb{P}(\mathbf{X}_n = s) = (\mu^n)_{s_0 s}$, i.e. the n -fold matrix product of μ evaluated at s_0 and s . A further generalization is to set the law of \mathbf{X}_0 to an arbitrary distribution π . This yields $\mathbb{P}(\mathbf{X}_n = s) = (\pi^t \mu^n)_s$, whereby the t indicates transposition.

We say that a distribution π is an *invariant distribution* of μ if $\pi^t \mu = \pi^t$. Thus, transitioning according to μ doesn't affect π . Once a link in a Markov chain with transition kernel μ follows the law π , all subsequent links will do likewise. In this case, the chain is considered to be in *equilibrium*. Equilibrium can be enforced by setting the distribution of \mathbf{X}_0 to π , but it may also be reached (approximately) in the long run through convergence.

The foundation of MCMC sampling is that under some circumstances the distributions of the n -th links of a Markov chain converge towards an invariant distribution regardless of the initial state. Thus, by simulating such a chain until equilibrium is reached sufficiently, we may obtain an approximate sample of this invariant distribution. On these grounds, MCMC methods provide schemes to build Markov chains with a predefined unique invariant distribution.

3.1 Convergence towards and existence of invariant distributions

This section first considers the convergence of the distributions of the n -th links of certain Markov chains. This convergence forms the very basis of MCMC sampling. Afterwards, we provide a version of the Perron-Frobenius theorem, which gives further insights into the

existence of invariant distributions. The following two properties of μ play a fundamental role.

μ is called *irreducible* if for every $z, s \in \mathcal{S}$ there exists an $n > 0$ such that $(\mu^n)_{zs} > 0$. Thus, regardless of the state the Markov chain starts in, every state can eventually be reached with positive probability.

μ is called *aperiodic* if there exists an N such that $(\mu^n)_{zs} > 0$ for all $n > N$ and all $z, s \in \mathcal{S}$. Thus, regardless of the state the Markov chain starts in, in the long run, it can always reach any state in one step with positive probability.

The name irreducible suggests that the Markov chain does not divide \mathcal{S} into separate, mutually inaccessible classes. In turn, aperiodicity excludes the case where parts of \mathcal{S} are reached in a periodic fashion, for example, only through either an even or odd number of transitions. Aperiodicity implies irreducibility, but not vice versa.

It is obvious that a periodic behavior may impede convergence of the distributions of the n -th links. For a general convergence theorem, aperiodicity is thus a necessary condition. We will now see that it is also sufficient. The following theorem is a simplified version of a convergence theorem for Markov chains over countable state spaces provided inter alia by Koenig (2005).

Theorem 1. *For an aperiodic stochastic matrix μ with invariant distribution π , we get that $\lim_{n \rightarrow \infty} (\mu^n)_{zs} = \pi_s$ for all $z, s \in \mathcal{S}$.*

Proof. Assume that $\mathbf{X} = (\mathbf{X}_0, \mathbf{X}_1, \dots)$ is a Markov chain with transition kernel μ that starts with an arbitrary but fixed state s_0 . Furthermore, consider the Markov chain $\mathbf{Z} = (\mathbf{Z}_0, \mathbf{Z}_1, \dots)$ with transition kernel μ and initial distribution π , i.e. $\mathbb{P}(\mathbf{Z}_0 = s) = \pi_s$ for all $s \in \mathcal{S}$. \mathbf{X} and \mathbf{Z} are supposed to be independent of each other.

Let \mathbf{T} be the random variable that represents the first n where \mathbf{X} and \mathbf{Z} equal, i.e. $\mathbf{T} = \min\{n \mid \mathbf{X}_n = \mathbf{Z}_n\}$. We want to show that \mathbf{T} is finite with probability one. Due to the aperiodicity of μ , we may choose an N such that μ^N has solely positive entries. Let

$\epsilon > 0$ be the smallest entry of μ^N and consider that

$$\begin{aligned} \mathbb{P}(nN < \mathbf{T}) &= \mathbb{P}(\mathbf{X}_i \neq \mathbf{Z}_i \text{ for all } i \leq nN) \leq \mathbb{P}(\mathbf{X}_{i \cdot N} \neq \mathbf{Z}_{i \cdot N} \text{ for all } i \leq n) \\ &= \sum_{z_0 \neq s_0} \pi_{z_0} \sum_{s_1 \in \mathcal{S}} \mu_{s_0 s_1}^N \sum_{z_1 \neq s_1} \mu_{z_0 z_1}^N \cdots \sum_{s_n \in \mathcal{S}} \mu_{s_{n-1} s_n}^N \underbrace{\sum_{z_n \neq s_n} \mu_{z_{n-1} z_n}^N}_{\leq 1 - \epsilon} \leq (1 - \epsilon)^n \xrightarrow{n \rightarrow \infty} 0 \end{aligned}$$

Define further $\mathbf{Y}_n = \mathbb{1}\{n \leq \mathbf{T}\} \mathbf{X}_n + \mathbb{1}\{n > \mathbf{T}\} \mathbf{Z}_n$. The Markov chain $\mathbf{Y} = (\mathbf{Y}_0, \mathbf{Y}_1, \dots)$ starts with copying \mathbf{X} and switches to \mathbf{Z} as soon as both equal the first time. We are interested in the distribution of \mathbf{Y} . To this end, consider an arbitrary path $s_{1:n} \in \mathcal{S}^n$, define $\mathbf{p}_{s_{j:i}} = \prod_{\ell=j}^i \mu_{s_{\ell-1} s_\ell}$ for $j, i = 1, \dots, n$ and observe that

$$\begin{aligned} \mathbb{P}(\mathbf{Y}_{0:n} = s_{0:n}) &= \mathbb{P}(\mathbf{Y}_{0:n} = s_{0:n}, n < \mathbf{T}) + \sum_{\ell=0}^n \mathbb{P}(\mathbf{Y}_{0:n} = s_{0:n}, \mathbf{T} = \ell) \\ &= \mathbf{p}_{s_{1:n}} \mathbb{P}(\mathbf{Z}_{0:n} \neq s_{0:n}) + \sum_{\ell=0}^n \mathbf{p}_{s_{1:\ell}} \mathbb{P}(\mathbf{Z}_\ell = s_\ell, \mathbf{Z}_{0:\ell-1} \neq s_{0:\ell-1}) \mathbf{p}_{s_{\ell+1:n}} = \mathbf{p}_{s_{1:n}} \end{aligned}$$

whereby we used that $\mathbb{P}(\mathbf{Z}_{0:n} \neq s_{0:n}) + \sum_{\ell=0}^n \mathbb{P}(\mathbf{Z}_\ell = s_\ell, \mathbf{Z}_{0:\ell-1} \neq s_{0:\ell-1}) = 1$. This shows that \mathbf{Y} is a Markov chain with transition kernel μ and initial state s_0 .

Altogether, we may state that

$$\begin{aligned} (\mu^n)_{s_0 s} &= \mathbb{P}(\mathbf{Y}_n = s) = \mathbb{P}(\mathbf{Y}_n = s, n \leq \mathbf{T}) + \mathbb{P}(\mathbf{Y}_n = s, n > \mathbf{T}) \\ \pi_s &= \mathbb{P}(\mathbf{Z}_n = s) = \mathbb{P}(\mathbf{Z}_n = s, n \leq \mathbf{T}) + \mathbb{P}(\mathbf{Y}_n = s, n > \mathbf{T}) \end{aligned}$$

and $|(\mu^n)_{s_0 s} - \pi_s| = |\mathbb{P}(\mathbf{Y}_n = s, n \leq \mathbf{T}) - \mathbb{P}(\mathbf{Z}_n = s, n \leq \mathbf{T})| \xrightarrow{n \rightarrow \infty} 0$ for all $s_0 \in \mathcal{S}$. \square

Given a distribution π , MCMC methods seek an aperiodic transition kernel μ with invariant distribution π . Thus, it is possible to sample approximately from π by simulating a Markov chain with transition kernel μ until equilibrium is reached to a sufficient extend. The last link within this chain is then taken as a single approximate sample from π . In particular, this procedure is independent of the state the Markov chain has started in. The pace by which equilibrium is approached is referred to as the *mixing time*.

Now, we consider a version of the well-known Perron-Frobenius theorem (Frobenius, 1912). It is usually stated in a more general context and corresponding proofs can be fairly

complicated. In turn, we provide our own convenient proof based on simple arithmetics and matrix algebra.

Theorem 2 (Perron-Frobenius Theorem). *An irreducible transition kernel μ has a unique invariant distribution π .*

Proof. Since any stochastic matrix has a right eigenvector with corresponding eigenvalue 1, it also has such a left eigenvector. In particular, any such left 1-eigenvector exhibits non-zero elements. Let x be a left 1-eigenvector of μ . If x has only non-negative or non-positive entries, we can immediately derive an invariant distribution π of μ through normalizing x , i.e. $\pi = x / \sum_{s \in \mathcal{S}} x_s$.

Assume now that x exhibits positive entries for $s \in \bar{N}$ and negative entries for $s \in N$. The following applies

$$\begin{aligned} \sum_{z \in \bar{N}} x_z \mu_{zs} + \sum_{z \in N} x_z \mu_{zs} = x_s &\Rightarrow \sum_{z \in \bar{N}} x_z \sum_{s \in \bar{N}} \mu_{zs} + \sum_{z \in N} x_z \sum_{s \in \bar{N}} \mu_{zs} = \sum_{s \in \bar{N}} x_s \\ \Leftrightarrow \underbrace{\sum_{z \in \bar{N}} x_z \sum_{s \in \bar{N}} \mu_{zs}}_{\leq 0} = \sum_{s \in \bar{N}} x_s - \sum_{z \in N} x_z \sum_{s \in \bar{N}} \mu_{zs} = \sum_{z \in \bar{N}} x_z \left(1 - \sum_{s \in \bar{N}} \mu_{zs} \right) = \underbrace{\sum_{z \in \bar{N}} x_z \sum_{s \notin \bar{N}} \mu_{zs}}_{\geq 0} \end{aligned} \quad (2)$$

Hence, the l.h.s and r.h.s. of (2) have to be zero, which implies that $\mu_{zs} = 0$ for all $z \in \bar{N}$ and $s \notin \bar{N}$. Consequently, $(\mu^n)_{zs} = 0$ for all $z \in \bar{N}$, $s \notin \bar{N}$ and $n \in \mathbb{N}$.

Since the existence of positive and negative entries implies reducibility, we conclude that irreducibility implies that any left 1-eigenvector has either solely non-positive or non-negative entries. Thus, an irreducible transition kernel μ exhibits at least one invariant distribution π .

Finally, assume that there is a second invariant distribution $\pi' \neq \pi$. In order to represent a distribution, not all components of π can be either larger or smaller than the components of π' . Thus, $\pi - \pi'$ must have positive and negative entries. However, $\pi - \pi'$ is a left 1-eigenvector of μ and thus, μ can't be irreducible, which contradicts the existence of π' . \square

The Perron-Frobenius theorem shows that invariant distributions can certainly be found for an abundance of stochastic matrices, especially for aperiodic ones. In the context of MCMC, it is, however, only a nice-to-have result and not utterly necessary. In fact, there is great freedom in choosing aperiodic transition kernels that exhibit a prespecified invariant distribution and each MCMC method provides its very own approach to do so.

4 A note on the Metropolis-Hastings acceptance probabilities for mixture spaces

4.1 Outline

This work is driven by the ubiquitous dissent over the abilities and contributions of the Metropolis-Hastings and reversible jump algorithm within the context of transdimensional sampling. In a mathematically coherent and at times somewhat journalistic elaboration, we aim to demystify this topic by taking a deeper look into the implementation of Metropolis-Hastings acceptance probabilities with regard to general mixture spaces.

Whilst unspectacular from a theoretical point of view, mixture spaces gave rise to challenging demands concerning their effective exploration. An often applied but not extensively studied tool for transitioning between distinct spaces are so-called translation functions.

We give an enlightening treatment of this topic that yields a generalization of the reversible jump algorithm and unveil a further important translation technique. Furthermore, by reconsidering the well-known Metropolis-within-Gibbs paradigm, we come across a dual strategy to develop Metropolis-Hastings sampler. We underpin our findings and compare the performances of our approaches by means of a changepoint example. Thereafter, in a more theoretical context, we revitalize the somewhat forgotten concept of maximal acceptance probabilities. This allows for an interesting classification of Metropolis-Hastings algorithms and gives further advice on their usage. A review of some errors in reasoning that have led to the aforementioned dissent concludes this whole section.

4.2 Introduction

The Metropolis-Hastings algorithm is one of the most well-known MCMC methods. It traverses through the state space by means of a user defined proposal distribution. Each proposed state undergoes an accept-reject step, which decides whether the proposed state or the previous link in the chain is chosen to be the next link. This step alone secures the

invariance of the Metropolis-Hastings Markov kernel towards the target distribution.

Originally, Metropolis-Hastings proposals were designed conveniently through kernel densities w.r.t. the same measure that runs the density of the target distribution (Metropolis et al., 1953; Hastings, 1970). In this case, the acceptance probability used in the accept-reject step is determined by the likelihood ratio of the transition in backward and forward direction. However, upcoming applications like variable selection (Mitchell and Beauchamp, 1988), point processes (Geyer and Møller, 1994) and changepoint analysis (Fearnhead, 2006; Siems et al., 2019) raised higher demands.

Common for these applications is that the elements of the state spaces are inhomogeneous in their dimension, i.e. the spaces are transdimensional. Obviously, such mixtures of different spaces are not straightforwardly accessible through standard techniques like random walk proposals. The first methods to conduct a change in dimension were plain births and deaths (Geyer and Møller, 1994). However, this is prone to ignore the relations shared among several coordinates and therewith promotes poor acceptance rates. Consequently, the exploration of the entire state space performs differently within the same and across the dimensions impairing the overall mixing time.

Subsequently, it became utterly popular to utilize functions that translate between points of different dimensions. Green (1995) pioneered this approach through the reversible jump algorithm. It was developed for purely continuous spaces and describes a particular class of proposals that act detached from the target space. As a result, the density of the target distribution and the kernel density of the proposal do not necessarily share the same underlying measure anymore. Nevertheless, Green (1995) was able to derive correct acceptance probabilities through a nifty application of the change of variable theorem.

According to Google scholar, Green (1995) has been cited over 5000 times. Unfortunately, Green (1995) and others caused a misperception about the abilities of the Metropolis-Hastings algorithm, which is nowadays ubiquitous. It is agreed among a significant number of scientific writings that the reversible jump algorithm has somehow made trans dimensional sampling possible and that other MCMC algorithms like Metropolis-Hastings are not applicable in these scenarios. Green (1995), Carlin and Chib (1995) and

Chen et al. (2012) even substantiate this wrong conclusion, which was acknowledged by Besag (2001); Waagepetersen and Sorensen (2001); Green (2003); Green and Hastie (2009); Sisson (2005); Sambridge et al. (2006) and many others. A thorough search for papers which oppose this stance indirectly, in one way or another, revealed only Geyer and Møller (1994); Tierney (1998); Andrieu et al. (2001); Godsill (2001); Jannink and Fernando (2004) and Roodaki et al. (2011).

This poses a significant division within the MCMC community, which requires a thorough treatment. To this end, we take up on the original task of exploring mixture spaces by means of the Metropolis-Hastings paradigm, whereby the very focus lies on the computation of acceptance probabilities. At the appropriate places, we will investigate the results of Tierney (1998) and Roodaki et al. (2011) and make conclusions beyond that. Furthermore, the discussion gives a more detailed consideration of the aforementioned claims of Green (1995); Carlin and Chib (1995) and Chen et al. (2012).

We will see that the foundation of the computation of acceptance probabilities in its most theoretical to most practical form is the compatibility of the involved densities towards a common underlying measure. For example, the generalized Metropolis-Hastings algorithm from Tierney (1998), which is applicable to virtually any combination of target distribution and proposal, derives the required densities through the Radon-Nikodym theorem (Nikodym, 1930). Unfortunately, the nonconstructive nature of this theorem yields only little practical value.

Therefore, subject to the aforementioned compatibility, the design of proposals comprises a trade-off between mutability and feasibility. As a consequence, each of the following methods exhibits its very own conditions and possibilities that need to be considered carefully in order to build the right Metropolis-Hastings sampler.

Similar to Green (1995), we examine so-called translation functions, which convey between pairs of spaces. However, we consider two different ways to apply them: before and after proposing new states and refer to these concepts as ad-hoc and post-hoc translations, respectively.

In accordance with the reversible jump algorithm, the main feature of post-hoc transla-

tions is the detachment of the proposal from the space the target distribution acts on. This entails an integral transformation which imposes strict conditions and requirements on the translation functions. Consequently, a sophisticated design in practical and mathematical terms is essential here.

In return, post-hoc translations support parsimony regarding the number of random variables that need to be generated. This enables very tightly adapted and even deterministic transitions. Furthermore, they grant access to certain proposals defining intractable compound distributions like convolutions, marginal distributions or factor distributions.

In contrast, ad-hoc translations support arbitrary translation functions, as long as they act on the mixture space exclusively. We may, for example, deploy partial maximizers of the likelihood. This represents a straightforward way to improve the acceptance rates of difficult transitions.

On the downside, ad-hoc translations require that the proposal kernel densities comply with the underlying measure space. Thus, the freedom in choosing translation functions comes at the prize of a smaller adaptability in terms of the generation of random states.

Whilst the name ad-hoc results from the ability to use simple and purposive translation functions, the term post-hoc refers to the moment when the translation function is applied.

The well-known Metropolis-within-Gibbs approach confines the set of possible transitions through conditioning and is therewith able to steer the exploration of the state space in specific ways. Consequently, it employs conditional densities defined on suitable measure spaces. It turns out that proposals set up directly on these measure spaces yield straightforward acceptance probabilities.

This shifts the challenges of creating a Metropolis-Hastings sampler to deriving conditional densities of the target distribution. We consider this in some way novel and innovative perspective to be dual to the traditional one that puts the definition of proposals first.

Metropolis-within-Gibbs can be applied in the usual way to reduce the number of coordinates that need to be considered and is therewith capable of relaxing the terms of ad-hoc translations. In addition, it captures far more difficult cases like so-called semi

deterministic translations, an important kind of post-hoc translations.

Our methods are scrutinized by means of a changepoint example. For this sake, we compare the acceptance rates of several birth and death like proposals. It turns out that ad-hoc proposals (combined with Metropolis-within-Gibbs) or post-hoc proposals are vital here. While the post-hoc translations are derived from delicate properties of the target distribution, the ad-hoc translations just trace for high likelihood values. Despite their strong differences, both methods achieve comparable acceptance rates.

The notion of maximal acceptance probabilities as introduced by Peskun (1973) and Tierney (1998) doesn't seem to have gained a lot of attention so far. The reason for this might be that there is usually only one choice for the acceptance probability in place. We revitalize this dusted property and show that post-hoc translations do not necessarily yield maximal acceptance probabilities. Therewith, they may sometimes differ from their maximal counterpart.

Transitions between pairs of different spaces are usually performed by separate proposals, which are combined to a mixture proposal. This allows for a clear modular design and is thus pursued primarily here.

As stated in Tierney (1998) already, acceptance probabilities can either be computed based on unique pairs of the individual proposals or from the mixture proposal as a whole. Albeit unintentionally, Roodaki et al. (2011) elaborates conditions upon when both approaches yield maximality. We will investigate their main result and put them into our context.

This work is structured as follows. At first, we consider the Metropolis-Hastings approach in applied terms in Section 4.3. This involves an introduction of the important detailed balance condition together with the primal Metropolis-Hastings algorithm. Furthermore, in Section 4.3.1 we talk about mixture spaces, mixture proposals and translations. Section 4.3.2 is concerned with the Metropolis-within-Gibbs approach. Subsequently, we elaborate a changepoint example in Section 4.4. The general consideration of Metropolis-Hastings and maximal acceptance probabilities is pursued in Section 4.5. Section 4.5.1 further transfers these observations to mixture proposals. We conclude with a discussion

and a brief literature review in Section 4.6.

4.3 The Metropolis-Hastings algorithm in applied terms

In this section, we deal with practical implementations of the Metropolis-Hastings algorithm. To this end, we develop ways to design proposals based on kernel densities w.r.t. sigma finite measure spaces. These comprise so-called ad-hoc and post-hoc translations and a new perspective for the Metropolis-within-Gibbs framework.

The Metropolis-Hastings algorithm is usually introduced on the basis of densities w.r.t. a common measure.⁶ For this sake, we are given a sigma finite measure space $(\mathcal{S}, \mathcal{A}, \lambda)$ and a probability measure π with a density \mathbf{p} such that $\pi(ds) = \mathbf{p}(s)\lambda(ds)$. The user provides a Markov kernel K from $(\mathcal{S}, \mathcal{A})$ to $(\mathcal{S}, \mathcal{A})$ in form of a kernel density $\mathbf{k}(s', s)$ such that $K(s', ds) = \mathbf{k}(s', s)\lambda(ds)$. This Markov kernel is referred to as the *proposal* and \mathbf{k} as the *proposal kernel density*.

Algorithm 1 (Metropolis-Hastings). (I) Choose an initial state $s_0 \in \mathcal{S}$ (II) In step $k = 1, 2, \dots$, given the previous state $s_{k-1} = s'$, propose a new state s according to $K(s', \cdot)$ and set $s_k = s$ with probability

$$\mathbf{a}_{s's} = \min \left\{ 1, \frac{\mathbf{p}(s)\mathbf{k}(s, s')}{\mathbf{p}(s')\mathbf{k}(s', s)} \right\} \quad (3)$$

otherwise $s_k = s'$. We set $\mathbf{a}_{s's}$ to zero wherever it is not defined.

We refer to (3) as the *acceptance probability* and to the second argument within the braces as the *acceptance ratio*.

Let further

$$\mu(s', \cdot) = \int \mathbf{a}_{s's} \delta_s(\cdot) + (1 - \mathbf{a}_{s's}) \delta_{s'}(\cdot) K(s', ds) \quad (4)$$

μ represents the Markov kernel whose application constitutes step (II) of Algorithm 1. Consequently, executing Algorithm 1 corresponds to simulating successive links of a Markov chain with initial state s_0 and Markov kernel μ .

The crucial point here is that, subject to some conditions, the unconditional distributions of these links converge in total variation to π (Tierney, 1994). Thus, after simulating the chain for a sufficiently long time, the last link can be regarded as an approximate sample of π .

In this work, we mainly focus on the most basic of these conditions, the invariance of π towards μ . This invariance can be accomplished through the acceptance probability alone. To this end, we will examine different families of proposals and how closed solutions for the acceptance probabilities are calculated. The other conditions to secure convergence rely on the specific form of μ and have to be checked rather individually from case to case (Tierney, 1994).

The required invariance proofs are facilitated by means of the following condition.

Definition 1. *The Markov kernel μ preserves the detailed balance condition w.r.t. π if*

$$\pi \otimes \mu(A \times B) = \pi \otimes \mu(B \times A)$$

for all $A, B \in \mathcal{A}$.

If μ preserves the detailed balance condition w.r.t. π , π is an invariant distribution of μ since $\pi \otimes \mu(\mathcal{S} \times \cdot) = \pi \otimes \mu(\cdot \times \mathcal{S}) = \pi$. The opposite implication does not hold in general.

Markov chains build from Markov kernels that preserve the detailed balance w.r.t. another distribution are called *reversible*. This is because, they exhibit same probabilities in forward and backward direction once a link in the chain follows the law of this distribution. Consequently, MCMC methods that preserve the detailed balance condition are also called reversible.

Lemma 1 (Hastings (1970)). *The Markov kernel μ of Equation (4) with the acceptance probability shown in (3) preserves the detailed balance w.r.t. π .*

Proof. There is nothing to prove for $A, B \in \mathcal{A}$ with $A = B$. Since $A \times B$ is the disjoint union of $(A \cap B) \times (A \cap B)$, $(A \setminus B) \times (B \setminus A)$, $(A \setminus B) \times (A \cap B)$ and $(A \cap B) \times (B \setminus A)$ we

can w.l.o.g. assume that $A \cap B = \emptyset$. We get

$$\begin{aligned} \pi \otimes \mu(A \times B) &= \int_A \mu(s', B) \pi(ds') = \int_A \int_B \mathbf{a}_{s's} \mathbf{k}(s', s) \mathbf{p}(s') \lambda(ds) \lambda(ds') \\ &\stackrel{*}{=} \int_B \int_A \mathbf{a}_{ss'} \mathbf{k}(s, s') \mathbf{p}(s) \lambda(ds') \lambda(ds) = \int_B \mu(s, A) \pi(ds) = \pi \otimes \mu(B \times A) \end{aligned}$$

whereby we have used that $\mathbf{a}_{s's} \mathbf{k}(s', s) \mathbf{p}(s') = \mathbf{a}_{ss'} \mathbf{k}(s, s') \mathbf{p}(s)$ and Fubini's theorem in $*$. \square

By looking at $\mathbf{a}_{s's}$, we see that normalization constants w.r.t. \mathbf{p} cancel out. Thus, the algorithm may deal with conditional versions of \mathbf{p} by deploying \mathbf{p} unchanged. This greatly facilitates data processing by means of conditioning.

4.3.1 Mixture state spaces and translations

Now, we elaborate sampling across alternating measurable spaces with the help of functions that translate between these spaces. We consider two naturally arising approaches, the so-called ad-hoc and post-hoc translations. Both exhibit their very own challenges and advantages in terms of feasibility and adaptability.

Let \mathcal{D} be a finite or countable set and let $(\mathcal{S}_i, \mathcal{A}_i, \lambda_i)_{i \in \mathcal{D}}$ be a family of disjoint sigma finite measure spaces. The *mixture measure space* $(\mathcal{S}, \mathcal{A}, \lambda)$ of $(\mathcal{S}_i, \mathcal{A}_i, \lambda_i)_{i \in \mathcal{D}}$ is defined through

$$\mathcal{S} = \bigcup_{i \in \mathcal{D}} \mathcal{S}_i, \quad \mathcal{A} = \sigma \left(\bigcup_{i \in \mathcal{D}} \mathcal{A}_i \right), \quad \lambda = \sum_{i \in \mathcal{D}} \lambda_i \circ \text{id}_i^{-1}$$

whereby \mathcal{A} is the smallest sigma algebra that comprises all \mathcal{A}_i 's. The sole purpose of the identities $\text{id}_i : \mathcal{S}_i \rightarrow \mathcal{S}$ with $\text{id}_i(s) = s$ is to lift states from the component spaces into the mixture space.

As before, we assume that π exhibits a density \mathbf{p} w.r.t. λ . A Metropolis-Hastings algorithm for π can now be build as usual. A natural way to set up proposals on the mixture space is by pursuing a modular design through a mixture of Markov kernels, whereby each component solely transitions within a pair of the spaces. We want to use the term *move* to feature the available transitions.

Let \mathcal{M} be a finite or countable set, i.e. the set of *moves*. In order to build a mixture proposal, we define measurable functions $\mathbf{q}_\ell : \mathcal{S} \rightarrow [0, 1]$ with $\sum_{\ell \in \mathcal{M}} \mathbf{q}_\ell(s') = 1$ for all $s' \in \mathcal{S}$. $\mathbf{q}_\ell(s')$ is the probability of choosing move ℓ while being in s' . Furthermore, each move ℓ exhibits a kernel density \mathbf{k}_ℓ that transitions between a pair of the spaces and complies with the measure of the target space.

For the sake of technical correctness, we need to lift each \mathbf{k}_ℓ into the mixture space by setting it to 0 for unsupported arguments. Therewith, the mixture proposal \mathbf{K} that operates on the mixture space $(\mathcal{S}, \mathcal{A})$ reads

$$\mathbf{K}(s', ds) = \sum_{\ell \in \mathcal{M}} \mathbf{q}_\ell(s') \mathbf{K}_\ell(s', ds)$$

with $\mathbf{K}_\ell(s', ds) = \mathbf{k}_\ell(s', s) \lambda(ds)$.

This time, the transition kernel for the Metropolis-Hastings algorithm for mixtures reads

$$\mu(s', \cdot) = \sum_{\ell \in \mathcal{M}} \mathbf{q}_\ell(s') \int \mathbf{a}_{s's}^\ell \delta_s(\cdot) + (1 - \mathbf{a}_{s's}^\ell) \delta_{s'}(\cdot) \mathbf{K}_\ell(s', ds) \quad (5)$$

Given the previous link s' , this kernel involves choosing a move ℓ according to $(\mathbf{q}_k(s'))_{k \in \mathcal{M}}$ at first, then proposing an s according to $\mathbf{K}_\ell(s', \cdot)$ and finally accepting or rejecting s through $\mathbf{a}_{s's}^\ell$.

This is in accordance with Algorithm 1, however, with the difference that we do not consider the mixture proposal as a whole. Instead, we incorporate the moves explicitly into the acceptance probability (see also Section 4.5.1).

Lemma 2. *Let $\mathbf{r}_\ell : \mathcal{M} \rightarrow \mathcal{M}$ be a bijection. The transition kernel μ of (5) with acceptance probability*

$$\mathbf{a}_{s's}^\ell = \min \left\{ 1, \frac{\mathbf{p}(s) \mathbf{k}_{\mathbf{r}_\ell}(s, s')}{\mathbf{p}(s') \mathbf{k}_\ell(s', s)} \cdot \frac{\mathbf{q}_{\mathbf{r}_\ell}(s)}{\mathbf{q}_\ell(s')} \right\}$$

preserves the detailed balance w.r.t. π .

Proof. For $A \in \mathcal{A}_j$, $B \in \mathcal{A}_i$, we get

$$\begin{aligned}
& \sum_{\ell \in \mathcal{M}} \int_A \int_B \mathbf{a}_{s's}^\ell \mathbf{p}(s') \mathbf{q}_\ell(s') \mathbf{k}_\ell(s', s) \lambda_i(ds) \lambda_j(ds') \\
&= \sum_{\ell \in \mathcal{M}} \int_B \int_A \mathbf{a}_{ss'}^{\mathbf{r}_\ell} \mathbf{p}(s) \mathbf{q}_{\mathbf{r}_\ell}(s) \mathbf{k}_{\mathbf{r}_\ell}(s, s') \lambda_j(ds') \lambda_i(ds) \\
&\stackrel{*}{=} \sum_{\ell \in \mathcal{M}} \int_B \int_A \mathbf{a}_{ss'}^\ell \mathbf{p}(s) \mathbf{q}_\ell(s) \mathbf{k}_\ell(s, s') \lambda_j(ds') \lambda_i(ds)
\end{aligned}$$

whereby in * we have used the uniqueness of \mathbf{r}_ℓ . \square

\mathbf{r}_ℓ stands for the unique backward or reverse move of ℓ . Consequently, if move ℓ transitions from \mathcal{S}_j to \mathcal{S}_i , move \mathbf{r}_ℓ should transition in the opposite direction from \mathcal{S}_i to \mathcal{S}_j . Please note that the uniqueness of \mathbf{r}_ℓ doesn't imply that the transitions performed by move ℓ can only be reversed by move \mathbf{r}_ℓ . In fact, the pairs of spaces are not supposed to be unique to the moves and the possible transitions determined by related move pairs may overlap arbitrarily.

Concerning the design of the moves, it is often not clear how to transition away from a point in one space to a point in another space. Especially random walk proposals pose a problem if there is no suitable measure of distance between the spaces available. This gives rise to two distinct approaches that employ functions as a device for translation.

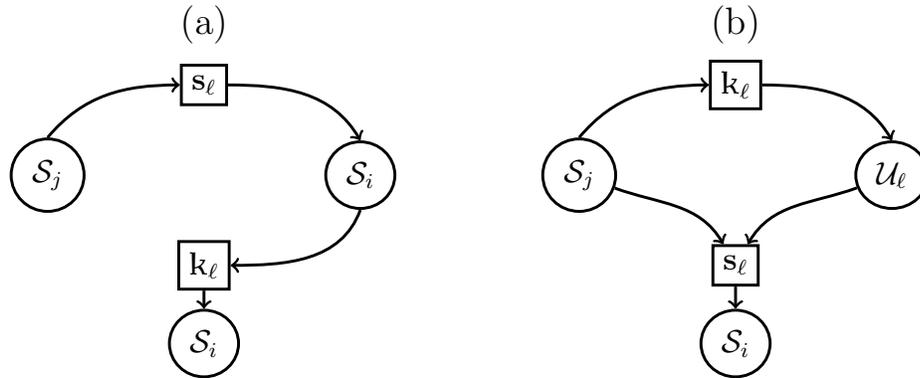


Figure 6: Transition graphs for a move $\ell \in \mathcal{M}$ that transitions from \mathcal{S}_j to \mathcal{S}_i in the fashion of (a) an ad-hoc translation and (b) a post-hoc translation

In order to transition from \mathcal{S}_j to \mathcal{S}_i in move ℓ , our first approach employs a measurable function $\mathbf{s}_\ell : \mathcal{S}_j \rightarrow \mathcal{S}_i$ that translates between the spaces and a kernel density \mathbf{k}_ℓ from $(\mathcal{S}_i, \mathcal{A}_i)$ to $(\mathcal{S}_i, \mathcal{A}_i, \lambda_i)$. Functions like \mathbf{s}_ℓ are referred to as *translation functions*. Given the previous link in the chain $s' \in \mathcal{S}_j$, new states are proposed by virtue of $\mathbf{K}_\ell(s', \cdot) = \int \cdot \mathbf{k}_\ell(\mathbf{s}_\ell(s'), s) \lambda_i(ds)$. We call this procedure *ad-hoc translation*. Figure 6(a) illustrates it graphically.

Corollary 1. *The transition kernel μ of (5) with acceptance probability*

$$\mathbf{a}_{s's}^\ell = \min \left\{ 1, \frac{\mathbf{p}(s) \mathbf{k}_{\mathbf{r}_\ell}(\mathbf{s}_{\mathbf{r}_\ell}(s), s')}{\mathbf{p}(s') \mathbf{k}_\ell(\mathbf{s}_\ell(s'), s)} \cdot \frac{\mathbf{q}_{\mathbf{r}_\ell}(s)}{\mathbf{q}_\ell(s')} \right\}$$

preserves the detailed balance w.r.t. π in ad-hoc translations.

Ad-hoc translations pose a flexible framework that is directly accessible through non-mathematicians with a basic understanding of the Metropolis-Hastings algorithm. In practice, however, another technically more demanding Metropolis-Hastings method has become the quasi gold-standard for sampling within mixed spaces. It is commonly referred to as the reversible jump algorithm (Green, 1995). We will allocate this approach to another type of translation and give a novel result that slightly generalizes this technique.

In contrast to ad-hoc translations, now, we propose first and then apply a translation to the proposed state. This allows us to detach the target space from the space where the proposal kernel density acts on. Thus, in move $\ell \in \mathcal{M}$ we may transition between a pair of the spaces, say from \mathcal{S}_j to \mathcal{S}_i , by detouring over another user-defined sigma finite measure space $(\mathcal{U}_\ell, \mathcal{B}_\ell, \nu_\ell)$. To this end, we design a kernel density \mathbf{k}_ℓ from $(\mathcal{S}_j, \mathcal{A}_j)$ to $(\mathcal{U}_\ell, \mathcal{B}_\ell, \nu_\ell)$ and a measurable function $\mathbf{s}_\ell : \mathcal{S}_j \times \mathcal{U}_\ell \rightarrow \mathcal{S}_i$. Given the previous link $s' \in \mathcal{S}_j$, a random state $u' \in \mathcal{U}_\ell$ is drawn by virtue of $\int \cdot \mathbf{k}_\ell(s', u') \nu_\ell(du')$ and the newly proposed state $s \in \mathcal{S}_i$ is obtained as $s = \mathbf{s}_\ell(s', u')$. We call this procedure *post-hoc translation*. Figure 6(b) illustrates it graphically.

The case where $(\mathcal{S}_j, \mathcal{A}_j)$ and $(\mathcal{S}_i, \mathcal{A}_i)$ are discrete can be treated straightforwardly since

the acceptance probability of move ℓ then reads

$$\mathbf{a}_{s's}^\ell = \min \left\{ 1, \frac{\mathbf{p}(s) \int \delta_{\mathbf{s}_{\mathbf{r}_\ell}(s,u)}(\{s'\}) \mathbf{k}_{\mathbf{r}_\ell}(s,u) \nu_{\mathbf{r}_\ell}(du)}{\mathbf{p}(s') \int \delta_{\mathbf{s}_\ell(s',u')}(\{s\}) \mathbf{k}_\ell(s',u') \nu_\ell(du')} \cdot \frac{\mathbf{q}_{\mathbf{r}_\ell}(s)}{\mathbf{q}_\ell(s')} \right\} \quad (6)$$

This applies, because the proposals in forward and backward direction are summarized as discrete and are thus compatible to \mathbf{p} . Though, the required feasibility of the integrals in (6) slightly restricts the set of possible post-hoc translations here.

However, most interesting cases arise when the proposal $\int \delta_{\mathbf{s}_\ell(s',u')}(\cdot) \mathbf{k}_\ell(s',u') \nu_\ell(du')$ either doesn't exhibit a kernel density w.r.t. λ_i or when it is computationally infeasible. Unfortunately, our methods so far do rely on such densities calling for a different approach here.

In practice, it would be desirable to derive the acceptance probabilities directly from the given proposal kernel densities. Thus, the trick is to accept and reject $s = \mathbf{s}_\ell(s', u')$ within $\mathcal{S}_j \times \mathcal{U}_\ell$ yielding the actual post-hoc paradigm. In order for this procedure to be well-defined, any transition $(s', u') \in \mathcal{S}_j \times \mathcal{U}_\ell$ must exhibit a unique backward transition $(s, u) \in \mathcal{S}_i \times \mathcal{U}_{\mathbf{r}_\ell}$. Thus, we need to define additional measurable functions $\mathbf{u}_\ell : \mathcal{S}_j \times \mathcal{U}_\ell \rightarrow \mathcal{U}_{\mathbf{r}_\ell}$ with

$$(\mathbf{s}, \mathbf{u})_\ell(s', u') = (s, u) \quad \Leftrightarrow \quad (\mathbf{s}, \mathbf{u})_{\mathbf{r}_\ell}(s, u) = (s', u')$$

for all $(s', u') \in \mathcal{S}_j \times \mathcal{U}_\ell$ and $(s, u) \in \mathcal{S}_i \times \mathcal{U}_{\mathbf{r}_\ell}$.

This reasoning shows that the possible options for choosing post-hoc translations this way are somewhat limited. Therewith, the seemingly restrictive conditions of the next lemma can be considered as fairly universal.

Lemma 3. *Assume that for move $\ell \in \mathcal{M}$ with transitions from \mathcal{S}_j to \mathcal{S}_i and corresponding unique backward move \mathbf{r}_ℓ , there exist measurable functions $\mathbf{u}_\ell : \mathcal{S}_j \times \mathcal{U}_\ell \rightarrow \mathcal{U}_{\mathbf{r}_\ell}$ and $f_\ell : \mathcal{S}_j \times \mathcal{U}_\ell \rightarrow \mathbb{R}_{\geq 0}$ with*

$$\text{(I)} \quad (\mathbf{s}, \mathbf{u})_\ell = (\mathbf{s}, \mathbf{u})_{\mathbf{r}_\ell}^{-1}$$

$$\text{(II)} \quad \lambda_i \otimes \nu_{\mathbf{r}_\ell}((\mathbf{s}, \mathbf{u})_\ell(ds', du')) = f_\ell(s', u') \lambda_j \otimes \nu_\ell(ds', du')$$

The transition kernel

$$\mu(s', \cdot) = \sum_{\ell \in \mathcal{M}} \mathbf{q}_\ell(s') \int \mathbf{a}_{s'u'}^\ell \delta_{\mathbf{s}_\ell}(\cdot) + (1 - \mathbf{a}_{s'u'}^\ell) \delta_{s'}(\cdot) \mathbf{k}_\ell(s', u') \nu_\ell(du')$$

with

$$\mathbf{a}_{s'u'}^\ell = \min \left\{ 1, \frac{\mathbf{p}(\mathbf{s}_\ell) \mathbf{k}_{\mathbf{r}_\ell}((\mathbf{s}, \mathbf{u})_\ell)}{\mathbf{p}(s') \mathbf{k}_\ell(s', u')} \cdot \frac{\mathbf{q}_{\mathbf{r}_\ell}(\mathbf{s}_\ell)}{\mathbf{q}_\ell(s')} \cdot f_\ell(s', u') \right\}$$

preserves the detailed balance w.r.t. π in post-hoc translations.

Proof. **(I)** implies that each $(\mathbf{s}, \mathbf{u})_\ell : \mathcal{S}_j \times \mathcal{U}_\ell \rightarrow \mathcal{S}_i \times \mathcal{U}_{\mathbf{r}_\ell}$ is a bijection with inverse $(\mathbf{s}, \mathbf{u})_{\mathbf{r}_\ell}$ and **(II)** guides through an integral transformation by means of the densities f_ℓ . Since

$$\begin{aligned} \lambda_j \otimes \nu_\ell(ds', du') &= \lambda_j \otimes \nu_\ell(\mathbf{s}_{\mathbf{r}_\ell}(\mathbf{s}_\ell(ds', du'))) \\ &= f_\ell(\mathbf{s}_\ell(ds', du')) \lambda_i \otimes \nu_{\mathbf{r}_\ell}(\mathbf{s}_\ell(ds', du')) = f_\ell(\mathbf{s}_\ell(s', u')) f_{\mathbf{r}_\ell}(s', u') \lambda_j \otimes \nu_\ell(ds', du') \end{aligned}$$

we see that $f_\ell(s', u') f_{\mathbf{r}_\ell}(s, u) = 1$ if $\mathbf{s}_\ell(s', u') = (s, u)$ and vice versa. Furthermore, for $A \in \mathcal{A}_j$ and $B \in \mathcal{B}_i$, we get

$$\begin{aligned} \pi \otimes \mathbf{K}_{\mathbf{r}_\ell}(B \times A) &= \int \delta_s(B) \delta_{\mathbf{s}_{\mathbf{r}_\ell}}(A) \mathbf{p}(s) \mathbf{k}_{\mathbf{r}_\ell}(s, u) \lambda_i \otimes \nu_{\mathbf{r}_\ell}(ds, du) \\ &= \int \delta_{\mathbf{s}_\ell}(B) \delta_{s'}(A) \mathbf{p}(\mathbf{s}_\ell) \mathbf{k}_{\mathbf{r}_\ell}((\mathbf{s}, \mathbf{u})_\ell) \lambda_i \otimes \nu_{\mathbf{r}_\ell}((\mathbf{s}, \mathbf{u})_\ell(ds', du')) \\ &= \int \delta_{s'}(A) \delta_{\mathbf{s}_\ell}(B) \mathbf{p}(\mathbf{s}_\ell) \mathbf{k}_{\mathbf{r}_\ell}((\mathbf{s}, \mathbf{u})_\ell) f_\ell(s', u') \lambda_j \otimes \nu_\ell(ds', du') \end{aligned}$$

The rest of the proof now follows the same scheme as before. \square

In mathematically precise terms, it is incorrect to refer to \mathbf{k}_ℓ as a proposal kernel density because it is generally not a kernel density of the proposal. However, in order to maintain a consistent language, we still want to keep this naming. Thus, the proof of Lemma 3 applies an integral transformation that secures the compatibility of the (joint) densities w.r.t. the transitions in forward and backward direction.

If all spaces are discrete coordinate spaces, we may choose arbitrary bijections $(\mathbf{s}, \mathbf{u})_\ell$ that meet **(I)** and apply

$$\mathbf{a}_{s'u'}^\ell = \min \left\{ 1, \frac{\mathbf{p}(\mathbf{s}_\ell) \mathbf{k}_{\mathbf{r}_\ell}((\mathbf{s}, \mathbf{u})_\ell)}{\mathbf{p}(s') \mathbf{k}_\ell(s', u')} \cdot \frac{\mathbf{q}_{\mathbf{r}_\ell}(\mathbf{s}_\ell)}{\mathbf{q}_\ell(s')} \cdot \frac{\lambda_i \otimes \nu_{\mathbf{r}_\ell}((\mathbf{s}, \mathbf{u})_\ell(ds', du'))}{\lambda_j \otimes \nu_\ell(ds', du')} \right\}$$

(Fronk and Giudici, 2004) deploy post-hoc translation in this way.

It is common to require that $\dim(\mathcal{S}_j) + \dim(\mathcal{U}_\ell) = \dim(\mathcal{S}_i) + \dim(\mathcal{U}_{r_\ell})$ if move ℓ transitions from \mathcal{S}_j to \mathcal{S}_i . Green (1995) pioneered the post-hoc translation approach and introduced this condition as *dimension matching*. His intention was to embed the spaces into a single superordinate one in order to avoid the inhomogeneity of the dimension. Around that time the nowadays ubiquitous term *transdimensional sampling* was coined.

The *reversible jump algorithm* of Green (1995) employs post-hoc translations over purely continuous coordinate spaces. It requires the dimension matching condition and that each $(\mathbf{s}, \mathbf{u})_\ell$ is a diffeomorphism with Jacobi determinant \mathbf{J}_ℓ . If **(I)** is met, the resulting acceptance probability reads

$$\mathbf{a}_{s'u'}^\ell = \min \left\{ 1, \frac{\mathbf{p}(\mathbf{s}_\ell) \mathbf{k}_{r_\ell}((\mathbf{s}, \mathbf{u})_\ell)}{\mathbf{p}(s') \mathbf{k}_\ell(s', u')} \cdot \frac{\mathbf{q}_{r_\ell}(\mathbf{s}_\ell)}{\mathbf{q}_\ell(s')} \cdot |\mathbf{J}_\ell| \right\}$$

This follows from an integration by substitution for multiple variables since $ds' du' |\mathbf{J}_\ell| = ds du$ if we substitute $(s, u) = (\mathbf{s}, \mathbf{u})_\ell(s', u')$.

Post-hoc translations are highly advanced, require very careful implementations and as it stands do only support certain kinds of state spaces. In return, they give the opportunity to generate random states within spaces of lower dimension than the target space. Therewith, they pave the way for tightly adapted, parsimonious sampling schemes where even deterministic transitions are viable.

Moreover, post-hoc translations may act on a space of higher dimension than the target space, which grants access to intractable compound distributions. To see this, consider the following minimal example to ordinarily transition from \mathbb{R} to \mathbb{R} by proposing values in \mathbb{R}^2 through the kernel density \mathbf{k} . By employing the diffeomorphism $(s', u_1, u_2) \mapsto (u_1 + u_2, s' - u_2, u_2)$ we end up with equal forward and backward moves and

$$\mathbf{a}_{s', u_1, u_2} = \min \left\{ 1, \frac{\mathbf{k}(u_1 + u_2, s' - u_2, u_2) \mathbf{p}(u_1 + u_2)}{\mathbf{k}(s', u_1, u_2) \mathbf{p}(s')} \right\}$$

The corresponding proposal on the target space, i.e. $\int \delta_{u_1+u_2}(\cdot) \mathbf{k}(s, u_1, u_2) du_1 du_2$, is the convolution of $\mathbf{k}(s, u_1, u_2)$ w.r.t u_1 and u_2 . If \mathbf{K} exhibits an intractable density on \mathbb{R} , the post-hoc paradigm proves functional here. Nevertheless, there is a little price to pay

when we use the post-hoc paradigm this way. We lose the maximality of the acceptance probability. This topic will be discussed in Section 4.5.

4.3.2 Metropolis-within-Gibbs

Imagine a coordinate state space, whereby each step of the Metropolis-Hastings algorithm solely modifies one single coordinate. This evokes the Metropolis-within-Gibbs framework. Here we investigate this approach further. We will see that a generalization allows for the identification of new measure spaces to run proposals on. This creates an innovative perspective for an important class of post-hoc translations.

Given a countable set of moves \mathcal{M} , consider a family of measurable functions $(\mathbf{g}_\ell)_{\ell \in \mathcal{M}}$ from $(\mathcal{S}, \mathcal{A})$ into a measurable space $(\mathcal{G}_\ell, \mathcal{B}_\ell)$ with $\{g\} \in \mathcal{B}_\ell$ for all $g \in \mathcal{G}_\ell$. Each \mathbf{g}_ℓ determines a partition of the state space \mathcal{S} into measurable sets that yield the same values under \mathbf{g}_ℓ . Given the previous link $s' \in \mathcal{S}$, the idea is to choose an ℓ out of \mathcal{M} and to propose a new state exclusively within the ensuing set

$$\mathfrak{R}_\ell(s') = \{s \in \mathcal{S} \mid \mathbf{g}_\ell(s) = \mathbf{g}_\ell(s')\}$$

This entails the use of conditional versions of π . For this sake, assume that \mathbf{X} is a random variable with $\mathbf{X} \sim \pi$ and $\mathbb{P}(\mathbf{X} \in ds \mid \mathbf{g}_\ell(\mathbf{X}) = g) = \mathbf{p}_\ell(g, s)\lambda_\ell(ds)$ with a measure λ_ℓ over \mathcal{A} and a kernel density \mathbf{p}_ℓ from $(\mathcal{G}_\ell, \mathcal{B}_\ell)$ to $(\mathcal{S}, \mathcal{A}, \lambda_\ell)$. Furthermore, we stipulate that for all $s' \in \mathcal{S}$, the restriction of λ_ℓ to $\mathfrak{R}_\ell(s')$ is sigma finite.

On these grounds, we employ proposal kernel densities \mathbf{k}_ℓ from $(\mathcal{S}, \mathcal{A})$ to $(\mathcal{S}, \mathcal{A}, \lambda_\ell)$ such that $\mathbf{K}_\ell(s', ds) = \mathbf{k}_\ell(s', s)\lambda_\ell(ds)$ and $\mathbf{K}_\ell(s', \mathfrak{R}_\ell(s')) = 1$ for all $\ell \in \mathcal{M}$ and $s' \in \mathcal{S}$. We further define move probabilities $\mathbf{q}_\ell : \mathcal{S} \rightarrow [0, 1]$ as before.

Lemma 4. *The transition kernel μ of (5) with acceptance probability*

$$\mathbf{a}_{s's}^\ell = \min \left\{ 1, \frac{\mathbf{p}_\ell(g, s)\mathbf{k}_\ell(s, s')}{\mathbf{p}_\ell(g, s')\mathbf{k}_\ell(s', s)} \cdot \frac{\mathbf{q}_\ell(s)}{\mathbf{q}_\ell(s')} \right\} \quad (7)$$

preserves the detailed balance w.r.t. π .

Proof. Consider the following equality

$$\pi(\cdot) = \int \int \mathbf{p}_\ell(g, s') \lambda_\ell(ds') \pi \circ \mathbf{g}_\ell^{-1}(dg)$$

For $A, B \in \mathcal{A}$ with $A \cap B = \emptyset$ we get

$$\begin{aligned} \pi \otimes \mu(A \times B) &= \int \sum_{\ell \in \mathcal{M}} \int_A \int_B \mathbf{a}_{s's}^\ell \mathbf{q}_\ell(s') \mathbf{k}_\ell(s', s) \lambda_\ell(ds) \mathbf{p}_\ell(g, s') \lambda_\ell(ds') \pi \circ \mathbf{g}_\ell^{-1}(dg) \\ &= \int \sum_{\ell \in \mathcal{M}} \int_A \int_B \mathbf{1} \{ \mathbf{g}_\ell(s') = \mathbf{g}_\ell(s) = g \} \mathbf{a}_{s's}^\ell \mathbf{q}_\ell(s') \mathbf{k}_\ell(s', s) \lambda_\ell(ds) \mathbf{p}_\ell(g, s') \lambda_\ell(ds') \pi \circ \mathbf{g}_\ell^{-1}(dg) \\ &\stackrel{*}{=} \int \sum_{\ell \in \mathcal{M}} \int_B \int_A \mathbf{a}_{ss'}^\ell \mathbf{q}_\ell(s) \mathbf{k}_\ell(s, s') \lambda_\ell(ds') \mathbf{p}_\ell(g, s) \lambda_\ell(ds) \pi \circ \mathbf{g}_\ell^{-1}(dg) = \pi \otimes \mu(B \times A) \end{aligned}$$

whereby we have used Fubini's theorem together with the partial sigma finiteness of λ_ℓ in *.

The major application scenario for Metropolis-within-Gibbs is to simplify sampling within coordinate spaces. Let $\mathcal{S} = \mathcal{S}_1 \times \dots \times \mathcal{S}_n$ be an arbitrary coordinate space together with a product sigma algebra \mathcal{A} and for $\ell = 1, \dots, n$ define $\mathbf{g}_\ell(s) = (s_1, \dots, s_{\ell-1}, s_{\ell+1}, \dots, s_n)$. Move ℓ describes transitions where all but the ℓ -th coordinate remain fixed.

Let further $\pi(ds) = \mathbf{p}(s) \nu(ds)$ with a product measure $\nu = \nu_1 \otimes \dots \otimes \nu_n$ over \mathcal{A} . We define

$$\mathbf{p}_\ell(\mathbf{g}_\ell(s'), s) = \frac{\mathbf{p}(s)}{\int \mathbf{p}(s) \nu_\ell(ds_\ell)} \mathbf{1} \{ \mathbf{g}_\ell(s') = \mathbf{g}_\ell(s) \} \quad \text{and} \quad \lambda_\ell(ds) = \nu_\ell(ds_\ell) \prod_{j \neq \ell} \#(ds_j)$$

with the counting measure $\#$. $\int \mathbf{p}(s) \nu_\ell(ds_\ell)$ is negligible here since it cancels out in (7). \mathbf{K}_ℓ is essentially build from a kernel density w.r.t. ν_ℓ . If we chose plain probabilities \mathbf{q}_ℓ for $\ell \in \mathcal{M}$, we may write $\mathbf{a}_{s's}^\ell = \min \left\{ 1, \frac{\mathbf{p}(s) \mathbf{k}_\ell(s, s')}{\mathbf{p}(s') \mathbf{k}_\ell(s', s)} \right\}$ for transitions (s', s) that differ in coordinate ℓ only.

The Metropolis-within-Gibbs approach exhibits a tremendous potential and even captures particular post-hoc translations. To see this, consider a translation function $\mathbf{s} : \mathcal{S}' \rightarrow \mathcal{S}$ between two measure spaces $(\mathcal{S}', \mathcal{A}', \lambda')$ and $(\mathcal{S}, \mathcal{A}, \lambda)$. The corresponding

post-hoc translation from \mathcal{S}' to \mathcal{S} applies \mathbf{s} in a deterministic manner. However, if \mathbf{s} is not injective, the backward move is usually not deterministic. In the following, we refer to this sort of moves as semi deterministic translations (SDT's). In practice, SDT's are the most commonly applied post-hoc translations.

Define

$$\mathbf{g}(s) = \begin{cases} \mathbf{s}(s) & s \in \mathcal{S}' \\ s & s \in \mathcal{S} \end{cases}$$

This yields Metropolis-within-Gibbs moves that mimic post-hoc translations. Even though transitions within one space and the same are not per se excluded here, we deem them as unsupported in this particular move.

In contrast to Lemma 3, \mathbf{s} is not subject to restrictions. Lemma 4 ensures that the acceptance probabilities are correct as long as we work within measure spaces that are in line with $\mathbb{P}(\mathbf{X} \in \cdot \mid \mathbf{g}(\mathbf{X}) = g)$. The following lemma represents a reconciliation of Lemma 3 and 4.

Lemma 5. *Assume that there exist a sigma finite measure space $(\mathcal{U}, \mathcal{B}, \nu)$ and measurable functions $\mathbf{u} : \mathcal{S}' \rightarrow \mathcal{U}$ and $f : \mathcal{S} \times \mathcal{U} \rightarrow \mathbb{R}_{\geq 0}$ such that*

- (I) $(\mathbf{s}, \mathbf{u}) : \mathcal{S}' \rightarrow \mathcal{S} \times \mathcal{U}$ is a bijection
- (II) $\lambda'((\mathbf{s}, \mathbf{u})^{-1}(ds, du)) = f(s, u)\lambda \otimes \nu(ds, du)$

For $s \in \mathcal{S}$, we may write

$$\mathbb{P}(\mathbf{X} \in \cdot \mid \mathbf{g}(\mathbf{X}) = s) \propto \mathbf{p}(s)\delta_s(\cdot) + \int \delta_{(\mathbf{s}, \mathbf{u})^{-1}(\cdot)} \mathbf{p}((\mathbf{s}, \mathbf{u})^{-1}) f(s, u)\nu(du) \quad (8)$$

Proof. We have to show that $\int_B \mathbb{P}(\mathbf{X} \in A \mid \mathbf{g}(\mathbf{X}) = s) \pi \circ \mathbf{g}^{-1}(ds) = \pi(A \cap \mathbf{g}^{-1}(B))$ for all

$B \in \mathcal{A}$ and $A \in \mathcal{A} \otimes \mathcal{A}'$. To this end, consider that

$$\begin{aligned}
\pi \circ \mathbf{g}^{-1} &= \int_{\cdot} \mathbf{p}(s) \lambda(ds) + \int \delta_{\mathbf{s}}(\cdot) \mathbf{p}(s') \lambda'(ds') \\
&= \int_{\cdot} \mathbf{p}(s) \lambda(ds) + \int \delta_{\mathbf{s}}(\cdot) \mathbf{p}((\mathbf{s}, \mathbf{u})^{-1}) f(u, s) \nu \otimes \lambda(du, ds) \\
&= \int_{\cdot} \mathbf{p}(s) \lambda(ds) + \underbrace{\int_{\cdot} \int \mathbf{p}((\mathbf{s}, \mathbf{u})^{-1}) f(u, s) \nu(du) \lambda(ds)}_{=\mathbf{b}(s)} = \int_{\cdot} \mathbf{p}(s) + \mathbf{b}(s) \lambda(ds)
\end{aligned}$$

This allows us to write

$$\begin{aligned}
&\int_B \mathbb{P}(\mathbf{X} \in A \mid \mathbf{g}(\mathbf{X}) = s) (\mathbf{p}(s) + \mathbf{b}(s)) \lambda(ds) \\
&\stackrel{*}{=} \int_B \mathbf{p}(s) \delta_s(A) + \int \delta_{(\mathbf{s}, \mathbf{u})^{-1}}(A) \mathbf{p}((\mathbf{s}, \mathbf{u})^{-1}) f(s, u) \nu(du) \lambda(ds) \\
&= \pi(A \cap B) + \int \delta_s(B) \delta_{(\mathbf{s}, \mathbf{u})^{-1}}(A) \mathbf{p}((\mathbf{s}, \mathbf{u})^{-1}) \lambda'((\mathbf{s}, \mathbf{u})^{-1}(ds, du)) \\
&= \pi(A \cap B) + \int \delta_{\mathbf{s}}(B) \delta_{s'}(A) \mathbf{p}(s') \lambda'(ds') = \pi(A \cap B) + \pi(A \cap \mathbf{s}^{-1}(B)) = \pi(A \cap \mathbf{g}^{-1}(B))
\end{aligned}$$

whereby in $*$, we have set the proportionality constant in (8) to $1/(\mathbf{p}(s) + \mathbf{b}(s))$. \square

Consequently, in accordance with Lemma 4, the proposal for transitioning from $s' \in \mathcal{S}'$ to \mathcal{S} complies with $\delta_{\mathbf{g}(s')}$ and thus applies \mathbf{s} deterministically. In turn, as in post-hoc translations, the backward move may exploit the form of $\int \delta_{(\mathbf{s}, \mathbf{u})^{-1}}(\cdot) \nu(du)$. To this end, we employ a kernel density \mathbf{k} from $(\mathcal{S}, \mathcal{A})$ to $(\mathcal{U}, \mathcal{B}, \nu)$. A small contemplation yields that the acceptance probability for transitions from $s' \in \mathcal{S}'$ to \mathcal{S} then reads

$$\min \left\{ 1, \frac{\mathbf{p}(s) \mathbf{k}(\mathbf{s}, \mathbf{u})}{\mathbf{p}(s') f(\mathbf{s}, \mathbf{u})} \cdot \frac{\mathbf{q}(s)}{\mathbf{q}(s')} \right\}$$

with move probability \mathbf{q} . This is the same acceptance probability as in the corresponding post-hoc approach.

Please note that the Metropolis-within-Gibbs algorithm doesn't capture all post-hoc translations. For translations that are not SDT's, there is hardly a suitable \mathbf{g} available that yields the same sampling scheme. To see this, consider the relation defined by the viable transitions under the two translation functions that belong to a move and its backward

move. \mathbf{g} is build from the partition of the two state spaces derived from the implied equivalence relation. However, for non-SDT's this partition can easily contain the spaces as a whole rendering them useless.

4.4 A changepoint example

In this section, we scrutinize our theory by means of a small changepoint example. To this end, we consider three different implementations for so-called birth and death moves which either add or remove changepoints. Firstly, we employ very plain and unsophisticated proposals. Secondly, we apply ad-hoc translations that incorporate maximizers of certain partial likelihoods. Finally, we use tightly adapted post-hoc translations in the fashion of the reversible jump algorithm.

It turns out that the ad-hoc and post-hoc approaches perform equally well on this example in terms of their computation times and acceptance rates. In contrast, the plain approach performs poorly, which justifies the need of sophisticated ad-hoc and post-hoc translations.

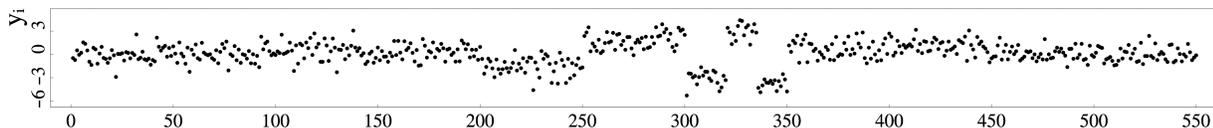


Figure 7: Gaussian change in mean example dataset.

Figure 7 shows an artificial dataset. The $n = 550$ datapoints were drawn independently from a normal distribution with variance 1 and mean values that where subject to 9 changes.

To build an exemplary Bayesian model here, we choose a prior for the changepoint locations and mean values, i.e. the segmentation and its heights. The time from one changepoint to the next is geometrically distributed with $q = 3/550$ and the segment heights in case of a jump and at the beginning are distributed according to $\mathcal{N}(0, 25)$. The

data at timepoint $i = 1, \dots, n$ also follows a normal distribution with its respective mean at i and variance equal to 1.

The sampling starts with no changepoints and an overall segment height of 0. Subsequently, new changepoint locations may be found or discarded through *birth* and *death* moves. Additionally, we employ *shift* moves to shift single changepoints and *adjust* moves to adjust single segment heights.

Each move only manipulates a subset of the segment heights or a single changepoint and leaves the rest as it is. The Metropolis-within-Gibbs framework as well as post-hoc translations support such moves and they can be applied directly without any extra effort.

Please note that the backward moves for adjust and shift are again adjust and shift. Similarly, death and birth will be reversed by birth and death, respectively. This arises naturally here, since each of the moves has its very own and unique complementary move, e.g. a birth can only be reversed by a death.

An adjust move relocates the old height h' of a uniformly chosen segment according to $\mathcal{N}(h', 0.5)$. In turn, a shift move shifts the location of a uniformly chosen changepoint uniformly to a new position within the neighboring changepoints.

Let $\phi(x; \mu, \sigma^2)$ be the density of the univariate normal distribution with mean μ and variance σ^2 evaluated at x . Adjusting the height h' of a segment S to h yields an acceptance probability of

$$\min \left\{ 1, \frac{\phi(h; 0, 25)}{\phi(h'; 0, 25)} \prod_{j \in S} \frac{\phi(y_j; h, 1)}{\phi(y_j; h', 1)} \right\}$$

Note that the probabilities of choosing the shift move, the segment and the density values of the datapoints within the untouched segments, cancel out. Furthermore, since the normal distribution is symmetric w.r.t. its mean, the proposal densities cancel out as well. What remains are the density values for the datapoints within the adjusted segment and the priors for the segment heights.

Consider 1 and $n + 1$ to be immutable auxiliary changepoints. In the shift move, we shift an existing changepoint, say at location i , uniformly to a new location, say j , within the neighboring changepoints. Let ℓ and k be the changepoints to the left and right of the

change point at i and let h_1 respectively h_2 be the corresponding segment heights. The acceptance probability for the shift move reads

$$\min \left\{ 1, \frac{\prod_{m=\ell}^{j-1} \phi(y_m ; h_1, 1) \prod_{m=j}^{k-1} \phi(y_m ; h_2, 1)}{\prod_{m=\ell}^{i-1} \phi(y_m ; h_1, 1) \prod_{m=i}^{k-1} \phi(y_m ; h_2, 1)} \right\}$$

Here, the probabilities of choosing the move and the change point, the priors for the heights and the density values of the data within the untouched segments cancel out.

4.4.1 A plain implementation of birth and death

In the death move, we choose a change point uniformly, remove it and propose a new height for the remaining segment according to $\mathcal{N}(0, 25)$. In the birth move, we propose a new change point uniformly among the timepoints without a change point. The two heights left and right of the new change point are then proposed independently according to $\mathcal{N}(0, 25)$.

Let p_d respectively p_b be the probabilities to choose a death or a birth move, respectively. We perform a death move. Therefore, let i be the timepoint of the change point that is to be removed and let ℓ and k be the change points to the left and right of the change point at i and let h_1 and h_2 be the corresponding segment heights. Finally, let h be the ensuing single segment height. The acceptance probability reads $\min \left\{ 1, \frac{n-c}{c+1} \cdot \frac{p_b}{p_d} \cdot L_{\ell,i,k} \right\}$, whereby $c + 1$ is the number of change points without 1 and $n + 1$ and

$$L_{\ell,i,k} = \frac{\prod_{m=\ell}^{k-1} \phi(y_m ; h, 1)}{\prod_{m=\ell}^{i-1} \phi(y_m ; h_1, 1) \prod_{m=i}^{k-1} \phi(y_m ; h_2, 1)}$$

In order to derive the acceptance probability for the corresponding birth move, we just take the reciprocal of the acceptance ratio.

4.4.2 An ad-hoc implementation of birth and death

Here, we utilize parts of the likelihood of the model. The likelihood of the data within a single segment can be maximized w.r.t. its segment height just by choosing the mean of

the involved datapoints. Thus, by proposing new segment heights close to these means we might propose sensible states that yield good acceptance rates. To this end, our sampling approach proposes new segment heights through a normal distribution with variance 0.01 and mean equal to the empirical mean of the associated data.

Let $\tau^2 = 0.01$ and for $1 \leq \ell \leq k \leq n$ define $\mu_{\ell:k} = \sum_{m=\ell}^{k-1} \frac{y_m}{(k-\ell)}$. By using the notation of Section 4.4.1, the acceptance probability for the death move reads

$$\min \left\{ 1, \frac{n-c}{c+1} \cdot \frac{p_b}{p_d} \cdot L_{\ell,i,k} \cdot \frac{\phi(h; 0, 25)}{\phi(h_1; 0, 25)\phi(h_2; 0, 25)} \cdot \frac{\phi(h_1; \mu_{\ell:i}, \tau^2)\phi(h_2; \mu_{i:k}, \tau^2)}{\phi(h; \mu_{\ell:k}, \tau^2)} \right\}$$

4.4.3 A post-hoc implementation of birth and death

Now, we design birth and death moves by virtue of the reversible jump paradigm. It is crucial hereby to find a carefully adapted diffeomorphism. As we have argued already, a good estimator for a segment height is the empirical mean of the observations belonging to it. If h is this mean and we split the segment, say from ℓ to k , into two parts through a changepoint at i , such that the lengths of the resulting segments are $n_1 = i - \ell$ and $n_2 = k - i + 1$ and their associated mean values are h_1 and h_2 , we see that $h = \frac{n_1 h_1 + n_2 h_2}{n_1 + n_2}$.

Furthermore, in order to build a diffeomorphism thereof, we introduce an auxiliary variable u . A sense of parsimony is achieved by stipulating that $u \in \mathbb{R}$. We treat u as part of the target space by setting $u = h_2$. That way, (h, u) describes a diffeomorphism with Jacobi determinant $\frac{n_1}{n_1 + n_2}$.

The opportunity to embed the auxiliary variables directly into the target space often remains unnoticed in practical implementations. Instead, u would be sampled independently, which is prone to forfeit acceptance rate by proposing unsuited states.

This yields an SDT, whereby we rely on the reasoning of Section 4.4.2 and draw u from $\mathcal{N}(m_{i,k}, \tau)$ in the birth move. Therewith, the acceptance probability for the death move reads $\min \left\{ 1, \frac{n_1}{n_1 + n_2} \cdot \frac{n-c}{c+1} \cdot \frac{p_b}{p_d} \cdot L_{\ell,i,k} \cdot \frac{\phi(h; 0, 25)}{\phi(h_1; 0, 25)\phi(h_2; 0, 25)} \cdot \phi(h_2, m_{i,k}, \tau) \right\}$.

4.4.4 Convergence

Here, we want to verify the convergence of our three different Markov chains towards the posterior distribution, say π , that is determined through the changepoint model and data. For this sake, we check irreducibility and aperiodicity as required by Theorem 1 in (Tierney, 1994). Additionally, Harris recurrence follows from Corollary 2 there.

We stipulate that the move probabilities are all positive. Consequently, the acceptance probabilities for all proposed transitions will as well be positive. Therefore, given an arbitrary state, in all three cases, any changepoint configuration can be occupied with positive probability through a finite amount of birth, death and shift moves.

In order to prove irreducibility, at first, we consider events, say E , whose elements all exhibit the same changepoint configuration. For E to be positive under π , the set of possible segment heights for each segment must be positive under an arbitrary univariate normal distribution. Thus, starting from a state that exhibits the very same changepoint configuration as E , transitioning into E can be done with positive probability by applying adjust moves to each segment height correspondingly. Again, this requires only a finite number of steps.

Now, take an arbitrary event that is positive under π and partition it by the different changepoint configurations present. Due to the finiteness of this partition, one of the ensuing subsets must be positive under π . Given an arbitrary state, we reach the changepoint configuration of this subset and then transition into this subset, both with positive probability through a finite amount of moves. This shows the irreducibility.

To see the aperiodicity, consider the two sets of states that can be reached through one or two successive applications of the adjust move. These sets overlap significantly, contradicting periodicity. Thus, the chain is aperiodic.

Together with the invariance, ensured by the particular forms of the acceptance probabilities, we conclude that all three Markov chains converge in total variation towards π , irrespective of the state they are started in.

4.4.5 Results and conclusions

We want to compare our proposals empirically in terms of their runtime and acceptance rates. To this end, we set all four move probabilities to 0.25, except in the boundary cases. If there is only one segment, birth and adjust have the same probability of 0.75 and if each timepoint accommodates a changepoint then death and adjust have the probabilities 0.5 respectively 0.25.

Acceptance rates	Death move	Birth move	Shift move	Adjust move
Plain proposals	0.0021	0.0022	0.0681	0.2896
Ad-hoc translations	0.0588	0.0594	0.0678	0.2904
Post-hoc translations	0.0639	0.0645	0.0681	0.2899

Table 1: Acceptance rates for particular moves.

For each of the three Markov chains, we simulated 10^7 successive links. On my computer (AMD 64 with a 3200 GHZ CPU), the sampling algorithm for the plain proposals and the post-hoc translations needed twelve seconds and the ad-hoc approach took 13 seconds.

Table 1 shows the acceptance rates for the individual moves. It becomes apparent that the plain birth and death moves struggle with finding alternate changepoint configurations. Thus, the Markov chain performs poorly in exploring the state space, which impairs the mixing time.

In contrast, our ad-hoc and post-hoc translations compare well with each other, though both exhibit fundamentally different strategies in proposing new segment heights. Several experiments revealed that an acceptance rate for birth and death moves of more than 6.5% is infeasible as long as the new changepoint locations are chosen in a plain uniform manner. Thus, the achieved rates are sound, however, higher rates are usually considered as better (Bedard, 2008).

This shows that sophisticated approaches like ad-hoc and post-hoc translations are key for sampling within mixed spaces. Whilst our post-hoc proposals can be regarded as ambitiously developed, the ad-hoc proposals constitute a simple independence sampler

that was derived from straightforward ideas. This indicates the great potential of ad-hoc translations. In turn, post-hoc translations may reduce the number of random variables to be generated and therewith result in more tightly adapted proposals that exhibit a reduced computational runtime.

4.5 The Metropolis-Hastings algorithm in abstract terms

So far, we have considered different strategies to implement the Metropolis-Hastings algorithm based on densities and measure spaces. Conversely, Tierney (1998), provides the theoretic foundations for the Metropolis-Hastings algorithm to work with arbitrary choices of π and \mathbf{K} directly. At the same time, it introduces the notion of a maximal Metropolis-Hastings algorithm. Even though this concept remained widely unrecognized, it gives important advice on the design of proposals and allows for a classification of Metropolis-Hastings approaches.

In this section, we will summarize the main outcomes of Tierney (1998) and put them in line with the results of Roodaki et al. (2011) for mixture proposals. Subsequently, we examine maximality with regards to our own methods.

Given a measurable space $(\mathcal{S}, \mathcal{A})$, for a set $A \in \mathcal{A} \otimes \mathcal{A}$ we define the *transpose* of A as $A^t = \{(s, s') \mid (s', s) \in A\}$. We leave it to the reader to show that $A^t \in \mathcal{A} \otimes \mathcal{A}$ and that $\pi \otimes \mu(\cdot) = \pi \otimes \mu(\cdot^t)$ if and only if μ preserves the detailed balance w.r.t. π .

In the following, we are given an arbitrary probability space $(\mathcal{S}, \mathcal{A}, \pi)$ and a Markov kernel \mathbf{K} from $(\mathcal{S}, \mathcal{A})$ to $(\mathcal{S}, \mathcal{A})$. The idea is to use

$$\mathbf{r}_{s's} = \frac{\pi \otimes \mathbf{K}(ds, ds')}{\pi \otimes \mathbf{K}(ds', ds)}$$

as the acceptance ratio. $\mathbf{r}_{s's}$ is the density of $\pi \otimes \mathbf{K}(\cdot^t)$ w.r.t. $\pi \otimes \mathbf{K}$ evaluated at (s', s) .

At first, we need to find a set where $\mathbf{r}_{s's}$ is guaranteed to exist. To this end, consider $\xi(\cdot) = \pi \otimes \mathbf{K}(\cdot) + \pi \otimes \mathbf{K}(\cdot^t)$. $\pi \otimes \mathbf{K}$ and $\pi \otimes \mathbf{K}(\cdot^t)$ are absolutely continuous w.r.t. ξ . Thus, the theorem of Radon-Nikodym (Nikodym, 1930) ensures the existence of $\frac{\pi \otimes \mathbf{K}(ds', ds)}{\xi(ds', ds)}$

and $\frac{\pi \otimes \mathbf{K}(ds, ds')}{\xi(ds', ds)}$. Let further

$$\mathfrak{R} = \left\{ (s', s) \mid \frac{\pi \otimes \mathbf{K}(ds', ds)}{\xi(ds', ds)} > 0, \frac{\pi \otimes \mathbf{K}(ds, ds')}{\xi(ds', ds)} > 0 \right\}$$

On \mathfrak{R} the density $\mathbf{r}_{s's}$ exists and equals the ratio of $\frac{\pi \otimes \mathbf{K}(ds, ds')}{\xi(ds', ds)}$ and $\frac{\pi \otimes \mathbf{K}(ds', ds)}{\xi(ds', ds)}$.

Lemma 6 (Tierney (1998)). *The transition kernel μ in (4) with acceptance probability*

$$\mathbf{a}_{s's} = \begin{cases} \min\{1, \mathbf{r}_{s's}\} & (s', s) \in \mathfrak{R} \\ 0 & \text{otherwise} \end{cases} \quad (9)$$

preserves the detailed balance w.r.t. π . Furthermore, $\mathbf{a}_{s's}$ is ξ almost surely maximal over all such acceptance probabilities.

Proof. By considering the symmetry of ξ and \mathfrak{R} , the detailed balance can be proven in the same fashion as Lemma 1. Furthermore, any proper acceptance probability $\hat{\mathbf{a}}_{s's}$ satisfies $\frac{\pi \otimes \mathbf{K}(ds', ds)}{\xi(ds', ds)} \hat{\mathbf{a}}_{s's} = \frac{\pi \otimes \mathbf{K}(ds, ds')}{\xi(ds', ds)} \hat{\mathbf{a}}_{ss'}$ almost surely w.r.t. ξ . Thus, on \mathfrak{R} we get $\hat{\mathbf{a}}_{s's} = \mathbf{r}_{s's} \hat{\mathbf{a}}_{ss'} \leq \min\{1, \mathbf{r}_{s's}\} = \mathbf{a}_{s's}$ and a small contemplation yields that outside of \mathfrak{R} we get $\hat{\mathbf{a}}_{s's} = 0$. \square

Lemma 6 allows us to speak of the *maximal Metropolis-Hastings algorithm* for the proposal \mathbf{K} when we use $\mathbf{a}_{s's}$ as in (9).

\mathfrak{R} is ξ almost surely uniquely defined and comprises all transitions an arbitrary Metropolis-Hastings sampler with proposal \mathbf{K} and target distribution π can carry out.

Take, for example, an SDT with translation function \mathbf{s} and a state s such that $\pi(\{s\}) = 0$ but $\pi(\mathbf{s}^{-1}(s)) > 0$. Let $T = \mathbf{s}^{-1}(s) \times \{s\}$. Since $\xi(T) > 0$ and $\pi \otimes \mathbf{K}(T^t) = 0$, we see that \mathfrak{R} and T are ξ almost surely disjoint. Thus, the transitions defined by T are not accessible for this SDT. In order to take advantage of such degenerated translation functions, we should resort to the ad-hoc paradigm instead. Section 4.4.2 provides an example for this.

4.5.1 Mixture proposals in general and maximality in particular

So far, we have build the acceptance probabilities for mixture proposals from unique pairs of forward and backward moves. However, the computation of the acceptance probability

for the maximal Metropolis-Hastings algorithm involves the mixture proposal as a whole. Consequently, if the two approaches are different, the pairwise one is prone to yield smaller acceptance rates.

Nevertheless, the pairwise approach is an important measure of convenience. It exonerates us from evaluating the full, perhaps infinite range of component proposals. This gives rise to the following lemma that provides sufficient conditions for the pairwise approach to satisfy the maximality.

Lemma 7 (Roodaki et al. (2011)). *If for all $\ell \in \mathcal{M}$, there exist disjoint $Z_\ell \in \mathcal{A} \otimes \mathcal{A}$ with*

$$\begin{aligned} \text{(I)} \quad & \int_{\mathcal{S} \times \mathcal{S} \setminus Z_\ell} \mathbf{q}_\ell(s') \pi \otimes \mathbf{K}_\ell(ds', ds) = 0 \\ \text{(II)} \quad & Z_{r_\ell} = Z_\ell^t \end{aligned}$$

we get

$$\frac{\pi \otimes \mathbf{K}(ds, ds')}{\pi \otimes \mathbf{K}(ds', ds)} = \frac{\mathbf{q}_{r_\ell}(s) \pi \otimes \mathbf{K}_{r_\ell}(ds, ds')}{\mathbf{q}_\ell(s') \pi \otimes \mathbf{K}_\ell(ds', ds)} \quad (10)$$

for $(s', s) \in Z_\ell$.

Proof. Define $\mathbf{r}_{s's}^\ell = \frac{\mathbf{q}_{r_\ell}(s) \pi \otimes \mathbf{K}_{r_\ell}(ds, ds')}{\mathbf{q}_\ell(s') \pi \otimes \mathbf{K}_\ell(ds', ds)}$. For $(s', s) \in Z_\ell$ we get

$$\mathbf{r}_{s's}^\ell \pi \otimes \mathbf{K}(ds', ds) = \mathbf{r}_{s's}^\ell \mathbf{q}_\ell(s') \pi \otimes \mathbf{K}_\ell(ds', ds) = \mathbf{q}_{r_\ell}(s) \mathbf{K}_{r_\ell}(ds, ds') = \pi \otimes \mathbf{K}(ds, ds')$$

This holds since **(I)** and **(II)** state that we can set $\mathbf{q}_k(s') \pi \otimes \mathbf{K}_k(ds', ds) = \mathbf{q}_{r_k}(s) \pi \otimes \mathbf{K}_{r_k}(ds, ds') = 0$ for all $k \neq \ell$. \square

Strangely enough, Roodaki et al. (2011) seem to make a different statement with their theorem. They point out when the acceptance probabilities can be derived from pairs of component moves. Therewith, they appear to miss that this is always viable as long as the pairs are unique, a statement that has already been known from (Tierney, 1998).

Under the assumption that the conditions **(I)** and **(II)** of Lemma 7 are met, we may ask for the maximality of the approaches considered in this work. As a rule, it is sufficient to show that for each move ℓ the acceptance ratios correspond to the r.h.s. of (10). This is

straightforward for the primal Metropolis-Hastings algorithm and ad-hoc translations, and also for (6) and (7).

However, post-hoc translations that accept within the auxiliary space are deviant. In the context of Lemma 3, we define

$$\mathbf{r}_{s'u'}^\ell = \frac{\mathbf{p}(\mathbf{s}_\ell)\mathbf{k}_{\mathbf{r}_\ell}((\mathbf{s}, \mathbf{u})_\ell)}{\mathbf{p}(s')\mathbf{k}_\ell(s', u')} \cdot f_\ell(s', u')$$

Furthermore, let $\mathbf{K}_\ell(s', \cdot) = \int \delta_{\mathbf{s}_\ell}(\cdot) \mathbf{k}_\ell(s', u') \nu_\ell(du')$ be the proposal of move ℓ .

Lemma 8. *For move ℓ with transitions from $s' \in \mathcal{S}_j$ to $s \in \mathcal{S}_i$, we get*

$$\frac{\pi \otimes \mathbf{K}_{\mathbf{r}_\ell}(ds, ds')}{\pi \otimes \mathbf{K}_\ell(ds', ds)} = \mathbb{E} \left(\mathbf{r}_{s'\mathbf{U}}^\ell \mid \mathbf{s}_\ell(s', \mathbf{U}) = s \right)$$

whereby \mathbf{U} is a random variable distributed according to $\int \cdot \mathbf{k}_\ell(s', u') \nu_\ell(du')$.

Proof. This holds since for $A \in \mathcal{A}_j$ and $B \in \mathcal{A}_i$

$$\begin{aligned} & \int_A \int_B \mathbb{E} \left(\mathbf{r}_{s'\mathbf{U}}^\ell \mid \mathbf{s}_\ell(s', \mathbf{U}) = s \right) \mathbf{K}_\ell(s', ds) \pi(ds') \\ &= \int_A \int \mathbb{E} \left(\mathbf{r}_{s'\mathbf{U}}^\ell \mid \mathbf{s}_\ell(s', \mathbf{U}) = \mathbf{s}_\ell(s', u') \right) \delta_{\mathbf{s}_\ell}(B) \mathbf{k}_\ell(s', u') \nu_\ell(du') \pi(ds') \\ &\stackrel{*}{=} \int_A \int \delta_{\mathbf{s}_\ell}(B) \mathbf{r}_{s'u'}^\ell \mathbf{k}_\ell(s', u') \mathbf{p}(s') \nu_\ell(du') \lambda(ds') \\ &= \int_A \delta_{\mathbf{s}_\ell}(B) \int_B \mathbf{p}(\mathbf{s}_\ell) \mathbf{k}_{\mathbf{r}_\ell}((\mathbf{s}, \mathbf{u})_\ell) f_\ell(s', u') \nu_\ell(du') \lambda(ds') \\ &\stackrel{**}{=} \int_B \int \delta_{\mathbf{s}_{\mathbf{r}_\ell}}(A) \mathbf{k}_{\mathbf{r}_\ell}(s, u) \mathbf{p}(s) \nu_{\mathbf{r}_\ell}(du) \lambda(ds) = \pi \otimes \mathbf{K}_{\mathbf{r}_\ell}(B \times A) \end{aligned}$$

whereby we have used the basic properties of the conditional expectation in $*$ and condition **(I)** and **(II)** of Lemma 3 in $**$. \square

Thus, the acceptance probability for move ℓ of the maximal Metropolis-Hastings algorithm for the post-hoc translations of Lemma 3 reads

$$\mathfrak{a}_{s's}^\ell = \min \left\{ 1, \frac{\mathbf{q}_{\mathbf{r}_\ell}(s)}{\mathbf{q}_\ell(s')} \cdot \mathbb{E} \left(\mathbf{r}_{s'\mathbf{U}}^\ell \mid \mathbf{s}_\ell(s', \mathbf{U}) = s \right) \right\}$$

This states that given a previous link s' , we accept or reject the newly proposed state s based on an expectation drawn from $\mathbf{r}_{s'u'}$ over those values of u' that yield $\mathbf{s}_\ell(s', u') = s$.

If $\mathbf{s}_\ell(s', \cdot)$ is injective, $\mathbf{r}_{s'u'}^\ell = \mathbb{E}\left(\mathbf{r}_{s'\mathbf{U}}^\ell \mid \mathbf{s}_\ell(s', \mathbf{U}) = \mathbf{s}_\ell(s', u')\right)$ holds and the post-hoc approach for this move is guaranteed to be maximal. This is the case for all SDT's in the fashion of Lemma 3 and 5.

On the downside, consider a transition (s', s) where the maximal Metropolis-Hastings algorithm exhibits an acceptance probability of one, i.e. $\int \mathbf{r}_{s'u'} \mathbb{P}(\mathbf{U} \in du' \mid \mathbf{s}_\ell(s', \mathbf{U}) = s) \geq 1$. The corresponding post-hoc translation approach accepts (s', s) with probability one if and only if $1 = \mathbb{P}(\mathbf{r}_{s'\mathbf{U}} \geq 1 \mid \mathbf{s}_\ell(s', \mathbf{U}) = s)$. In non-injective cases, these two conditions can be very distinct and lead to significantly different results.

4.6 Discussion and literature review

In this section, we examined several practical and theoretical approaches to derive acceptance probabilities for the Metropolis-Hastings algorithm. Due to their tricky characteristics, post-hoc translations received the most attention. In accordance with Green (1995), key to their implementation is that the accept-reject step is performed within the auxiliary space rather than the target space. As a consequence, the translation functions need to be supplemented by auxiliary functions determining for each transition its unique backward transition. On top of this bijective completion, further conditions are imposed in order to enable a certain integral transformation.

In an abstract consideration of the Metropolis-Hastings algorithm, we came across the notion of maximality as introduced by Peskun (1973) and Tierney (1998). Post-hoc translations may lose this maximality if they are not carefully implemented. This renders SDT one of the most vital post-hoc approaches. Interestingly, we found that SDT's are also accessible through the Metropolis-within-Gibbs paradigm.

Tierney (1998) was able to identify the superset of possible transitions a Metropolis-Hastings sampler is able to carry out. Therewith, we understood the limits and possibilities of SDT's better. We found that translation functions that exhibit a certain degeneracy are ineffective in post-hoc approaches, a point that underlines the utter relevance of ad-hoc translations.

All considered methods made intense use of mixture proposals. They feature a modular design and are therefore of high practical value. It is further beneficial to compute the acceptance probabilities solely based on unique pairs of moves. However, we should avoid an overlap of their supports since this may forfeit maximality.

So far, we solely considered mixture proposals w.r.t. a countable set of moves. However, the set of moves can alternatively be part of any sigma finite measure space, whereby the move probabilities are replaced by a kernel density. Consequently, summation over mixture components is replaced by integration. As before, we need to define unique pairs of forward and backward moves. An analogy to this is a post-hoc translation having an auxiliary variable that represents the move with an associated auxiliary function that determines the corresponding backward move.

Our purposive contemplation of ad-hoc as well as post-hoc translations, and Metropolis within Gibbs opens up innovative ways of choosing sampling schemes that satisfy various demands regarding simplicity and sophistication. The author of this work hopes that this helps to eliminate the ubiquitous confusions surrounding transdimensional sampling.

In the remainder of this section, we discuss a sequence of far-reaching errors in reasoning. They all share the conclusion that ordinary MCMC methods struggle with transdimensional sampling.

Green (1995) concludes wrongly that we have to use dimension matching in order to pursue transdimensional moves. On page 715 in Green (1995), the acceptance probability solely utilizes single components of the mixture proposal. This implies that the density used to perform a move must also be used for the corresponding backward move. For transdimensional moves, this is obviously not possible and made it inevitably necessary to introduce the dimension matching condition. Thus, Green's undeniably great idea appears to be the result of a simple error in reasoning.

Carlin and Chib (1995) claim, by referring to Tierney (1994), that transdimensional moves create absorbing states and therefore violate the convergence of Markov chains. Unfortunately, Tierney (1994) doesn't seem to provide this statement. In order to circumvent this putative problem, they propose to work on the product space $\prod_{i \in \mathcal{D}} \mathcal{S}_i$ instead.

Frankly speaking, this adds a huge burden just to avoid the inhomogeneous nature of the spaces. According to Google scholar, Carlin and Chib (1995) was cited over 1000 times, though a few of them question the necessity of this approach, e.g. Godsill (2001); Green (2003); Green and Hastie (2009); Green and O’Hagan (1998).

In a more measure theoretic context, Chen et al. (2012) comes to the conclusion that usual MCMC cannot transition across spaces of different dimensions. They argue with the lack of dominating measures and therewith ignore the fact that each countable collection of sigma finite measures indeed exhibits a common dominating measure.

As we have seen, the acceptance probability of the reversible jump algorithm is not necessarily maximal in the sense of Lemma 6. Interestingly, Green (1995) refers to Peskun (1973) in order highlight the maximality of his choice of acceptance probability.

These misperceptions have introduced a significant burden to the field of transdimensional sampling. The concerning literature is difficult to overview due to inconsistent claims and full of poor mathematical language due to the absurd overemphasis on dimensionality. As a result, it sometimes feels like reading some sort of star trek novel.

Finally, the name “reversible jump” apparently stems from Green’s wrong conviction that he constructed a method that overcomes the non-reversibility of transdimensional moves in the Metropolis-Hastings algorithm. Thus, this name lacks a proper meaning and collides awkwardly with “reversible MCMC” making these topics even harder to grasp for newcomers.

5 Exact inference in Bayesian changepoint models

A sequence of posterior samples obtained from a changepoint model and data allows us to derive various quantities like the expected number of changepoints, a changepoint histogram and many more. Siekmann et al. (2011) gives an impressive demonstration of what can be inferred from posterior samples based on ion channel data.

However, the number of possible changepoint configurations grows exponentially w.r.t. the data size and thus, only fractions of them can be covered by samples. This might often be sufficient as there are usually comparatively few relevant ones. Nonetheless, we can only be sure of the accuracy of our approaches if we apply exact inference algorithms, which is indeed feasible at times.

In this section, we elaborate exact inference strategies for a certain, but still broad, class of Bayesian changepoint models. A range of such algorithms already exists. However, across those papers the mathematical notations are quite diverse and in parts they convey similar contents.

The following list provides the Bayesian changepoint papers that are most relevant to this section. Fearnhead and Liu (2007); Adams and MacKay (2007) and Lai and Xing (2011), elaborate, inter alia, online changepoint detection algorithms. We can also find retrospective inference strategies like exact sampling (Fearnhead, 2006), changepoint entropy (Guédon, 2015), changepoint histograms (Nam et al., 2012; Rigaiil et al., 2012; Turner et al., 2010; Aston et al., 2012), and approaches for parameter estimation (Lai and Xing, 2011; Yildirim et al., 2013; Bansal et al., 2008).

On this basis, we will develop novel inference algorithms. Though, in order to provide a comprehensive understanding of Bayesian changepoint analysis and its capabilities, we will implement the bulk of the existing approaches into our own lightweight and simple notational framework. This framework resembles that of Fearnhead and Liu (2007) most.

Meanwhile, we will keep an eye on inference in hidden Markov models (Rabiner and Juang, 1986), hidden semi-Markov models (Yu, 2010) and the Kalman-filter (Kalman, 1960; Rauch et al., 1965). This helps clarifying that inference in Bayesian changepoint models is

closely related to inference in these more basic models. General terms for corresponding inference strategies comprise Bayes filtering and smoothing (Särkkä, 2013).

The basic innovation provided in this section concerns the development of very efficient pointwise inference. This paves the way for the computation of a wide range of pointwise expectations w.r.t. all timepoints at once with an overall linear complexity.

These pointwise statements greatly facilitate the understanding of posterior distributions and enable a novel EM algorithm (Dempster et al., 1977). This algorithm finds local maximizers of the likelihood marginalized over the unknown segmentation and is used for parameter estimation or model fitting in general. Under good conditions, each EM step exhibits a linear space and time complexity.

The reader will learn how to derive the formulas needed to build an EM algorithm based on a broad class of changepoint models. These formulas employ elementary expectations, which may also be approximated through sampling via MCMC. By this means, we are even able to maximize a pivotal likelihood locally without knowing its functional form. Approximate sampling strategies related to changepoint models via MCMC are discussed in Section 4.

This whole section is accompanied by a Laplacian change in median example. Though computationally and statistically attractive, the Laplace distribution appears to be fairly unconsidered within the Bayesian changepoint community. We make intense use of vivid visualization schemes to reveal its strengths. It turns out that the Laplace distribution can be employed efficiently and suits changepoint data that exhibits strong outliers. In our example, these outliers even appear in a highly asymmetric fashion. This poses no problem since the resulting posterior distributions address this by developing skewness.

The majority of the considered algorithms are dynamic programming algorithms (Bellman, 1966) implemented in purely discrete terms. This fact accounts for their efficiency. Under good conditions they exhibit a quadratic complexity w.r.t. the data size. By a slight approximation through a pruning approach, similar to that of Fearnhead (2006), the complexity may further reduce to linear. This renders our algorithms fit for big data purposes.

In turn, the non-discrete part of our algorithms can be traced back completely to elementary integrals. Therewith, our approach is accessible to all sorts of methods for recursive Bayesian estimation or numerical integration.

Our inference algorithms are applicable to a broad class of distributions, most notably the exponential family. However, in order to draw inference in an exact manner, we need to put limitations on the changepoint model.

We assume that the segment heights and lengths are mutually independent and that the observations are independent given the segment heights. The independence assumptions imply that changepoints partition the whole random process into independent parts. This greatly reduces the computational complexity of the algorithms.

The structure of this section is as follows. In Section 5.1, we put our Bayesian changepoint model and its random variables in concrete terms and display the corresponding graphical model. Section 5.2 develops the forward inference and discusses it in connection with the exponential family and the Laplace distribution. It further provides a pruning scheme and applies the results to the well-log data set. Section 5.3 is concerned with backward inference. Therefore, it develops algorithms to compute the MAP estimator, posterior samples and entropy. Section 5.4 elaborates novel pointwise inference and applies its results to the well-log example. Section 5.5 elaborates parameter estimation in changepoint models via the likelihood. To this end, it develops an EM algorithm and provides a thorough approach to estimate certain parameters in the well-log example. Thereafter, the EM algorithm will be scrutinized in a more frequentist fashion. Section 5.6 concludes with a discussion.

5.1 Model assumption

Let (y_1, \dots, y_n) represent our observed data and $t_1 < \dots < t_n$ the corresponding observational timepoints. For the sake of simplicity, we assume that changepoints can only occur within the set $\{t_1, \dots, t_n\}$. Consequently, we refer to these timepoints solely by indexes from 1 to n .

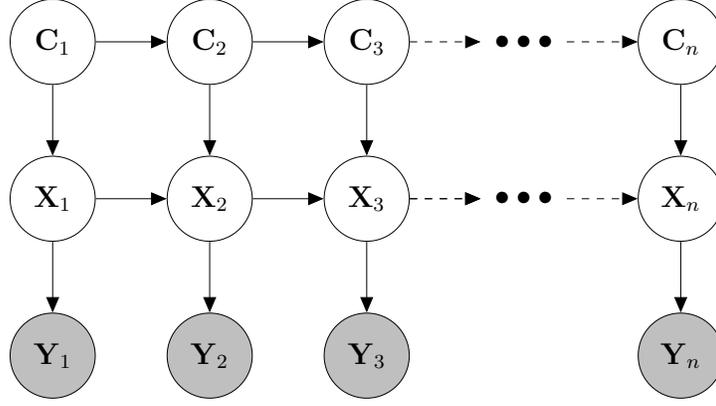


Figure 8: The Bayesian network of our changepoint model. Each node represents a random variable and the arrows determine how the model factorizes into the initial distributions. The gray nodes highlight the random variables that are observed.

For the timepoints $i = 1, \dots, n$, the position of the most recent changepoint prior to i is represented by the discrete random variable $\mathbf{C}_i \in \{0, \dots, i\}$, whereby $\mathbf{C}_i = 0$ represents the case where no changepoints occur within $\{1, \dots, i\}$.

For $0 \leq j < i \leq n$, we assume that there exist $0 \leq \mathbf{q}_{ji} \leq 1$ with

$$\mathbb{P}(\mathbf{C}_i = j) = \begin{cases} \mathbf{q}_{01} & i = 1 \text{ and } j = 0 \\ \mathbb{P}(\mathbf{C}_{i-1} = j) \mathbf{q}_{ji} & \text{otherwise} \end{cases} \quad (11)$$

\mathbf{q}_{ji} corresponds to $\mathbb{P}(\mathbf{C}_i = j \mid \mathbf{C}_{i-1} = j)$ and $\mathbb{P}(\mathbf{C}_i = i)$ ensues indirectly therefrom. The \mathbf{q}_{ji} 's can be derived from discrete length distributions as the complement of the hazard rate. Those length distributions are allowed to be non-stationary w.r.t. the time and may also depend on the distances between successive timepoints. When necessary, the \mathbf{q}_{0i} 's implement residual times.

Let $(\mathbb{X}, \mathcal{X})$ and $(\mathbb{Y}, \mathcal{Y})$ be standard Borel spaces. The segment height at i is modeled through the $(\mathbb{X}, \mathcal{X})$ -valued random variable \mathbf{X}_i with

$$(\mathbf{X}_1 \mid \mathbf{C}_1) \sim \mathcal{J} \quad \text{and} \quad (\mathbf{X}_i \mid \mathbf{C}_i = j, \mathbf{X}_{i-1} = x) \sim \begin{cases} \mathcal{J} & j = i \\ \delta_x & j < i \end{cases} \quad (12)$$

whereby \mathcal{J} is a probability measure over \mathcal{X} . \mathcal{J} may change its form w.r.t. i , but we leave this out to keep the notation simple.

Finally, the random variable \mathbf{Y}_i which is defined over the sigma finite measure space $(\mathbb{Y}, \mathcal{Y}, \psi)$ represents the observation at timepoint i . Its distribution is determined through the conditional density

$$\mathbf{p}(\mathbf{Y}_i = y \mid \mathbf{X}_i = x) = \mathbb{P}(\mathbf{Y}_i \in dy \mid \mathbf{X}_i = x) / \psi(dy)$$

for $y \in \mathbb{Y}$ and $x \in \mathbb{X}$.

We will later use these densities to derive conditional distributions based on the observations (y_1, \dots, y_n) , whereby we assume that $\int \prod_{\ell=j}^i \mathbf{p}(\mathbf{Y}_\ell = y_\ell \mid \mathbf{X}_\ell = x) \mathcal{J}(dx) > 0$ for all $0 < j \leq i \leq n$. This sort of elementary integrals are pivotal for the feasibility of the inference algorithms in this whole section.

The graphical model depicted in Figure 8 reiterates how the initial distributions are composed in order to form the complete changepoint model. It is of high value, because it allows us to quickly grasp conditional independence properties (Pearl, 1998; Bishop and Mitchell, 2014). In the following, it is assumed that the reader is capable of deducing these properties himself.

Nevertheless, context specific independence properties (Boutilier et al., 1996) result from the concrete distributions in use and can't be read from this graph. In our case, the most important is the independence induced by changepoints, i.e.

$$(\mathbf{X}_{1:i-1}, \mathbf{Y}_{1:i-1}, \mathbf{C}_{1:i-1} \perp\!\!\!\perp \mathbf{X}_{i:n}, \mathbf{Y}_{i:n}, \mathbf{C}_{i+1:n} \mid \exists \ell \geq i : \mathbf{C}_\ell = i)$$

for $i = 2, \dots, n$.

5.2 Forward inference

Forward inference processes the data in a one by one fashion, examples are the forward algorithm in hidden Markov models (Rabiner and Juang, 1986) or update and prediction in the Kalman-filter (Kalman, 1960). These algorithms can be summarized by the term

Bayes filter (Särkkä, 2013). Adams and MacKay (2007); Fearnhead and Liu (2007) and Lai and Xing (2011) elaborate forward inference in changepoint models.

It is used when data comes available in real time and needs to be processed straight away, e.g. for online detection of changepoints. In addition, it provides the foundation for the remaining inference.

The following two quantities are pivotal for forward inference. For $0 < i \leq n$ and $0 \leq j \leq i$ define

$$\mathbf{c}_{ji} = \mathbb{P}(\mathbf{C}_i = j \mid \mathbf{Y}_{1:i} = y_{1:i}) \quad (13)$$

$$\mathcal{H}_{ji} = \mathbb{P}(\mathbf{X}_i \in \cdot \mid \mathbf{C}_i = j, \mathbf{Y}_{1:i} = y_{1:i}) \quad (14)$$

\mathbf{c}_{ji} as well as \mathcal{H}_{ji} make statements concerning the i -th timepoint, whereby they incorporate all available observations hitherto. \mathbf{c}_{ji} expresses the probability that the most recent changepoint prior to i is j . \mathcal{H}_{ji} is the distribution of the segment height at i given that the segment starts at j . In both cases, the actual right end of the segment that contains i and all the subsequent observations remain unconsidered.

At first, we deal with the \mathbf{c}_{ji} 's. They rely on the important auxiliary quantities

$$\mathbf{r}_{ji} = \mathbb{P}(\mathbf{Y}_i \in dy_i \mid \mathbf{C}_i = j, \mathbf{Y}_{j:i-1} = y_{j:i-1}) / \psi(dy_i) \quad (15)$$

for $0 < i \leq n$ and $0 \leq j \leq i$. We will justify their existence and elaborate their computation in Lemma 11 and Lemma 10 later on. Given the \mathbf{r}_{ji} 's, the computation of the \mathbf{c}_{ji} 's is subject to a convenient iteration scheme as the following lemma shows.

Lemma 9. *For $0 < i \leq n$ and $0 \leq j \leq i$ we conclude that*

$$\mathbf{c}_{ji} = \frac{\mathbf{r}_{ji}}{\mathbf{Z}_i} \begin{cases} \mathbf{q}_{ji} \mathbf{c}_{ji-1} & j < i \\ 1 - \sum_{j=0}^{i-1} \mathbf{q}_{ji} \mathbf{c}_{ji-1} & j = i \end{cases}$$

$$\mathbf{Z}_i = \left(\sum_{j=0}^{i-1} \mathbf{r}_{ji} \mathbf{q}_{ji} \mathbf{c}_{ji-1} \right) + \mathbf{r}_{ii} \left(1 - \sum_{j=0}^{i-1} \mathbf{q}_{ji} \mathbf{c}_{ji-1} \right)$$

with $\mathbf{c}_{00} = 1$.

Proof. For $0 < j < i \leq n$, since $\mathbf{q}_{ji} = \mathbb{P}(\mathbf{C}_i = j \mid \mathbf{C}_{i-1} = j)$ and $(\mathbf{C}_i \perp\!\!\!\perp \mathbf{Y}_{1:i-1} \mid \mathbf{C}_{i-1})$, we see that

$$\mathbb{P}(\mathbf{C}_i = j \mid \mathbf{Y}_{1:i-1}) = \mathbb{P}(\mathbf{C}_{i-1} = j \mid \mathbf{Y}_{1:i-1}) \mathbb{P}(\mathbf{C}_i = j \mid \mathbf{C}_{i-1} = j) = \mathbf{c}_{ji-1} \mathbf{q}_{ji}$$

The proof now follows from $\mathbb{P}(\mathbf{C}_i = j, \mathbf{Y}_i \in dy_i \mid \mathbf{Y}_{1:i-1} = y_{1:i-1}) / \psi(dy_i) = \mathbf{c}_{ji-1} \mathbf{q}_{ji} \mathbf{r}_{ji}$, whereby we have used that $(\mathbf{Y}_i \perp\!\!\!\perp \mathbf{Y}_{1:j-1} \mid \mathbf{C}_i = j)$. \square

The proof of Lemma 9 reveals that the computation of \mathbf{c}_{ji} can be divided into two steps. In the first step, we go from $i-1$ to i through \mathbf{q}_{ji} without considering y_i . In the second step, we incorporate y_i through \mathbf{r}_{ji} and normalize with \mathbf{Z}_i . These two steps are an analogy of the prominent prediction and update in the Kalman-filter.

Remark 1. *The value of \mathbf{Z}_i in Lemma 9 corresponds to*

$$\mathbb{P}(\mathbf{Y}_i \in dy_i \mid \mathbf{Y}_{1:i-1} = y_{1:i-1}) / \psi(dy_i)$$

It can be used to compute the marginal loglikelihood of the observations marginalized over all segment heights and changepoint positions, i.e.

$$\ln \left(\mathbb{P}(\mathbf{Y}_{1:n} \in dy_{1:n}) / \psi(dy_{1:n}) \right) = \sum_{i=1}^n \ln(\mathbf{Z}_i)$$

Their existence follows from the existence of the \mathbf{r}_{ji} 's.

Pseudocode 1 demonstrates the resulting forward inference algorithm. Assuming that the computation of \mathbf{r}_{ji} and \mathbf{r}_{ii} in step 7 and 10 has in both cases a space and time complexity of $\mathcal{O}(h)$, the overall runtime complexity corresponds to $\mathcal{O}(hn^2)$ and the space complexity to $\mathcal{O}(n^2 + h)$.

Now, we elaborate two different approaches to compute the \mathcal{H}_{ji} 's and \mathbf{r}_{ji} 's. As we have seen in Lemma 9, \mathbf{r}_{ji} enables inference over changepoint locations. In turn, \mathcal{H}_{ji} can be used for inference over segment heights. The first approach takes advantage of certain elementary integrals, i.e. the segmental marginal likelihoods, as the following lemma shows.

Pseudocode 1 Forward inference

```

1: Input: Data  $(y_1, \dots, y_n)$ 
2: Output:  $Z_i$  and  $c_{ji}$  for  $0 < i \leq n$  and  $0 \leq j \leq i$ 
3:  $c_{00} = 1$ 
4: for  $i = 1, \dots, n$  do
5:    $\ell = 0, Z_i = 0$ 
6:   for  $j = 0, \dots, i - 1$  do
7:     Compute  $r_{ji}$  ▷ Complexity  $\mathcal{O}(h)$ 
8:      $\ell = \ell + c_{j,i-1}q_{ji}, \quad d_{ji} = c_{j,i-1}q_{ji}r_{ji}, \quad Z_i = Z_i + d_{ji}$ 
9:   end for
10:  Compute  $r_{ii}$  ▷ Complexity  $\mathcal{O}(h)$ 
11:   $d_{ii} = (1 - \ell)r_{ii}, \quad Z_i = Z_i + d_{ii}$ 
12:  for  $j = 0, \dots, i$  do
13:     $c_{ji} = d_{ji}/Z_i$ 
14:  end for
15: end for

```

Lemma 10. For $0 < j \leq i \leq n$ the following holds

$$r_{ji} = \frac{\int \prod_{\ell=j}^i \mathbf{p}(\mathbf{Y}_\ell = y_\ell \mid \mathbf{X}_\ell = x) \mathcal{J}(dx)}{\int \prod_{\ell=j}^{i-1} \mathbf{p}(\mathbf{Y}_\ell = y_\ell \mid \mathbf{X}_\ell = x) \mathcal{J}(dx)} \quad (16)$$

$$\mathcal{H}_{ji}(dx) = \frac{\prod_{\ell=j}^i \mathbf{p}(\mathbf{Y}_\ell = y_\ell \mid \mathbf{X}_\ell = x) \mathcal{J}(dx)}{\int \prod_{\ell=j}^i \mathbf{p}(\mathbf{Y}_\ell = y_\ell \mid \mathbf{X}_\ell = x) \mathcal{J}(dx)} \quad (17)$$

whereby $r_{0i} = r_{1i}$.

Proof. (16) follows from

$$r_{jj} \cdot \dots \cdot r_{ji} = \mathbb{P}(\mathbf{Y}_{j:i} = dy_{j:i} \mid \mathbf{C}_i = j) / \psi(dy_{j:i}) = \int \prod_{\ell=j}^i \mathbf{p}(\mathbf{Y}_\ell = y_\ell \mid \mathbf{X}_\ell = x) \mathcal{J}(dx)$$

for all $0 < i \leq n$ and $0 \leq j \leq i$. Hereby, we have used the requirement of Section 5.1 that $\int \prod_{\ell=j}^i \mathbf{p}(\mathbf{Y}_\ell = y_\ell \mid \mathbf{X}_\ell = x) \mathcal{J}(dx) > 0$ for all $0 < j \leq i \leq n$. This is necessary to ensure the existence of the r_{ji} 's.

Equation (17) can now be proved through induction by means of Lemma 11. In the base case, where $i = j - 1$, we set \mathcal{H}_{jj-1} to \mathcal{J} . In the induction step from $i - 1$ to i we divide $\mathbf{p}(\mathbf{Y}_i = y_i \mid \mathbf{X}_i = x) \mathcal{H}_{j,i-1}(dx)$ by r_{ji} to obtain $\mathcal{H}_{ji}(dx)$. \square

The integrals employed in Lemma 10 cannot per se be reused in an iterative fashion and might have to be computed every time anew. Since they involve all the observations from j and i , their computation then exhibits a non-constant complexity w.r.t. n increasing the complexity of the forward algorithm by at least a factor of n .

The second approach utilizes the following relation in order to compute the \mathcal{H}_{ji} 's and \mathbf{r}_{ji} 's in an iterative fashion.

Lemma 11. For $0 < i \leq n$ and $0 \leq j \leq i$

$$\mathbf{r}_{ji} \cdot \mathcal{H}_{ji}(dx) = \mathbf{p}(\mathbf{Y}_i = y_i \mid \mathbf{X}_i = x) \cdot \mathcal{H}_{ji-1}(dx) \quad (18)$$

with $\mathcal{H}_{i-1} = \mathcal{H}_{00} = \mathcal{J}$.

Proof. At first, we show that $(\mathbf{X}_i \mid \mathbf{C}_i = j, \mathbf{Y}_{1:i-1} = y_{1:i-1}) \sim \mathcal{H}_{ji-1}$ for $0 < j < i \leq n$. Therefore, consider that $(\mathbf{X}_{i-1} \perp\!\!\!\perp \mathbf{C}_i \mid \mathbf{C}_{i-1}, \mathbf{Y}_{1:i-1} = y_{1:i-1})$ and

$$\begin{aligned} (\mathbf{X}_i \mid \mathbf{C}_i = j, \mathbf{Y}_{1:i-1} = y_{1:i-1}) &\stackrel{\text{d}}{\sim} (\mathbf{X}_{i-1} \mid \mathbf{C}_i = j, \mathbf{C}_{i-1} = j, \mathbf{Y}_{1:i-1} = y_{1:i-1}) \\ &\stackrel{\text{d}}{\sim} (\mathbf{X}_{i-1} \mid \mathbf{C}_{i-1} = j, \mathbf{Y}_{1:i-1} = y_{1:i-1}) \sim \mathcal{H}_{ji-1} \end{aligned}$$

With this and $(\mathbf{Y}_i \perp\!\!\!\perp \mathbf{C}_i, \mathbf{Y}_{1:i-1} \mid \mathbf{X}_i)$, we may conclude that

$$\begin{aligned} \mathbf{p}(\mathbf{Y}_i = y_i \mid \mathbf{X}_i = x) \mathcal{H}_{ji-1}(dx) &= \mathbb{P}(\mathbf{X}_i \in dx, \mathbf{Y}_i \in dy_i \mid \mathbf{C}_i = j, \mathbf{Y}_{1:i-1} = y_{1:i-1}) / \psi(dy_i) \\ &= \mathbb{P}(\mathbf{Y}_i \in dy_i \mid \mathbf{C}_i = j, \mathbf{Y}_{1:i-1} = y_{1:i-1}) / \psi(dy_i) \cdot \mathcal{H}_{ji}(dx) = \mathbf{r}_{ji} \cdot \mathcal{H}_{ji}(dx) \end{aligned}$$

for $0 < i \leq n$ and $0 \leq j \leq i$. □

Lemma 11 reveals that \mathcal{H}_{ji} and \mathbf{r}_{ji} can be derived from $\mathbf{p}(\mathbf{Y}_i = y_i \mid \mathbf{X}_i = x)$ and \mathcal{H}_{ji-1} . For this sake, in step 7 and 10 of Pseudocode 1, we need to store \mathcal{H}_{ji} . It will then be reused to compute \mathcal{H}_{ji+1} and \mathbf{r}_{ji+1} . \mathcal{H}_{ii} and \mathbf{r}_{ii} , in turn, are derived from $\mathbf{p}(\mathbf{Y}_i = y_i \mid \mathbf{X}_i = x) \mathcal{J}(dx)$.

Many distributions, most namely the members of the exponential family, support these operations in closed form, for example, through conjugacy (Raiffa and Schlaifer, 1961). The stepwise evolution of the posterior distributions leaves their functional form invariant and only requires a transformation of the parameters. Consequently, these distributions

are stored by means of their parameters. This transformation can usually be done in constant time w.r.t. n leading to a forward algorithm with an overall quadratic complexity.

Finally, the joint distribution of the segment height at i and most recent changepoint prior to i conditioned on the observations up to i can be expressed through

$$\mathbb{P}(\mathbf{X}_i \in dx, \mathbf{C}_i = j \mid \mathbf{Y}_{1:i} = y_{1:i}) = \mathbf{c}_{ji} \mathcal{H}_{ji}(dx)$$

5.2.1 Pruning

With pruning we refer to the task of discarding certain computation steps w.r.t. j in Pseudocode 1. Having skipped the computation of, say \mathbf{c}_{ji} , from then on we will also skip all further computation steps that involve a changepoint at j , i.e. we skip the computation of $\mathbf{c}_{j\ell}$ for all $i < \ell \leq n$. In the optimal case, the computational complexity reduces thereby to linear w.r.t. n .

Imagine that at timepoint i it is certain that the most recent changepoint prior to i is j . This means that the most recent changepoint prior to i cannot be at a timepoint prior to j . Hence, all further computations involving a changepoint prior to j may safely be skipped. In this case, pruning doesn't even cause inaccuracies at all. This resembles the PELT method as discussed in Section 2.2.

However, since changepoints are normally not certain, a worthwhile pruning approach introduces inaccuracies to the distribution of changepoints. In order to keep these inaccuracies small, we develop a method that adjusts the intensity of pruning dynamically. Our approach is similar to the method proposed in Fearnhead (2006).

In accordance to the particles in particle filter approaches (Doucet et al., 2001), we call the quantity \mathbf{c}_{ji} the (j, i) -th particle.

At first, we introduce a certain minimal removal age T , i.e. particles (j, i) with $i - j < T$ must be retained. Therewith, we address the following issue. A timepoint j that doesn't appear to be the most recent changepoint prior to i can later gain more credibility. This is because inconspicuous changes in segment height require more datapoints to be processed in order to reveal their existence.

Particles whose age is at least T can be removed on the fly. Therefore, we employ another threshold T' , say $T' = 10^{-15}$. If the relative contribution of a particle to the current mass is smaller than T' , it will be removed.

Pseudocode 2 Forward inference with pruning

```

1: Input: Data  $(y_1, \dots, y_n)$  and thresholds  $T > 1$  and  $T' \geq 0$ 
2: Output:  $Z_i$  and a sparse representation of  $c_{ji}$  for  $0 < i \leq n$  and  $0 \leq j \leq i$ 
3:  $c_{00} = 1$ 
4: for  $i = 1, \dots, n$  do
5:    $\ell = 0, Z_i = 0$ 
6:   for  $j$  such that  $c_{ji-1}$  exists do
7:     Compute  $\mathbf{r}_{ji}$  ▷ Complexity  $\mathcal{O}(h)$ 
8:      $\mathbf{d}_{ji} = c_{ji-1} \mathbf{q}_{ji} \mathbf{r}_{ji}$ 
9:     if  $i - j < T$  or  $\mathbf{d}_{ji} \geq Z_i \cdot T'$  then ▷ Pruning condition
10:      Declare  $c_{ji}$  and set  $Z_i = Z_i + \mathbf{d}_{ji}$  and  $\ell = \ell + c_{ji-1} \mathbf{q}_{ji}$ 
11:    end if
12:  end for
13:  Compute  $\mathbf{r}_{ii}$  and declar  $c_{ii}$  ▷ Complexity  $\mathcal{O}(h)$ 
14:   $\mathbf{d}_{ii} = (1 - \ell) \mathbf{r}_{ii}$  and  $Z_i = Z_i + \mathbf{d}_{ii}$ 
15:  for  $j$  such that  $c_{ji}$  exists do
16:     $c_{ji} = \mathbf{d}_{ji} / Z_i$ 
17:  end for
18: end for

```

Pseudocode 2 shows the pseudocode of the pruned forward inference algorithm. The pruning conditions are implemented in Step 9. $\mathbf{d}_{ji} \geq Z_i \cdot T'$ states that the (j, i) -th particle is retained if the relative contribution of \mathbf{d}_{ji} to Z_i is at least T' .

Thus, our approach is capable of removing particles directly when they are dealt with justifying the term *on the fly pruning*. Therewith, Step 9 adjusts the number of pruned particles dynamically. If there is sufficient certainty over the locations of the last changepoints, earlier particles will not pass anymore. Hence, the efficiency of pruning is improved each time a new changepoint location pops up. However, it will refuse to prune if no new changepoint locations are spotted anymore.

In contrast to the non-pruned algorithm, here we use a data structure that stores exclusively the existing particles. It needs to be capable of iterating through the particles (j, i) w.r.t. j , but an index based access is not required. This can be achieved efficiently through lists or a sparse matrix implementation.

The actual complexity of the pruning scheme relies on how the number of changepoints increases with n . In the worst case, where the number of changepoints stays constant or the model is extremely uncertain about the changepoint locations, this algorithm may not improve the complexity whatsoever.

However, in practice the number of changepoints usually grows with n . In this case, if we hold j and increase i , more and more reasonable changepoint locations between j and i will pop up. This will eventually sort out j and also any timepoint before j to be the most recent changepoint prior to i .

Hence, a sufficient growth in the number of changepoints in relation to n , e.g. linear, yields an upper bound for the segment length w.r.t. the retained particles. In this case, even if the complexity of computing \mathbf{r}_{ji} and \mathcal{H}_{ji} is a function of the segment length, $i - j$, our pruning approach is able to achieve a time and space complexity of $\mathcal{O}(n)$ for Pseudocode 2.

5.2.2 The exponential family

Now, we show how to compute the \mathbf{r}_{ji} 's and \mathcal{H}_{ji} 's for a certain but rich family of distributions. This amounts to improved versions of Pseudocode 1 and 2.

Let $\mathbb{Y} \subseteq \mathbb{R}^r$ and $\mathbb{X} \subseteq \mathbb{R}^s$, and $\mathbf{f}(y | x) = \mathbf{p}(\mathbf{Y}_i = y | \mathbf{X}_i = x)$ for all $i = 1, \dots, n$. \mathbf{f} may also depend on i , but we skip this to keep the notation simple. Further assume that \mathcal{J} exhibits a density \mathbf{g} w.r.t. a measure ξ over \mathcal{X} . The functional form of \mathbf{f} and \mathbf{g} is defined through

$$\begin{aligned} \mathbf{f}(y | x) &= h(y) \exp(\eta(x)^t T(y) - A(x)) \\ \mathbf{g}(x; \nu, \chi) &= \tilde{h}(x) \exp\left(\chi^t \eta(x) - \nu A(x) - \tilde{A}(\nu, \chi)\right) \end{aligned}$$

whereby $T : \mathbb{Y} \rightarrow \mathbb{R}^u$, $\eta : \mathbb{X} \rightarrow \mathbb{R}^u$, $h : \mathbb{Y} \rightarrow \mathbb{R}$, $\tilde{h} : \mathbb{X} \rightarrow \mathbb{R}$ and $A : \mathbb{X} \rightarrow \mathbb{R}$ are measurable

functions. $\nu \in \mathbb{R}$ and $\chi \in \mathbb{R}^u$ are pivotal parameters. The functions $A : \mathbb{R}^s \rightarrow \mathbb{R}$ and $\tilde{A} : \mathbb{R}^{u+1} \rightarrow \mathbb{R}$ ensure that the densities \mathbf{f} and \mathbf{g} are correctly normalized. Additionally, we require that $h(y_i) > 0$ for all $i = 1, \dots, n$.

\mathbf{f} and \mathbf{g} are members of the exponential family w.r.t. to x and (ν, χ) , respectively. Moreover, \mathbf{g} is conjugate to \mathbf{f} . This means that applying Bayes' theorem (Bayes, 1763) to these densities results in a density that is of the same form as \mathbf{g} , but with updated parameters. More precisely

$$\begin{aligned} & \mathbf{f}(y | x) \cdot \mathbf{g}(x | \nu, \chi) \\ &= \mathbf{g}(x ; \nu + 1, T(y) + \chi) \cdot h(y) \exp\left(-\tilde{A}(\nu, \chi) + \tilde{A}(\nu + 1, T(y) + \chi)\right) \end{aligned} \quad (19)$$

Lemma 12. For $0 < j \leq i \leq n$ there exist $\chi_{ji} \in \mathbb{R}^s$ and $\nu_{ji} \in \mathbb{R}$ with

$$\mathcal{H}_{ji}(dx) = \mathbf{g}(x ; \nu_{ji}, \chi_{ji})\xi(dx)$$

These parameters can be computed recursively through $\nu_{ji} = \nu_{j,i-1} + 1$ and $\chi_{ji} = \chi_{j,i-1} + T(y_i)$, whereby $\nu_{i,i-1} = \nu_{00} = \nu$ and $\chi_{i,i-1} = \chi_{00} = \chi$. Furthermore,

$$\mathbf{r}_{ji} = h(y_i) \exp\left(-\tilde{A}(\nu_{j,i-1}, \chi_{j,i-1}) + \tilde{A}(\nu_{j,i-1} + 1, T(y_i) + \chi_{j,i-1})\right)$$

Proof. These inference formulas are a direct consequence of Lemma 11. The existence of $\mathcal{H}_{ji}(dx)/\xi(dx)$ follows from the fact that $\int \mathbf{f}(y_i | x)\mathbf{g}(x | \nu, \chi)\xi(dx) > 0$ for all $\nu \in \mathbb{R}$, $\chi \in \mathbb{R}^u$ and $i = 1, \dots, n$. \square

This allows for a recursive computation of \mathbf{r}_{ji} that can easily be incorporated into Pseudocode 1. There, we just need to modify Step 7 and 10 to compute and save ν_{ji}, χ_{ji} . This has a constant complexity w.r.t. n and thus, the overall space and time complexity of the unpruned forward algorithm becomes $\mathcal{O}(n^2)$.

There are useful summaries for concrete members of the exponential family available, e.g. Wikipedia Contributors (2018). In this context, the quantities involved in Lemma 11 receive particular names. $\mathcal{H}_{j,i-1}$ is called (conjugate) prior, \mathbf{f} likelihood, \mathcal{H}_{ji} posterior and \mathbf{r}_{ji} posterior predictive or compound.

5.2.3 The Laplacian change in median model

Let now $\mathbb{Y} = \mathbb{X} = \mathbb{R}$ and $\psi(dy) = dy$, and

$$\begin{aligned}\mathcal{J}(dx) &= \frac{1}{2\tau} \exp\left(-\frac{|x - \mu|}{\tau}\right) dx \\ \mathbf{p}(\mathbf{Y}_i = y \mid \mathbf{X}_i = x) &= \frac{1}{2\sigma} \exp\left(-\frac{|y - x|}{\sigma}\right)\end{aligned}$$

with $\sigma, \tau > 0$, $\mu \in \mathbb{R}$. Therewith, the segment heights at changepoint locations and observations follow a Laplace distribution with scale τ and σ , respectively. Each observational distribution ties its median to the corresponding segment height and the segment heights themselves employ μ as their median. Thus, the changepoint model infers changes in median, or more generally speaking, changes in location. All parameters are allowed to depend on i , but we leave this out to keep the notation simple.

\mathcal{J} effectively contributes an observation at μ to each segment, however with a particular scale τ . These two parameters should be chosen such that \mathcal{J} easily covers all likely segment heights, but also inhibits unlikely ones. Similarly, σ 's task is to promote and prevent joint accommodation of observations within single segments and across multiple segments, respectively.

Jafari et al. (2016) considers the Laplacian change in median model for the single changepoint case. Apart from that, this model seems to be disregarded within the Bayesian changepoint literature.

However, it is of interest, because it is more robust against outliers than, for example, the normal distribution, but computationally more attractive than, for example, Student's t-distribution. The Laplace distribution serves as an example for a distribution that is not part of the exponential family w.r.t. its location parameter.

Now, we elaborate an analytical way to compute the \mathbf{r}_{ji} 's and \mathcal{H}_{ji} 's in this model. Therefore, let

$$\mathbf{Z}_{ji}^m = \int x^m \exp\left(-\frac{|x - \mu|}{\tau} - \sum_{\ell=j}^i \frac{|y_\ell - x|}{\sigma}\right) dx \quad (20)$$

for $0 < j \leq i \leq n$. Lemma 10 shows that for $0 < j \leq i \leq n$

$$\mathbf{r}_{ji} = \frac{\mathbf{Z}_{ji}^0}{2\sigma\mathbf{Z}_{ji-1}^0} \quad \text{and} \quad \mathcal{H}_{ji}(dx) = \frac{1}{\mathbf{Z}_{ji}^0} \exp\left(-\frac{|x-\mu|}{\tau} - \sum_{\ell=j}^i \frac{|y_\ell-x|}{\sigma}\right) dx \quad (21)$$

whereby $\mathbf{Z}_{00}^0 = 2\tau$. Hence, the computation of the ratio \mathbf{r}_{ji} and \mathcal{H}_{ji} comes down to computing the integral in Equation (20) for $m = 0$. Key to this computation is the fact that $-\frac{|x-\mu|}{\tau} - \sum_{\ell=j}^i \frac{|y_\ell-x|}{\sigma}$ can be described through a continuous piecewise linear function.

To this end, let (z_1, \dots, z_r) be a sequence of real numbers which is sorted in ascending order and $(\sigma_1, \dots, \sigma_r)$ a sequence of positive real numbers. By decomposing each of the absolute value functions into their positive and negative parts we see that

$$-\sum_{\ell=1}^r \frac{|x-z_\ell|}{\sigma_\ell} = \begin{cases} \mathbf{d}_0 \cdot x + \mathbf{e}_0 & x < z_1 \\ \mathbf{d}_j \cdot x + \mathbf{e}_j & x \in [z_j, z_{j+1}) \\ \mathbf{d}_r \cdot x + \mathbf{e}_r & x \geq z_r \end{cases} \quad (22)$$

with

$$\mathbf{d}_j = -\sum_{\ell=1}^j 1/\sigma_\ell + \sum_{\ell=j+1}^r 1/\sigma_\ell \quad \text{and} \quad \mathbf{e}_j = \sum_{\ell=1}^j z_\ell/\sigma_\ell - \sum_{\ell=j+1}^r z_\ell/\sigma_\ell$$

Moreover, due to the continuity mentioned above, we see that $\mathbf{d}_{\ell-1}z_\ell + \mathbf{e}_{\ell-1} = \mathbf{d}_\ell z_\ell + \mathbf{e}_\ell$ for all $\ell = 1, \dots, r$.

Corollary 2. *Let (z_1, \dots, z_r) and $(\sigma_1, \dots, \sigma_r)$ be as before. For $m \in \mathbb{N}$, the following holds*

$$\int x^m \exp\left(-\sum_{\ell=1}^r \frac{|x-z_\ell|}{\sigma_\ell}\right) dx = \sum_{\ell=1}^r \exp(\mathbf{d}_\ell z_\ell + \mathbf{e}_\ell) (\mathbf{b}_{\ell-1}^\ell - \mathbf{b}_\ell^\ell) \quad (23)$$

whereby for $j = 0, \dots, r$ and $k = 1, \dots, r$ we define

$$\mathbf{b}_j^k = \begin{cases} \sum_{\ell=0}^m \frac{(-1)^\ell z_k^{m-\ell} m!}{\mathbf{d}_j^{\ell+1} (m-\ell)!} & \mathbf{d}_j \neq 0 \\ z_k^{m+1}/(m+1) & \mathbf{d}_j = 0 \end{cases}$$

Proof. Equation (23) follows from a piecewise integration w.r.t. to the partitioning of \mathbb{R} that is implied by $(-\infty, z_1, \dots, z_r, \infty)$. The \mathbf{b}_j^k 's are the result of a repeated integration by parts conducted on these pieces. \square

Corollary 2 requires a sorting of the sequence (μ, y_1, \dots, y_n) . Sorting can be done by successively adding observations. Each insertion costs $\mathcal{O}(\ln(n))$. Alternatively, we may use the approach of Weinmann et al. (2015), which was developed in another context. It utilizes a certain data structure to set up all the sorted sequences beforehand. This preprocessing step has a time complexity of $\mathcal{O}(n^2)$ instead of $\mathcal{O}(n^2 \ln(n))$. However, since this preprocessing thwarts the benefits of pruning, we won't consider it further.

The main drawback in using the Laplace distribution is that the integration concerning the (j, i) -particle involves a sum of $i - j + 2$ summands. Unfortunately, a stepwise computation in the fashion of Lemma 11 does not lead to any simplification here. This is because the piecewise linear functions as described in (22) change their appearance dramatically with new observations. Hence, in each step we need to reconsider all previous observations within the concerned segment and not just the new one. Therefore, the complexity of the unpruned forward algorithm is $\mathcal{O}(n^3 \ln(n))$ making changepoint problems with $n > 10^5$ a very challenging task.

However, the pruning scheme of Section 5.2.1 is able to reduce this complexity tremendously. If we assume that there is a constant $\ell < \infty$ with $i - j < \ell$ for each existing particle (j, i) , the time and space complexity reduces to $\mathcal{O}(n)$.

It is essential to improve the numerical properties of the integration by pulling out $\mathbf{m}_{ji} = \max_{x \in \mathbb{R}} \left\{ -\frac{|x-\mu|}{\tau} - \sum_{\ell=j}^i \frac{|y_\ell - x|}{\sigma} \right\}$ or a similar value from the exponent of the r.h.s. of Equation (23). It would then be reintroduced to \mathbf{r}_{ji} by multiplying with $\exp(\mathbf{m}_{ji-1} - \mathbf{m}_{ji})$ in Equation (21). \mathcal{H}_{ji} does not require a reintroduction of \mathbf{m}_{ji} . This is very effective in avoiding division by values close to zero.

Within the existing literature you can find several generalizations of the Laplace distribution, most notably a multivariate and an asymmetric (Kotz et al., 2012). However, these two are different to the univariate extension that is represented by \mathcal{H}_{ji} .

In order to facilitate the understanding in later chapters, we want to elaborate some of its properties briefly by means of an example. Therefore, consider the unnormalized density $\exp\left(-\sum_{i=1}^5 |x - z_i|\right)$ with $(z_1, \dots, z_5) = (-7, -5, 0, 1.2, 1.3)$.

$\sum_{i=1}^5 |x - z_i|$ is depicted in Figure 9(a). In the context of Gibbs measures (Georgii,

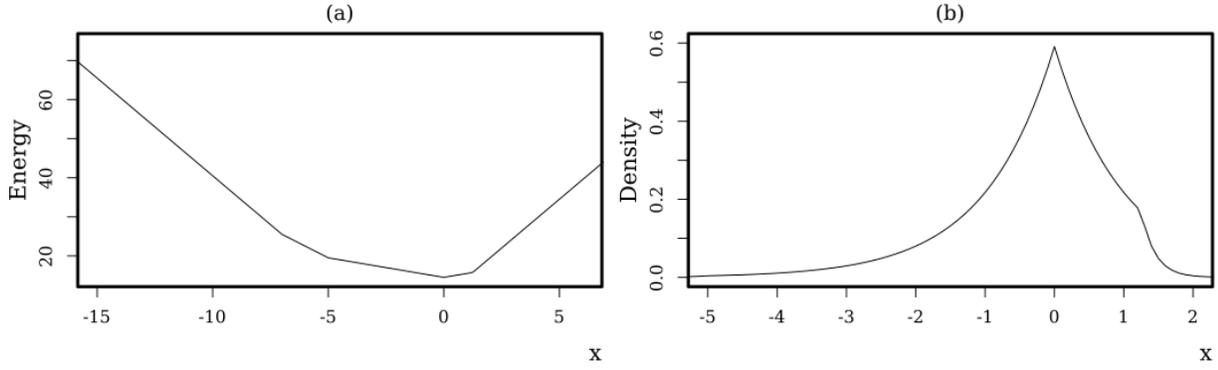


Figure 9: (a) The graph of $\sum_{i=1}^5 |x - z_i|$. (b) Normalized version of $\exp(-\sum_{i=1}^5 |x - z_i|)$.

2011), we would call it energy function. In accordance with Equation (22), it has the form of a continuous piecewise linear function being monotonically decreasing from $-\infty$ to the median of the z_i 's and monotonically increasing thereafter. Before and after the possibly non-unique medians, we even get strong monotony.

If we follow the graph from the median to minus or plus ∞ , the slope decreases or increases, respectively, each time we pass a z_i . This yields a progressive increase of energy in both directions, which is controlled by the occurrences of the z_i 's.

However, the higher the energy the lower the likelihood, i.e. the density values. The corresponding density, i.e. the normalized version of $\exp(-\sum_{i=1}^5 |x - z_i|)$, is depicted in Figure 9(b). Its single mode matches the median of z_1, \dots, z_5 . We see that the quick occurrence of z_4 and z_5 immediately starts slashing the likelihood. In turn, z_1 and z_2 are comparatively widespread, which yields a more gradual reduction of the likelihood. However, in the long run, the left and right tails decay according to the same exponential distribution.

The standard deviation w.r.t. the density is 1.0208203. Standard deviation measures the extend to which the distribution scatters around its expectation, but doesn't reveal anything about the possibly asymmetric spread.

For this purpose, we may use skewness, which amounts to -1.2708606. Its negativity

confirms the apparent tendency to produce larger values to the left than to the right.

Thus, we conclude that the \mathcal{H}_{j_i} 's define skewed distributions over \mathbb{R} . The skewness is owed by the irregular spread of the observations around their median. Moderate observations significantly slash the likelihood in the direction of extremer observations and thus, thwart their impact. $\mathcal{H}_{j_i}(dx)/dx$ can be composed continuously by exponential functions and exhibits a single mode or a single plateau without further modes.

5.2.4 Well-log example

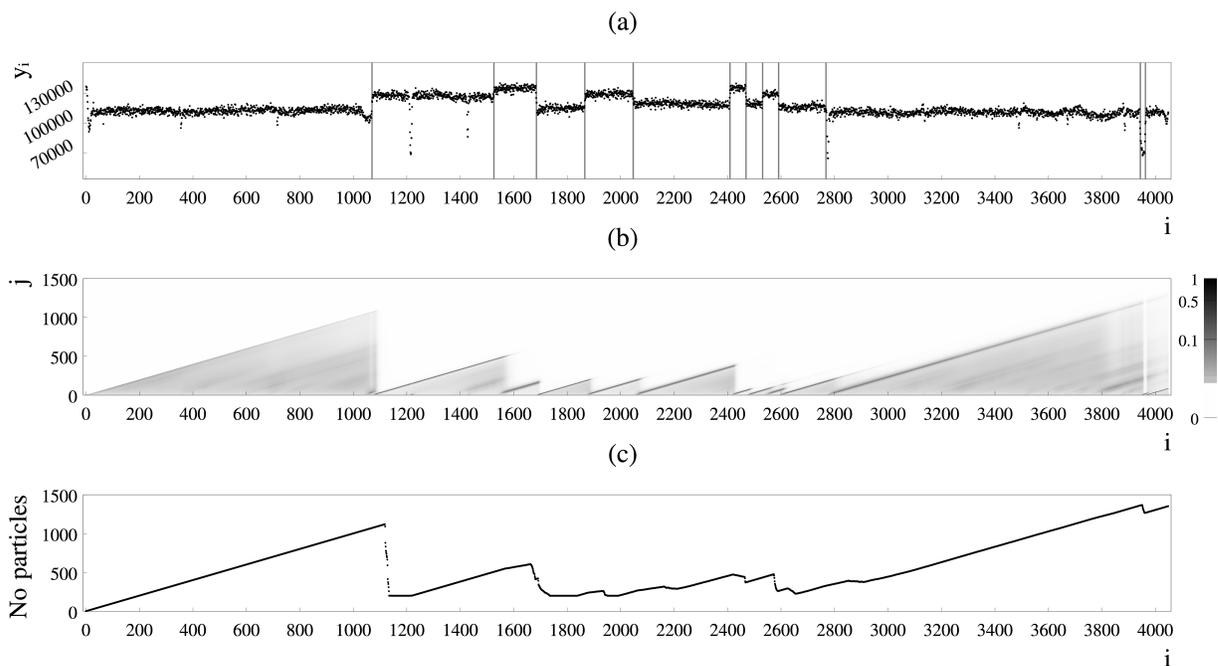


Figure 10: (a) Well-log data and changepoint estimates (through MAP estimation) as gray vertical lines. (b) Density plot of the run-lengths. For each point (i, j) it displays the value of c_{i-j_i} . (c) Illustration of the number of particles retained by the pruned forward algorithm.

Now we consider a well-log dataset (see Figure 10(a)) that stems from nuclear-magnetic response of underground rocks (Fearnhead and Clifford, 2003; Fearnhead, 2006; Fearnhead

and Rigail, 2019). Besides strong outliers, there are small and large changes in location present. Our task is to build a changepoint model that is sensitive towards the significant location changes that are marked by the gray vertical lines and specific otherwise. Therefore, we employ the Laplacian change in median model of Section 5.2.3. This example will further accompany us at several other occasions, i.e. Sections 5.2.4, 5.4.1, 5.5.3, 5.5.5 and 6.7.

We choose $\mu = 113854$, $\tau = 6879$ and $\sigma = 25000$ and the time from one changepoint to the next follows a negative binomial distribution with number of trials $r = 3$ and success probability $q = 0.01430724$. In this section, we solely discuss the values of μ and τ in more detail. The reasoning behind the other parameters requires a deeper understanding of changepoint inference and will be carried out in Section 5.5.3.

μ represents the median of all the possible segment heights. However, not all datapoints are plausible candidates, in particular those heights indicated by outliers should be excluded. Thus, the median of the whole data set, which is 113854, is a good choice here since it gives only little weight to outliers.

In a similar fashion, we choose the scale, τ , as the mean absolute deviation around the median w.r.t. the whole data set. As well as the median, it exhibits some robustness against outliers. Thus, our approach to choose μ and τ functions poses a supportive mean against strong perturbations in the data.

For reasons to be elaborated in detail in Section 5.5.2, we modify the length distribution of the first segment slightly. Let $\text{NB}(\cdot ; q, r)$ represent the probability mass function of the negative binomial distribution with success probability q and number of trials r . We assume that $q \leq r/(r + 1)$. For $0 < i \leq n$ we choose

$$\mathbf{q}_{ji} = \begin{cases} \sum_{\ell=i-j} \text{NB}(\ell ; q, r) / \sum_{\ell=i-j-1} \text{NB}(\ell ; q, r) & 0 < j < i \\ 1 - q' & j = 0 \end{cases}$$

whereby $q' = q/(r(1 - q))$. Thus, the length of the first segment now follows a geometric distribution with success probability q' .

Figure 10(b) visualizes the so-called *run-lengths*, first introduced by Adams and MacKay

(2007), as a density plot. For (i, j) it depicts the probability that the most recent changepoint prior to i is j timepoints away, that is c_{i-j} . At any timepoint i , by following the oblique lines down to zero, we may encounter the location of the most recent changepoint prior to i .

These locations usually appear slightly shifted to the right. This is because, in the forward algorithm, at each timepoint i , we only process the datapoints until i and not further. It may thus take a little while before the model pays attention to a changepoint. Our visualization of the run-lengths can be of great help when changepoints are to be detected on time.

Besides, this visualization technique can be used to get a rough picture of the models sensitivity and specificity towards changepoints at interesting spots. By this means, we may conclude that our Laplacian change in median model is pleasantly specific towards the outliers and small nuisance mean shifts present. Though, some spots, most notably the perturbation at around 3800 that takes the form of a dip, could be of concern.

In our computation we used the pruning scheme of Section 5.2.1 with $T = 200$ and $T' = 10^{-15}$ to compute the c_{ji} 's. We implemented our algorithms in C++ and the whole computation takes around 26 seconds on an AMD FX-8350 with 4.2 GHz. In comparison, for $T = 4000$ the computation takes around 361 seconds.

Figure 10(c) depicts the evolution of the particles. For each timepoint i , we see the number of particles that were retained by the pruning procedure. It shrinks dramatically in the presence of changepoints, which gives our pruning approach its convincing efficiency.

5.3 Backward inference

In the following, we consider retrospective or offline inference algorithms. Intrinsic to these algorithms is that they reprocesses the quantities inferred by the forward inference algorithm in reverse order, starting at n . Therewith they are an analogy to the backward algorithms in hidden Markov models (Rabiner and Juang, 1986).

Let

$$\begin{aligned}\tilde{\mathbf{c}}_{ji} &= \mathbb{P}(\mathbf{C}_i = j \mid \mathbf{C}_{i+1} = i + 1, \mathbf{Y}_{1:i} = y_{1:i}) && \text{for } 0 < i < n \text{ and } 0 \leq j \leq i \\ \tilde{\mathbf{c}}_{jn} &= \mathbf{c}_{jn} && \text{for } j = 0, \dots, n\end{aligned}$$

Since changepoints induce independence, we see that

$$\mathbb{P}(\mathbf{C}_i = j \mid \mathbf{C}_{i+1} = i + 1, \mathbf{Y}_{1:i} = y_{1:i}) = \mathbb{P}(\mathbf{C}_i = j \mid \mathbf{C}_{i+1} = i + 1, \mathbf{Y}_{1:n} = y_{1:n})$$

Thus, the $\tilde{\mathbf{c}}_{ji}$'s pose primal retrospective quantities that are used as a basic building block in our backward inference algorithms.

Lemma 13. *For $0 < i < n$ and $0 \leq j \leq i$ we get*

$$\tilde{\mathbf{c}}_{ji} = \frac{\mathbf{c}_{ji}(1 - \mathbf{q}_{j+1})}{\sum_{\ell=0}^i \mathbf{c}_{\ell i}(1 - \mathbf{q}_{\ell+1})}$$

Proof. Since $(\mathbf{C}_{i+1} \perp\!\!\!\perp \mathbf{Y}_{1:i} \mid \mathbf{C}_i)$ we may conclude that

$$\begin{aligned}\mathbb{P}(\mathbf{C}_i = j, \mathbf{C}_{i+1} = i + 1 \mid \mathbf{Y}_{1:i} = y_{1:i}) \\ = \mathbb{P}(\mathbf{C}_i = j \mid \mathbf{Y}_{1:i} = y_{1:i}) \mathbb{P}(\mathbf{C}_{i+1} = i + 1 \mid \mathbf{C}_i = j) = \mathbf{c}_{ji}(1 - \mathbf{q}_{j+1})\end{aligned}$$

A normalization w.r.t. i yields the proof. \square

In the pruned case, we solely use the existing particles to compute the $\tilde{\mathbf{c}}_{ji}$'s. Thereby, we declare $\tilde{\mathbf{c}}_{ji}$ only if \mathbf{c}_{ji} exists. Their computation has thus a time and space complexity which is linear in the number of particles.

Similarly, we deal with the segment height at i conditioned on the whole segment that contains i and all observations. However, their distribution does not differ from those of the segment heights with unknown right end, as the next lemma shows.

Lemma 14. *For $0 < i \leq n$ and $0 \leq j \leq i$ the following holds*

$$(\mathbf{X}_i \mid \mathbf{C}_i = j, \mathbf{C}_{i+1} = i + 1, \mathbf{Y}_{1:n} = y_{1:n}) \sim \mathcal{H}_{ji}$$

Proof. Since changepoints induce independence we get

$$(\mathbf{X}_i \mid \mathbf{C}_i = j, \mathbf{C}_{i+1} = i + 1, \mathbf{Y}_{1:n} = y_{1:n}) \stackrel{d}{\sim} (\mathbf{X}_i \mid \mathbf{C}_i = j, \mathbf{C}_{i+1} = i + 1, \mathbf{Y}_{1:i} = y_{1:i})$$

and finally $(\mathbf{X}_i \perp\!\!\!\perp \mathbf{C}_{i+1} \mid \mathbf{C}_i, \mathbf{Y}_{1:i})$ yields

$$(\mathbf{X}_i \mid \mathbf{C}_i = j, \mathbf{C}_{i+1} = i + 1, \mathbf{Y}_{1:i} = y_{1:i}) \stackrel{d}{\sim} (\mathbf{X}_i \mid \mathbf{C}_i = j, \mathbf{Y}_{1:i} = y_{1:i}) \sim \mathcal{H}_{ji}$$

□

5.3.1 Likelihood of changepoint locations and MAP estimator

The likelihood of a given set of changepoints $0 < \tau_1 < \dots < \tau_k \leq n$ conditioned on the whole set of observed data is expressed through

$$\tilde{c}_{0\tau_1-1} \left(\prod_{i=1}^{k-1} \tilde{c}_{\tau_i\tau_{i+1}-1} \right) \tilde{c}_{\tau_k n} \quad (24)$$

whereby $\tilde{c}_{00} = 1$.

The MAP (maximum a posteriori) estimator for the changepoint locations can be obtained with the Viterbi style algorithm of Section 2.2. Hereby, we would simply choose $f(y_{j:i}) = -\ln(\tilde{c}_{ji})$ for $0 \leq j \leq i \leq n$ with a dummy observation y_0 and $\gamma = 0$. This provides an element of

$$\operatorname{argmax}_{c_{1:n} \in \mathbb{N}^n} \left\{ \mathbb{P}(\mathbf{C}_{1:n} = c_{1:n} \mid \mathbf{Y}_{1:n} = y_{1:n}) \right\} \quad (25)$$

The gray vertical lines in Figure 10(a) show the result of this estimation in the well-log example.

After estimating the changepoint locations we may also estimate the corresponding segment heights by means of a Bayes or MAP estimator w.r.t. the \mathcal{H}_{ji} (see also Lemma 14).

5.3.2 Sampling

Now, we show how to obtain i.i.d. samples from $(\mathbf{C}_{1:n} \mid \mathbf{Y}_{1:n} = y_{1:n})$. Later on in Section 6 we deploy these samples for a novel approach to analyze changepoint models and changepoint data.

Pseudocode 3 Sampling

```
1: Input: A possibly sparse representation of  $\tilde{c}_{ji}$  for  $0 < i \leq n$  and  $0 \leq j \leq i$ .
2: Output: Sample  $S$ 
3:  $i = n$ 
4: while  $i > 0$  do
5:    $s = 0$ 
6:   Sample an  $u$  from the uniform distribution over  $[0, 1]$ 
7:   for  $j = i, \dots, 0$  such that  $\tilde{c}_{ji}$  exists do
8:      $s = s + \tilde{c}_{ji}$ 
9:     if  $s \geq u$  then
10:       $S \ll j$ 
11:       $i = j - 1$ 
12:      break the for loop
13:     end if
14:   end for
15: end while
```

Sampling from $(\mathbf{C}_{1:n} \mid \mathbf{Y}_{1:n} = y_{1:n})$ is done in a stepwise manner by means of the \tilde{c}_{jn} 's. We start at n and recursively obtain new changepoint locations until we arrive at 0. Pseudocode 3 shows a possible implementation of the sampling procedure also dealing with a sparse representation of the \tilde{c}_{ji} 's. This algorithm exhibits an overall complexity of $\mathcal{O}(n)$, because Step 8 can be executed $n + 1$ times at most.

Given a changepoint sample, we may further sample the segment heights conditioned on the whole data set and changepoint locations according to the \mathcal{H}_{ji} 's (see also Lemma 14). This finally yields a joined sample of the changepoint locations and segment heights conditioned on the whole data set, i.e. a sample of $(\mathbf{C}_{1:n}, \mathbf{X}_{1:n} \mid \mathbf{Y}_{1:n} = y_{1:n})$.

5.3.3 Entropy

Guédon (2015) is concerned with the entropy of changepoint locations. Entropy can be used as a measure of dispersion. We want to discuss its computation briefly. The entropy

of $(\mathbf{C}_{1:n} \mid \mathbf{Y}_{1:n} = y_{1:n})$ is defined as

$$- \sum_{c_{1:n} \in \mathbb{N}^n} \mathbb{P}(\mathbf{C}_{1:n} = c_{1:n} \mid \mathbf{Y}_{1:n} = y_{1:n}) \ln(\mathbb{P}(\mathbf{C}_{1:n} = c_{1:n} \mid \mathbf{Y}_{1:n} = y_{1:n}))$$

whereby we agree that $0 \cdot \ln(0) = 0$.

Its computation is done by a dynamic programming algorithm. Therefore let

$$\mathbf{b}(c_{1:i}) = \begin{cases} \mathbb{P}(\mathbf{C}_{1:i} = c_{1:i} \mid \mathbf{C}_{i+1} = i+1, \mathbf{Y}_{1:i} = y_{1:i}) & 0 < i < n \\ \mathbb{P}(\mathbf{C}_{1:n} = c_{1:n} \mid \mathbf{Y}_{1:n} = y_{1:n}) & i = n \end{cases}$$

and for $0 < i \leq n$ let $\mathbf{e}_i = \sum_{c_{1:i} \in \mathbb{N}^i} \mathbf{b}(c_{1:i}) \ln(\mathbf{b}(c_{1:i}))$. Ultimately, $-\mathbf{e}_n$ corresponds to the entropy of $(\mathbf{C}_{1:n} \mid \mathbf{Y}_{1:n} = y_{1:n})$. The following lemma shows that \mathbf{e}_n can be derived by computing all the \mathbf{e}_i 's one by one in a recursive fashion.

Lemma 15 (Guédon). *For $i = 1, \dots, n$ the following holds*

$$\mathbf{e}_i = \sum_{j=0}^i \tilde{\mathbf{c}}_{ji} (\mathbf{e}_{j-1} + \ln(\tilde{\mathbf{c}}_{ji}))$$

whereby we define $\mathbf{e}_{-1} = \mathbf{e}_0 = 0$.

Proof.

$$\begin{aligned} \mathbf{e}_i &= \sum_{c_{1:i} \in \mathbb{N}^i} \mathbf{b}(c_{1:i}) \ln(\mathbf{b}(c_{1:i})) = \sum_{j=0}^i \sum_{c_{1:i} \in \mathbb{N}^i \text{ with } c_i=j} \mathbf{b}(c_{1:i}) \ln(\mathbf{b}(c_{1:i})) \\ &= \tilde{\mathbf{c}}_{0i} \ln(\tilde{\mathbf{c}}_{0i}) + \tilde{\mathbf{c}}_{1i} \ln(\tilde{\mathbf{c}}_{1i}) + \underbrace{\sum_{j=2}^i \tilde{\mathbf{c}}_{ji} \sum_{c_{1:j-1} \in \mathbb{N}^{j-1}} \mathbf{b}(c_{1:j-1}) \ln(\mathbf{b}(c_{1:j-1}))}_{\mathbf{e}_{j-1}} + \tilde{\mathbf{c}}_{ji} \ln(\tilde{\mathbf{c}}_{ji}) \\ &= \sum_{j=0}^i \tilde{\mathbf{c}}_{ji} (\mathbf{e}_{j-1} + \ln(\tilde{\mathbf{c}}_{ji})) \end{aligned}$$

□

Since each $\tilde{\mathbf{c}}_{ji}$ is used only twice, the complexity of computing \mathbf{e}_n is linear in the number of particles.

5.4 Pointwise inference

In the following three sections, we utilize

$$\tilde{\mathbf{q}}_i = \mathbb{P}(\mathbf{C}_i = i \mid \mathbf{Y}_{1:n} = y_{1:n}) \quad \text{and} \quad (\mathbf{X}_i \mid \mathbf{Y}_{1:n} = y_{1:n})$$

for $i = 1, \dots, n$. These are pointwise quantities conditioned on the whole data set. The first expresses the probability of seeing a changepoint at i and the latter represents the random segment height at i . The consideration of these quantities can be referred to as smoothing (see, for example, Rauch et al. (1965) and Särkkä (2013)).

Rigaill et al. (2012) and Aston et al. (2012) employ the $\tilde{\mathbf{q}}_i$'s for model selection purposes. However, they can also be used to pursue more sophisticated tasks. We can, for example, use the standard deviation of $(\mathbf{X}_i \mid \mathbf{Y}_{1:n} = y_{1:n})$ to derive meaningful bands for the segment heights or display the trajectory of their expectations. Moreover, they pave the way for a novel EM algorithm for changepoint models (Section 5.5.1).

Lemma 16. *For $0 < i \leq n$ the following applies*

$$\tilde{\mathbf{q}}_i = \sum_{\ell=i}^{n-1} \tilde{\mathbf{c}}_{i\ell} \tilde{\mathbf{q}}_{\ell+1} + \tilde{\mathbf{c}}_{in}$$

Proof. For $0 < i < n$ we see that

$$\begin{aligned} \mathbb{P}(\mathbf{C}_i = i \mid \mathbf{Y}_{1:n} = y_{1:n}) &= \sum_{\ell=i}^{n-1} \mathbb{P}(\mathbf{C}_\ell = i, \mathbf{C}_{\ell+1} = \ell + 1 \mid \mathbf{Y}_{1:n} = y_{1:n}) + \tilde{\mathbf{c}}_{in} \\ &= \sum_{\ell=i}^{n-1} \mathbb{P}(\mathbf{C}_\ell = i \mid \mathbf{C}_{\ell+1} = \ell + 1, \mathbf{Y}_{1:n} = y_{1:n}) \tilde{\mathbf{q}}_{\ell+1} + \tilde{\mathbf{c}}_{in} = \sum_{\ell=i}^{n-1} \tilde{\mathbf{c}}_{i\ell} \tilde{\mathbf{q}}_{\ell+1} + \tilde{\mathbf{c}}_{in} \end{aligned}$$

There is nothing to show for $\tilde{\mathbf{q}}_n$. □

Remark 2. *The quantity $\tilde{\mathbf{c}}_{i\ell} \tilde{\mathbf{q}}_{\ell+1}$ represents the probability of seeing a segmentation with a segment that starts in i and ends in ℓ conditioned on the whole set of observations.*

Remark 3. $\sum_{i=1}^n \tilde{\mathbf{q}}_i$ *corresponds to the expected number of changepoints given the data.*

The computation of the $\tilde{\mathbf{q}}_i$'s yields a dynamic programming algorithm that starts with $\tilde{\mathbf{q}}_n$ and proceeds with $\tilde{\mathbf{q}}_{n-1}$ and so forth until $\tilde{\mathbf{q}}_1$ is reached. However, this algorithm requires a particular access to specific particles. Whilst in the previous algorithms, we iterate through the existing particles (j, i) w.r.t. j , now we need to iterate over i . Thus, in order to benefit from pruning here, we need to take care of this fact, e.g. by storing a second matrix or using an appropriate sparse representation. The computation of $\tilde{\mathbf{q}}_i$ for $i = 1, \dots, n$ touches each particle only once and its complexity is thus linear in the number of particles.

Lemma 17. *For $0 < i \leq n$ the following holds*

$$(\mathbf{X}_i \mid \mathbf{Y}_{1:n} = y_{1:n}) \sim \sum_{j=0}^i \mathcal{H}_{jn} \tilde{\mathbf{c}}_{jn} + \sum_{\ell=i}^{n-1} \mathcal{H}_{j\ell} \tilde{\mathbf{c}}_{j\ell} \tilde{\mathbf{q}}_{\ell+1}$$

Proof. A marginalization over all segments that contain i yields for $0 < i < n$

$$\begin{aligned} \mathbb{P}(\mathbf{X}_i \in \cdot \mid \mathbf{Y}_{1:n} = y_{1:n}) &= \sum_{j=0}^i \mathbb{P}(\mathbf{X}_i \in \cdot, \mathbf{C}_n = j \mid \mathbf{Y}_{1:n} = y_{1:n}) \\ &\quad + \sum_{\ell=i}^{n-1} \mathbb{P}(\mathbf{X}_i \in \cdot, \mathbf{C}_\ell = j, \mathbf{C}_{\ell+1} = \ell + 1 \mid \mathbf{Y}_{1:n} = y_{1:n}) \\ &\stackrel{*}{=} \sum_{j=0}^i \mathcal{H}_{jn} \tilde{\mathbf{c}}_{jn} + \sum_{\ell=i}^{n-1} \mathcal{H}_{j\ell} \tilde{\mathbf{c}}_{j\ell} \tilde{\mathbf{q}}_{\ell+1} \end{aligned}$$

In * we have used that the segment heights stay constant within a segment. \square

In the unpruned case, $\mathbb{P}(\mathbf{X}_i \in \cdot \mid \mathbf{Y}_{1:n} = y_{1:n})$ is a mixture with $\mathcal{O}(n^2)$ components. Therefore, in contrast to the marginal changepoint probabilities, it can pose a huge burden to store these quantities directly. However, Lemma 17 indicates that $\mathbb{P}(\mathbf{X}_i \in \cdot \mid \mathbf{Y}_{1:n} = y_{1:n})$ shares many components with $\mathbb{P}(\mathbf{X}_{i+1} \in \cdot \mid \mathbf{Y}_{1:n} = y_{1:n})$, a fact that can be exploited nicely.

Lemma 18. *For $0 < i < n$ the following holds*

$$\begin{aligned} \mathbb{P}(\mathbf{X}_i \in \cdot \mid \mathbf{Y}_{1:n} = y_{1:n}) &= \mathbb{P}(\mathbf{X}_{i+1} \in \cdot \mid \mathbf{Y}_{1:n} = y_{1:n}) \\ &\quad - \sum_{\ell=i+1}^{n-1} \mathcal{H}_{i+1\ell} \tilde{\mathbf{c}}_{i+1\ell} \tilde{\mathbf{q}}_{\ell+1} - \mathcal{H}_{i+1n} \tilde{\mathbf{c}}_{i+1n} + \sum_{\ell=0}^i \mathcal{H}_{\ell i} \tilde{\mathbf{c}}_{\ell i} \tilde{\mathbf{q}}_{i+1} \end{aligned}$$

The proof of Lemma 18 is straightforward. The intuition behind it is that the sets of segments that contain either i or $i + 1$, only differ in the segments that start at $i + 1$ and end at i . This is because, these are the only segments which do not include i and $i + 1$ simultaneously. These segments need to be either added or removed in order to transition from the set of segments that contain $i + 1$ to those that contain i . The initial timepoint to start with is n since $\mathbb{P}(\mathbf{X}_n \in \cdot \mid \mathbf{Y}_{1:n} = y_{1:n}) = \sum_{j=0}^n \tilde{\mathbf{c}}_{jn} \mathcal{H}_{jn}$ is already known.

In practice, we are less inclined to compute these probability distributions, we mostly utilize expectations w.r.t. to measurable functions $\mathbf{f} : \mathbb{X} \rightarrow \mathbb{R}^u$, i.e.

$$\mathbb{E}(\mathbf{f}(\mathbf{X}_i) \mid \mathbf{Y}_{1:n} = y_{1:n}) \quad (26)$$

Since the expectation w.r.t. a mixture distribution also decomposes into the sum of corresponding expectations, the basic requirement for computing (26) is the feasibility of

$$\mathcal{H}_{ji}^{\mathbf{f}} = \int \mathbf{f}(x) \mathcal{H}_{ji}(dx) \quad (27)$$

for $0 < i \leq n$ and $0 \leq j \leq i$. Together with Lemma 18, this allows for an iteration scheme that touches each of the segments only twice, once when it is added and once when it is removed.

Corollary 3. *For $0 < i < n$ the following holds*

$$\begin{aligned} & \mathbb{E}(\mathbf{f}(\mathbf{X}_i) \mid \mathbf{Y}_{1:n} = y_{1:n}) \\ &= \mathbb{E}(\mathbf{f}(\mathbf{X}_{i+1}) \mid \mathbf{Y}_{1:n} = y_{1:n}) - \sum_{\ell=i+1}^{n-1} \mathcal{H}_{i+1\ell}^{\mathbf{f}} \tilde{\mathbf{c}}_{i+1\ell} \tilde{\mathbf{q}}_{\ell+1} - \mathcal{H}_{i+1n}^{\mathbf{f}} \tilde{\mathbf{c}}_{i+1n} + \sum_{\ell=0}^i \mathcal{H}_{\ell i}^{\mathbf{f}} \tilde{\mathbf{c}}_{\ell i} \tilde{\mathbf{q}}_{i+1} \end{aligned}$$

Theorem 3. *Given a possibly sparse representation of $\tilde{\mathbf{c}}_{ji}$ for $0 < i \leq n$ and $0 \leq j \leq i$, the computation of $\tilde{\mathbf{q}}_i$ for all $0 < i \leq n$ amounts to a time complexity of $\mathcal{O}(\text{number of particles})$ and a space complexity of $\mathcal{O}(n)$. Further assume that we have computed $\mathcal{H}_{ji}^{\mathbf{f}}$ whenever particle $\tilde{\mathbf{c}}_{ji}$ exists. The time and space complexity of computing $\mathbb{E}(\mathbf{f}(\mathbf{X}_i) \mid \mathbf{Y}_{1:n} = y_{1:n})$ for all $0 < i \leq n$ at once reads $\mathcal{O}(\text{number of particles})$ and $\mathcal{O}(n)$, respectively.*

This poses a huge improvement to the straightforward approach that applies Lemma 17 directly. Under good conditions, with pruning in place, we are now able to compute

arbitrary marginal expectations, and therewith variances, skewnesses and many others, for all $i = 1, \dots, n$ at once with linear complexity w.r.t. n . In particular, this renders the EM algorithm of Section 5.5.1 feasible for big data purposes.

These inference formulas for segment height appear to be new within the changepoint literature. However, after developing these formulas, similar approaches dedicated to hidden semi-Markov models came to my attention. Ultimately, a careful comparison with Yu (2010) revealed a close algorithmic relationship between changepoint models and hidden semi-Markov models. The main difference between these models is that the latter employs only discrete distributions for the segment height.

5.4.1 Pointwise statements in the well-log example

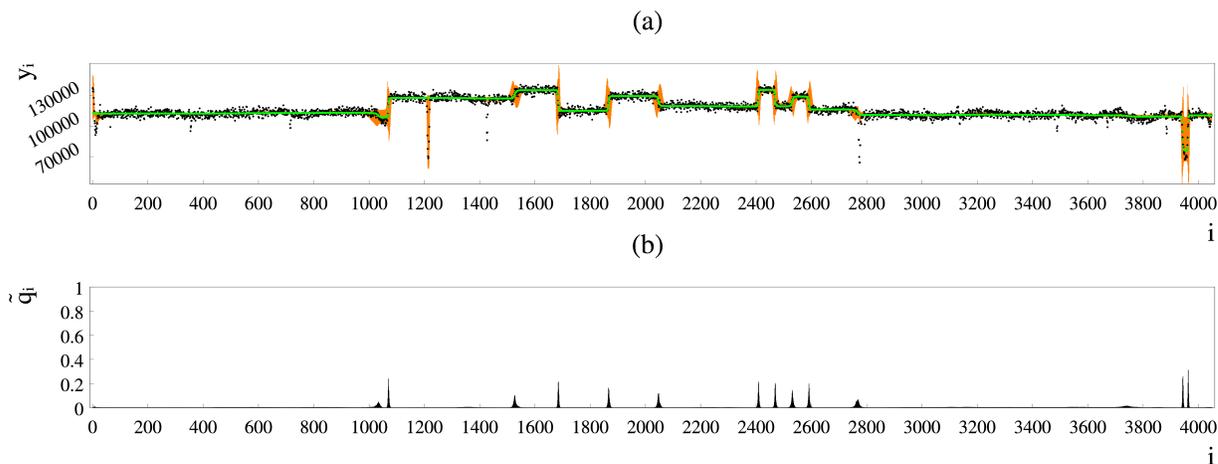


Figure 11: (a) Well-log data and posterior expectation of the segment heights in green and a depiction of the standard deviation and skewness in orange. (b) Marginal changepoint probabilities.

We are now interested in several posterior pointwise quantities with regard to our well-log example in Section 5.2.4. To this end, we compute the first three moments of the segment heights.

For $m \in \mathbb{N}$ we define $\mathbf{f}_m(x) = x^m$. According to Corollary 3, computing the m -th

moments of $(\mathbf{X}_i \mid \mathbf{Y}_{1:n} = y_{1:n})$ for $i = 1, \dots, n$ comes down to computing $\mathcal{H}_{ji}^{f_m}$ whenever particle (j, i) exists. Since $\mathcal{H}_{ji}^{f_m} = \mathbf{Z}_{ji}^m / \mathbf{Z}_{ji}^0$ (see Equation (20)), we achieve this by means of Corollary 2.

Figure 11(a) depicts the well-log data in black and in green it shows the expectations $\mathbf{m}_i = \mathbb{E}(\mathbf{X}_i \mid \mathbf{Y}_{1:n} = y_{1:n})$. We see that within a segment, the expectations do hardly vary even in the presence of outliers. However, they move away very rapidly at changepoint locations.

We take a simple and fairly arbitrary approach to visualize the standard deviations $\mathbf{sd}_i = \sqrt{\mathbb{E}\left((\mathbf{X}_i - \mathbf{m}_i)^2 \mid \mathbf{Y}_{1:n} = y_{1:n}\right)}$ and skewnesses $\mathbf{sk}_i = \mathbb{E}\left((\mathbf{X}_i - \mathbf{m}_i)^3 / \mathbf{sd}_i^3 \mid \mathbf{Y}_{1:n} = y_{1:n}\right)$ in a joint fashion. For each timepoint $i = 1, \dots, n$, (a) depicts the interval

$$[\mathbf{m}_i - 2 \cdot (\mathbf{sd}_i - \mathbf{sk}_i \mathbb{1}\{\mathbf{sk}_i < 0\}), \mathbf{m}_i + 2 \cdot (\mathbf{sd}_i + \mathbf{sk}_i \mathbb{1}\{\mathbf{sk}_i \geq 0\})]$$

in orange.

This shows that within a segment, the variations of the heights are fairly small. In turn, at changepoint locations their variations become much larger due to their ambiguity towards two different segment heights.

Whilst the mean doesn't follow the outliers, for the example at around 1200, the model expresses its tendencies through its skewness. Skewness also arises around changepoint locations. This is partially a consequence of the mixed form of the pointwise segmental distributions. However, as we have seen in Section 5.2.3, thanks to the Laplace distribution, single components of the mixture may also be skewed.

Thus, in the middle of longer segments, where only very few components dominate the posterior height distribution, the Laplace distribution is capable of maintaining skewness. In contrast, the normal distribution wouldn't allow this.

Figure 11(b) shows the marginal changepoint probabilities $\tilde{\mathbf{q}}_i$ for each timepoint $i = 1, \dots, n$. Within the Bayesian changepoint literature, these quantities are mainly used to justify changepoint models or to perform model selection in a manual fashion. By this means, we can obtain a rough impression of the spots where the model is inclined to place changepoints. According to Remark 3, the expected number of changepoints is 17.8.

Due to the location uncertainty, single variations in the data can usually be explained through changepoints w.r.t. more than one location. Within such a set of locations we usually see only one changepoint at a time. As a consequence, the probability of seeing changepoints there is shared among the corresponding $\tilde{\mathbf{q}}_i$'s. For this reason, the heights in Figure 11(b) appear to be quite disorganized and not very informative in terms of the sensitivity of the model to place changepoints.

Consider, for example, the dip at around 3800. The marginal changepoint probabilities appear suspiciously low and broad there. In fact, the (empirical) probability of seeing at least one changepoint between 3600 and 3900, inferred from 10^7 changepoint samples (see Section 5.3.2), amounts to 0.76, which is a highly unexpected result. We address this issue by means of a novel visualization scheme in Section 6.

Instead of the moments, we could also compute median and interquartile range. Even though this appears to be more suitable for the Laplace distribution, these quantities cannot be represented in terms of simple expectations and their implementation is in particular not straightforward for mixture distributions.

5.5 Parameter estimation via likelihood and the EM algorithm

In this section, we discuss means to estimate the parameters of a changepoint model. Thereby, we focus on the marginal loglikelihood function (Remark 1). Examples for parameters are μ, σ, τ, q and r in our well-log example (Section 5.2.4).

Assume that θ represents the underlying parameter (or parameters) of interest and Θ a real coordinate space that represents all possible values of θ . For the purpose of choosing a value for θ , we have to take into account that the latent changepoint process remains unknown. This is important because the model depends on the parameters and the segmentation jointly.

A primal approach could be to estimate the set of parameters, changepoint locations and segment heights simultaneously. However, this is very susceptible towards overfitting and therewith paves the way for degenerated models.

A workaround is to fit the parameters in accordance to a predetermined segmentation or if appropriate, a single representative segment. This takes advantage of prior knowledge, but doesn't guarantee that the model will later be sensitive towards this choice of segmentation or segment. More importantly, it ignores the uncertainty behind the choice of the segmentation.

Bayesian techniques are capable of incorporating uncertainty into the estimation. A common approach is to model θ through a random variable $\boldsymbol{\theta}$ and obtain a Bayes estimator, e.g. the posterior mean, $\mathbb{E}(\boldsymbol{\theta} \mid \mathbf{Y}_{1:n} = y_{1:n})$.

Such an estimator absorbs all possible segmentations through a marginalization over $\mathbf{X}_{1:n}$ and $\mathbf{C}_{1:n}$. This nicely takes care of the segmentation uncertainty. Furthermore, computing the expectation over $\boldsymbol{\theta}$ addresses the uncertainty behind the parameters.

In this course, however, we focus on the segmentation uncertainty. Thus, we leave the parameters non-random and consider the likelihood marginalized over all possible segmentations

$$L_\theta = \mathbb{P}(\mathbf{Y}_{1:n} \in dy_{1:n} ; \theta) / \psi(dy_{1:n})$$

The sign $;$ introduces the list of non-random parameters that are to be estimated. In Remark 1, we have already explained how to compute L_θ .

Under some conditions, finding a maximum of L_θ yields sensible estimates for θ . One of these conditions is that the distribution behind the data and the distribution employed in the changepoint model match sufficiently.

This can be an awkwardly strong obligation, especially, if it implies that the model is build from mixture distributions, e.g. where outliers are modeled separately. In fact, for changepoint inference, a tendency to organize the datapoints correctly into segments, is mostly sufficient. Thus, changepoint models just need to prove a tool not to reflect reality, a perspective that lies more in the heart of a Bayesian approach than a frequentist one. In this context, we may additionally resort to prior knowledge to steer the parameter estimation, by means of the marginal likelihood, into the right direction. Section 5.5.3 demonstrates this for the well-log example.

A simple grid search poses a first heuristic to find a maximizer of L_θ . However, if θ is of higher dimension, the required grid can easily become too huge to be scanned through. Moreover, the grid may be too rough at important sights, which promotes inaccurate results.

To this end, we could demand for more sophisticated algorithms like gradient descent (Cauchy, 1847), but we are usually unable to compute derivatives of L_θ . This is because L_θ is the product of an intricate marginalization that entangles θ and the observations in manifold ways.

Fortunately, the EM algorithm is able to overcome this issue by exploiting the fact that our changepoint models factorize into much smaller and simpler distributions.

In the following, we will first discuss the EM algorithm for changepoint models in general (Section 5.5.1). Afterwards, we discuss the EM algorithm for segment length for the geometric as well as the negative binomial distribution (Section 5.5.2). Section 5.5.3 elaborates a parameter estimation approach for the well-log example. Finally, we investigate our EM algorithm for segment height and observation and apply it to the Laplacian change in median model followed by a brief simulation study (Section 5.5.4 and 5.5.5).

5.5.1 An EM algorithm for changepoint models

The expectation-maximization (EM) algorithm (Dempster et al., 1977) is a popular way to gain estimators for the parameters of a statistical model that involves latent or unobserved variables. To this end, it provides local maximizers of the marginal likelihood. The great advantage of the EM algorithm is that it exploits the factorized structure many statistical models including changepoint models are subject to.

The EM algorithm works on the basis of densities. Therefore, assume that $(\mathbb{X}, \mathcal{X})$ is endowed with a measure ξ such that $\mathcal{J}(dx)/\xi(dx)$ exists.

For $c_{1:n} \in \mathbb{N}^n$ and $x_{1:n} \in \mathbb{X}^n$, let $\mathcal{L}_\theta(c_{1:n}, x_{1:n})$ be the complete loglikelihood of the

change point model w.r.t. $c_{1:n}$ and $x_{1:n}$, i.e.

$$\begin{aligned} \mathcal{L}_\theta(c_{1:n}, x_{1:n}) &= \sum_{i=1}^n \mathbf{1}\{c_i = i\} \mathbf{g}(x_i; \theta) + \mathbf{f}(y_i | x_i; \theta) \\ &+ \ln(\mathbb{P}(\mathbf{C}_1 = c_1; \theta)) + \sum_{i=2}^n \ln(\mathbb{P}(\mathbf{C}_i = c_i | \mathbf{C}_{i-1} = c_{i-1}; \theta)) \end{aligned} \quad (28)$$

whereby $\mathbf{g}(x; \theta) = \ln(\mathcal{J}(dx; \theta)/\xi(dx))$ and $\mathbf{f}(y_i | x; \theta) = \ln(\mathbf{p}(\mathbf{Y}_i = y_i | \mathbf{X}_i = x; \theta))$. For simplicity, we assume that $\mathbf{p}(\mathbf{Y}_i = y | \mathbf{X}_i = x; \theta)$ is of the same form for all $i = 1, \dots, n$ and agree that $\ln(0) = -\infty$ and $0 \cdot \ln(0) = 0$.

Algorithm 2 (EM algorithm). *Choose a $\theta_0 \in \Theta$ to start with. In step $\ell = 1, 2, \dots$ derive θ_ℓ as an element of*

$$\operatorname{argmax}_{\theta \in \Theta} \left\{ \mathbb{E}(\mathcal{L}_\theta(\mathbf{C}_{1:n}, \mathbf{X}_{1:n}) | \mathbf{Y}_{1:n} = y_{1:n}; \theta_{\ell-1}) \right\} \quad (29)$$

Keep iterating over ℓ until θ_ℓ converges.

This convergence is justified theoretically if L_θ is bounded (see, for example, Bishop and Mitchell (2014)). However, there is no guarantee that a global maximum of L_θ will be reached. Instead, we can only guarantee for a local maximum.

In practice, convergence is reached if the distances of successive values of θ_ℓ fall below a self defined threshold. If pruning is in place, instead of converging, the values of θ_ℓ may start to circulate. This needs to be taken care of in order to avoid unnecessary u-turns. Hoffman and Gelman (2014) describe such an approach in another context.

Since the expectation of a sum decomposes into a sum of expectations, we only need to compute the expectations w.r.t. to the single summands in (28). Furthermore, the desired parameters do usually solely affect either the distribution of time, observation or segment height. Thus, we consider these cases separately in order to focus on their particular characteristics.

The EM algorithm is an abstract algorithm, where each statistical model requires its own implementation. Several instances for particular models have been developed, most notably the Baum-Welch algorithm for (discrete) hidden Markov models (Rabiner and

Juang, 1986). Lately, interesting progress has been made in connection with continuous time hidden Markov models (Liu et al., 2015; Roberts and Ephraim, 2008). This concerns the estimation of the entries of the underlying generator matrix, given the observations.

Intrinsic to this estimation is the exploitation of endpoint-conditioned distributions (see also Hobolth and Stone (2009)). This is a decomposition of the posterior distribution, based on the discrete grid of observational timepoints. Even though, in some cases, this is applicable to changepoint models as well, in the following, we elaborate algorithms that require a more sophisticated approach. That is, a decomposition based on segmentations.

Until today, there are some particular instances of the EM algorithm for changepoint models available. It has been shown that an EM algorithm can be developed to detect changepoint locations (Keshavarz and Huang, 2014). Bansal et al. (2008), in turn, estimate the prior probabilities $\mathbb{P}(\mathbf{C}_i = i)$ in the single changepoint case by means of an EM algorithm. Furthermore, stochastic EM algorithms for parameter estimation in an online fashion have been developed (Yildirim et al., 2013).

We demonstrate implementations of our EM algorithm, which is concerned with parameter estimation, for a certain fundamental class of changepoint models and thereby provide a basic guide. At the same time, we discuss its limits in regard to feasibility.

We assume that we are in the middle of the iteration process of Algorithm 2. The current value of θ is θ_{old} and our task is to conduct the next EM step to provide a new value, say θ_{new} . In order to keep the notation uncluttered, we will mostly omit and not mention θ_{old} . However, it is important to keep in mind that the employed posterior distributions are build on this particular assignment of θ .

5.5.2 An EM algorithm for segment length

At first, we elaborate the case where the sojourn time is geometrically distributed. Through the iteration procedure of the EM algorithm, we want to compute the success probability q as part of θ , whereby q is not employed in any other distribution. We assume that q_{old} is the value of q within θ_{old} and q_{new} the value of q within θ_{new} . Subsequently, we explain how to obtain estimates for the success probability of the negative binomial distribution

when the number of successes is given.

In the geometric case we have

$$\begin{aligned}\ln(\mathbb{P}(\mathbf{C}_i = c_i \mid \mathbf{C}_{i-1} = c_{i-1})) &= \mathbb{1}\{c_i = i\} \ln(q) + \mathbb{1}\{c_i < i\} \ln(1 - q) \\ \ln(\mathbb{P}(\mathbf{C}_1 = c_1)) &= \mathbb{1}\{c_1 = 1\} \ln(q) + \mathbb{1}\{c_1 = 0\} \ln(1 - q)\end{aligned}\quad (30)$$

for $i = 2, \dots, n$.

Lemma 19. *We may set*

$$q_{new} = \frac{\tilde{\mathbf{q}}_1 + \dots + \tilde{\mathbf{q}}_n}{n}$$

Proof. We just need to find a global maximum of

$$\sum_{i=1}^n \tilde{\mathbf{q}}_i \ln(q) + (1 - \tilde{\mathbf{q}}_i) \ln(1 - q)$$

w.r.t. to q . □

According to Remark 3, in each step, this EM algorithm derives the expected number of changepoints and divides it by the number of timepoints.

Similar to Bansal et al. (2008), we may assume that the \mathbf{q}_{ji} are autonomous probabilities and independent of j , say $\mathbf{q}^i = \mathbf{q}_{ji}$. In this case, each step of the EM algorithm simply yields $\mathbf{q}_{new}^i = 1 - \tilde{\mathbf{q}}_i$ for $i = 1, \dots, n$.

Now, we discuss the EM algorithm for the success probability, q , of the negative binomial distribution with a predetermined r . We employ q_{new} and q_{old} as before.

This time, the expectation in Equation (29) will be composed based on all possible segments from j to i with $0 < i \leq n$ and $0 \leq j \leq i$. In this case, the negative binomial distribution contributes as $r \ln(q) + (i - j) \ln(1 - q)$, whereby we have skipped the binomial coefficient since it does not rely on q .

In the following, we distinguish three kinds of segments: introductory segments with $0 = j < i \leq n$, intermediate segment with $0 < j \leq i < n$ and final segments with $0 < j \leq i = n$.

Each segment accounts for a single summand and all summands together constitute the expectation in Equation (29). The intermediate segments are the easiest to deal with.

Lemma 20. *The summand for the intermediate segment from j to i with $0 < j \leq i < n$ can be represented as*

$$\tilde{\mathbf{c}}_{ji} \tilde{\mathbf{q}}_{i+1} (r \ln(q) + (i - j) \ln(1 - q))$$

Proof. For the segment from j to i we get

$$\mathbb{P}(\mathbf{C}_i = j, \mathbf{C}_{i+1} = i + 1 \mid \mathbf{Y}_{i:n} = y_{1:n}) \ln(q^r (1 - q)^{i-j}) = \tilde{\mathbf{c}}_{ji} \tilde{\mathbf{q}}_{i+1} \ln(q^r (1 - q)^{i-j})$$

whereby we have used Remark 2. □

We assume that the data observation was stopped in the middle of the data generating process that goes on forever. In order to facilitate the computation of the summands w.r.t. the final segments, we incorporate them as a whole and not only until timepoint n . Let \mathbf{S} be the length of the final segment, i.e. the segment that contains timepoint n .

Lemma 21. *We conclude that for $0 < j \leq n$ and $n - j \leq i$*

$$\mathbb{P}(\mathbf{S} = i, \mathbf{C}_n = j \mid \mathbf{Y}_{1:n} = y_{1:n}) = \frac{\tilde{\mathbf{c}}_{jn} \cdot NB(i; q_{old}, r)}{\sum_{\ell=n-j}^{\infty} NB(\ell; q_{old}, r)}$$

Proof. This follows from $(\mathbf{S} \perp\!\!\!\perp \mathbf{C}_{1:n-1}, \mathbf{Y}_{1:n}, \mathbf{X}_{1:n} \mid \mathbf{C}_n)$. □

Corollary 4. *We may use the following summand for the final segment that starts at $0 < j \leq n$*

$$\begin{aligned} & \sum_{i=n-j}^{\infty} \mathbb{P}(\mathbf{S} = i, \mathbf{C}_n = j \mid \mathbf{Y}_{1:n} = y_{1:n}) (r \ln(q) + i \ln(1 - q)) \\ &= \tilde{\mathbf{c}}_{jn} r \ln(q) + \tilde{\mathbf{c}}_{jn} \left(\sum_{\ell=n-j}^{\infty} \ell \cdot NB(\ell; q_{old}, r) \right) \Big/ \left(\sum_{\ell=n-j}^{\infty} NB(\ell; q_{old}, r) \right) \ln(1 - q) \end{aligned}$$

We further assume that the data was observed in the middle of the process which started infinitely long ago. Let now \mathbf{S} represent the random length from 0 to the end of the first segment. We get that $\mathbb{P}(\mathbf{S} = i) = \frac{q}{r(1-q)} \sum_{\ell=i+1}^{\infty} NB(\ell; q, r)$, whereby $\mathbf{S} = 0$ corresponds to $\mathbf{C}_1 = 1$ and $\mathbf{S} = i$ with $i > 0$ corresponds to $\mathbf{C}_i = 0$ and $\mathbf{C}_{i+1} = i + 1$. This distribution incorporates q in a fairly difficult manner rendering an exact computation of the desired summand for the EM algorithm infeasible.

While the negative binomial distribution may heavily concentrate its mass to specific lengths, the distribution of \mathbf{S} allows a much larger freedom. Intuitively speaking, it puts the most of its mass to all the lengths that are not longer than those lengths which are likely under the corresponding negative binomial distribution. Therefore, it is inappropriate to simply consider 1 as the first changepoint. Instead, we will use minor modifications of the length distribution to address these difficulties.

The distribution of \mathbf{S} has its mode at 0 and is monotonously decreasing. Thus, a geometric distribution appears to be a reasonable replacement. Its success probability, say q' , can be chosen as $q/r(1 - q)$. We come to this conclusion by looking at

$$\mathbb{P}(\mathbf{S} = i) = \frac{q}{r(1 - q)} \left(1 - \sum_{\ell=0}^i \binom{r - 1 + \ell}{\ell} q^r (1 - q)^\ell \right)$$

Considering $i = 0$, this means that $(q(1 - q^r))/(r(1 - q))$ appears to be the right candidate for q' . However, in order to keep things simple, we neglect q^r and set $q' = q/(r(1 - q))$ instead.

Please note that this approach precludes values of q that are larger than $r/(r + 1)$, because in this case, q' doesn't represent a probability anymore. Nevertheless, since for $q = r/(r + 1)$ we already get a nonsensical expected segment length of 1 under the negative binomial distribution, this restriction is insignificant.

Corollary 5. *By means of the above geometric approximation, the summand for the segment from 0 to i with $0 < i \leq n$ can be represented as*

$$\tilde{c}_{0i} \tilde{\mathbf{q}}_{i+1} (\ln(q) - \ln(1 - q) + i \ln(1 - q/(r(1 - q))))$$

Having computed the summands for all segments, the expectation in (29), which is subject to maximization, finally equates to

$$c_1 \ln(q) + c_2 \ln(1 - q) + c_3 \ln(1 - q/(r(1 - q))) \tag{31}$$

with constants c_1, c_2 and c_3 .

Lemma 22. *Let*

$$z_{1,2} = \frac{r(c_1 + c_2) + c_1(r + 1) + c_3 \pm \sqrt{(r(c_1 + c_2) + c_1(r + 1) + c_3)^2 + (r + 1)(c_1 + c_2)c_1r}}{4(r + 1)(c_1 + c_2)}$$

A maximizer for q in Equation (31) can be chosen uniquely from $\{z_1, z_2\}$.

Proof. Differentiating (31) and setting to 0 yields

$$\begin{aligned} c_1/q - c_2/(1 - q) - \frac{c_3}{r(1 - q)^2 - q(1 - q)} &= 0 \\ \Leftrightarrow q^2(r + 1)(c_1 + c_2) - q(r(c_1 + c_2) + c_1(r + 1) + c_3) + c_1r &= 0 \end{aligned} \quad (32)$$

whereby $q \notin \{0, 1, r/(r + 1)\}$. We can easily verify that $c_1 > 0$ and $c_3 \geq 0$. Since $0 < \mathbf{q}_{01} < 1$, we even get $c_3 > 0$. Substituting $q = 0$ into the l.h.s. of Equation (32) yields c_1r , which is positive, and for $q = r/(r + 1)$ we get $-c_3$, which is negative. Therefore, the l.h.s. of Equation (32) exhibits one single root between 0 and $r/(r + 1)$. Finally, $q \leq r/(r + 1)$ implies that this is the only solution. \square

The EM algorithm for the negative binomial distribution essentially parses through the existing \tilde{c}_{ji} 's and touches each of them exactly once. Thereby, it employs the marginal changepoint probabilities $\tilde{\mathbf{q}}_i$ for $i = 1, \dots, n$. Assume that we have computed \tilde{c}_{ji} for $0 \leq j \leq i$ and $0 < i \leq n$. By virtue of Theorem 3, we may state that the time and space complexity of each EM step for q in the geometric as well as in the negative binomial case is $\mathcal{O}(\text{number of particles})$.

5.5.3 Parameter estimation in the well-log example

In this section, we are concerned with determining the parameters σ , q and r in the well-log example of Section 5.2.4. μ and τ remain fixed here. Estimating the parameters of a statistical model is commonly referred to as training or learning. Once trained, the finalized model can be employed straight away whenever new data is available. To this end, we elaborate a practically adapted approach that employs the EM algorithm for segment length (Section 5.5.2) and the MAP (Section 5.3.1).

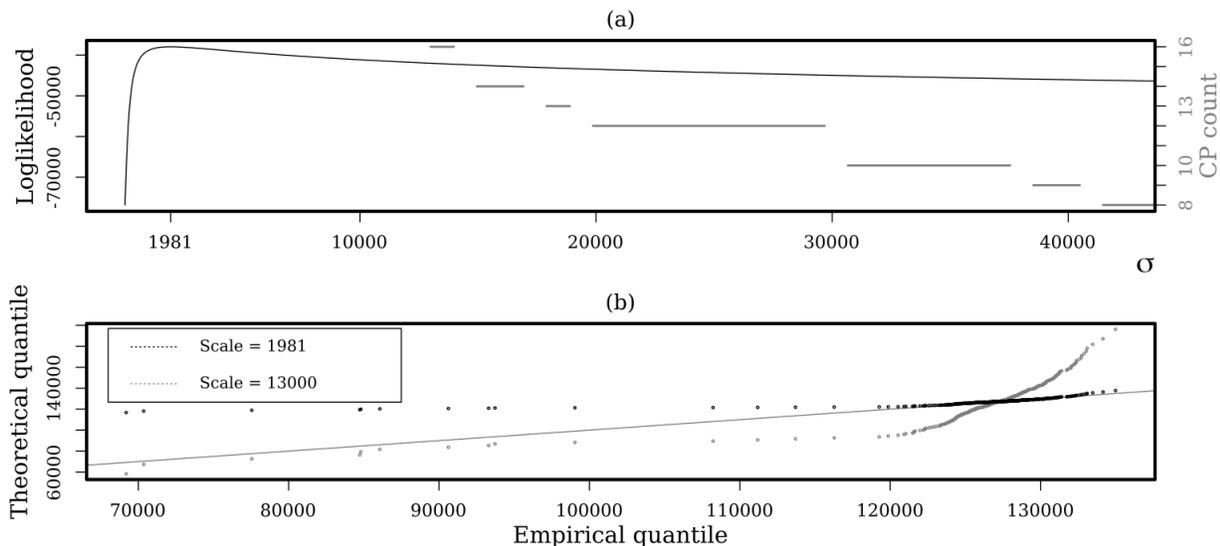


Figure 12: (a) Illustration of the loglikelihood (left axis) and number of changepoints inferred by the MAP (right axis) as a function of σ . (b) Q-Q plot of the empirical distribution of the datapoints from 1100 to 1500 against two Laplace distributions w.r.t. two different scale parameters. The straight line represents the diagonal.

We start with discussing σ . Figure 12(a) depicts the loglikelihood, $\ln(L_\sigma)$, of our changepoint model and the number of changepoints inferred by the MAP, both w.r.t. σ and with fixed $q = 0.01430724$ and $r = 3$.

L_σ exhibits a single maximum at 1981, which yields an unreasonable changepoint count of far more than 16. Thus, an estimator for σ , solely based on maximizing the likelihood, doesn't provide reasonable estimates here.

The seemingly right-skewed shape of $\ln(L_\sigma)$ suggests that the posterior mean w.r.t. σ may provide more suitable estimates. For this sake, we put a prior on σ , which is the uniform distribution from 50 to 100000. Furthermore, we use a piecewise linear approximation of $\ln(L_\sigma)$ to compute $\mathbb{E}(\sigma \mid \mathbf{Y}_{1:n} = y_{1:n})$ in the fashion of Corollary 2. Surprisingly, this yields a similarly useless value of approximately 1984. This is because, those values of σ near the peak dominate the whole calculation.

Another naturally arising estimation procedure is to infer σ by fitting a manually chosen

representative segment. This has the advantage that it can be performed independently of q .

We take the segment from timepoint 1100 to 1500. Figure 12(b) depicts two Q-Q plots and the diagonal. One Q-Q plot shows the empirical distribution against the Laplace distribution with scale equal to 1981 and location equal to the segment's median. The other Q-Q plot depicts the same distributions but this time with scale equal to 13000.

The fact that 1981 maximizes L_σ is reflected by the good fit of the intermediate quantile pairs to the diagonal. However, the lower quantile pairs are situated far above the diagonal, which indicates that the corresponding theoretical quantiles are much larger than the empirical ones. Thus, the large deviations present in the data are very unlikely for this particular Laplace distribution, which provokes many undesirable changepoints. In turn, $\sigma = 13000$ appears to be a more suitable choice, since the corresponding lower quantile pairs stay below the diagonal.

What we can learn thereof, is that the data within the segments does not stem from a Laplace distribution. Apparently, the data is strongly asymmetric, but more importantly, the portion of large deviations is simply too small. A likelihood based estimation does, however, rely on the assumption that the empirical and theoretical distributions match sufficiently. As a result, maximizing L_σ yields a far too small value for σ .

Having set σ to 13000, we would now choose $q = 0.0088$ yields an average length of 339 for the negative binomial distribution with $r = 3$ and thus around twelve changepoints on average within the 4049 datapoints.

Therewith, the MAP estimator does not detect any changepoints within the considered segment, but 2 extra changepoints, one at 1034 and one at 3744. The latter is highly undesirable, since it belongs to a perturbation in the data that takes the form of a dip.

However, the same reasoning, based on Q-Q plots, conducted on the segment that contains 3744 doesn't lead to a larger σ . Thus, in order to get rid of the additional changepoints, we either have to increase σ or decrease q based on a different approach.

Since the additional segments concerning the changepoint at 3744 have reasonable lengths, we shouldn't address this by decreasing q . Furthermore, the probability of segments

with a length of at most 50 is only 1% under the negative binomial distribution with $r = 3$ and $q = 0.0088$. Such segments are readily present in the data suggesting that we should not only increase σ , but also q .

Hence, our final approach is to estimate q and σ in a simultaneous fashion. Thereby, we resort to the likelihood to obtain appropriate values for q and the MAP to adjust σ . The MAP provides a most likely changepoint configuration and is therefore well-suited for controlling the tendencies of the posterior distributions. Fortunately, those MAP estimates having twelve changepoints seem to match sufficiently with our desired locations depicted in Figure 10(a).

It is reasonable to assume that our well-log data represents an average example for this application. Therefore, we aim at selecting a set of parameters that lies centrally within a region of suitable values for σ and q . This induces a sense of robustness, which also supports the reusability of the estimated parameters in case new data is observed.

Additionally, we need to take numerical aspects into account. The roles of q and σ are mutually intertwined. We can, for example, obtain longer segments by either decreasing q or increasing σ . Hence, smaller values of q might be combined with smaller values of σ and vice versa. This collinearity paves the way to nonsensical, extreme parameter choices, which may seriously impair the accuracy during the calculations. Aiming at high likelihood values wherever possible mitigates the risk concerning this issue.

To this end, we employ the function $\sigma \mapsto \hat{q}$ that returns the EM estimate for q given σ . By this means, we choose σ such that the MAP solidly yields twelve changepoints given that q is set to \hat{q} and r is set to 3.

Figure 13(a) demonstrates this. It shows the graph of \hat{q} w.r.t. σ and the corresponding changepoint count inferred by the MAP. $\sigma = 25000$ is roughly centered within those values that yield twelve changepoints. The corresponding value of \hat{q} reads 0.01430724.

Figure 13(b) shows the loglikelihood w.r.t. q while having set σ to 25000, and the corresponding changepoint count inferred by the MAP. We see that the value 0.01430724, indicated by the circle, is also roughly centered within those values that yield twelve changepoints. Moreover, the loglikelihood exhibits only one single global maximum being

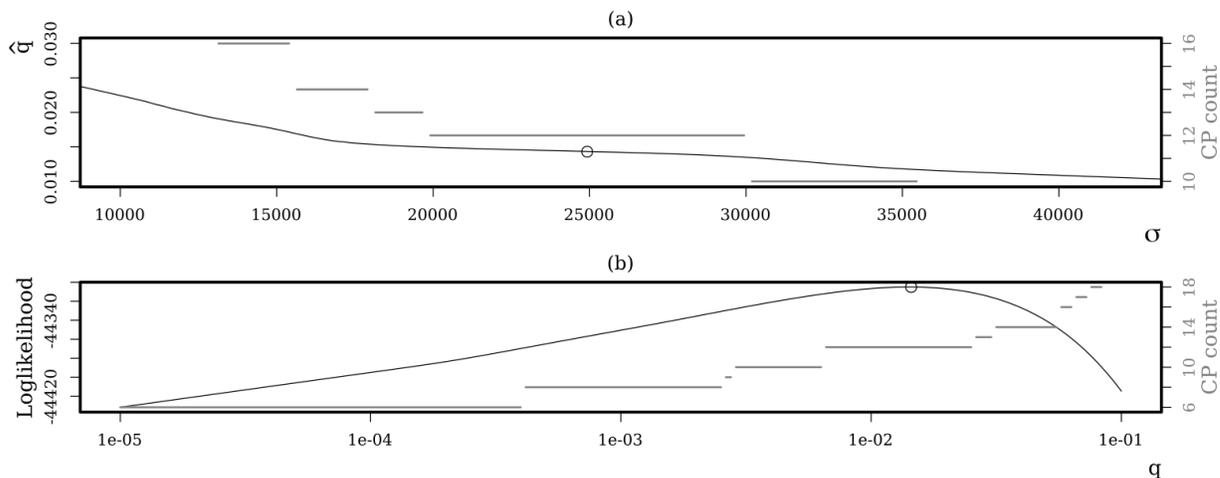


Figure 13: (a) Illustration of \hat{q} (left axis) and the number of changepoints inferred by the MAP (right axis) w.r.t. σ . The circle marks the chosen set of parameters for σ and q . (b) Illustration of the loglikelihood (left axis) and the number of changepoints inferred by the MAP (right axis) w.r.t. q and $\sigma = 25000$. The circle marks the global maximum of the likelihood.

correctly inferred by the EM algorithm.

Therewith, we have considered two orthogonal directions of the (σ, q) -plane. Although unlikely, this would be misleading if the MAP estimates spread very irregularly in the neighborhood of our selected set of parameter values.

Figure 1, however, suggests that this is not the case. It depicts an excerpt of the (σ, q) -plane and colored tiles rendered in accordance with the changepoint counts inferred by the MAP. The upper and lower boundaries of the tiles w.r.t. σ were computed by a simple binary search heuristic. The cross points to $\sigma = 25000$ and $q = 0.01430724$.

As we see, our set of parameters is situated centrally within a comparably thick and regular region. The unsteady or even incoherent nature of other regions indicates that there are dissimilar changepoint configurations with equal changepoint counts. The aforementioned collinearity between σ and q is reflected through the beveled appearance of the regions.

Finally, we discuss the choice of r . With increasing r , the mode of the length distribution shifts away from zero and thereby potentially increases the minimal size of the segments. This mitigates the influence of outliers and thus allows for a slight decrease of the segmental scale. Therewith, the sensitivity of the model is improved without impairing its specificity towards outliers.

However, we cannot increase r arbitrarily in the presence of comparably small segments. We found that r should not be larger than 3. This is because for $r > 3$ the negative binomial distribution, when fitted to the segment lengths, turns out to be too sharply peaked. As a result, the sensitivity towards unreasonable changepoints soars. We found that for $r = 5, 6$ and $q = \hat{q}$, we can't even obtain a σ that yields a MAP with twelve changepoints.

5.5.4 An EM algorithm for segment height and observation

Now, we elaborate the EM algorithm for segment height and observation by means of three exponential family alike families of distributions. They cover a good range of practical cases and convey the basic idea behind this nicely. We will further contemplate complexity and feasibility concerns.

We assume that \mathbf{f} and \mathbf{g} solely rely on their very own parameters θ^g and θ^f . θ^g and θ^f are defined on sets Θ^g and Θ^f , respectively, and are part of θ , but not linked to each other or any parameter that concerns the distribution of segment length.

Thus, we maximize the expectations of $\sum_{i=1}^n \mathbf{f}(y_i | \mathbf{X}_i; \theta^f)$ and $\sum_{i=1}^n \mathbf{1}\{\mathbf{C}_i = i\} \mathbf{g}(\mathbf{X}_i; \theta^g)$ separately (see Equation (28) and Algorithm 2). We use θ_{old}^f and θ_{old}^g as well as θ_{new}^f and θ_{new}^g to represent the values of θ^f and θ^g within θ_{old} and θ_{new} .

Lemma 23. *Assume that $\mathbf{g}(x; \theta^g) = \eta(\theta^g)^t T(x) + A(\theta^g)$ with functions $\eta : \Theta^g \rightarrow \mathbb{R}^u$, $A : \Theta^g \rightarrow \mathbb{R}$ and measurable $T : \mathbb{X} \rightarrow \mathbb{R}^u$. Any value of θ^g that maximizes*

$$(\tilde{q}_1 + \dots + \tilde{q}_n)A(\theta^g) + \sum_{\ell=1}^u \eta_{\ell}(\theta^g) \sum_{i=1}^n \mathbb{E}(T_{\ell}(\mathbf{X}_i) \mathbf{1}\{\mathbf{C}_i = i\} | \mathbf{Y}_{1:n} = y_{1:n}) \quad (33)$$

is a candidate for θ_{new}^g .

Proof. This follows directly from the particular form of \mathbf{g} . \square

The functional form of \mathbf{g} in Lemma 23 states that \mathcal{J} is a member of the exponential family w.r.t. θ^g . Since θ^g and x are employed in separate factors within \mathbf{g} , θ^g does not concern the expectations drawn from the $T_\ell(\mathbf{X}_i)\mathbb{1}\{\mathbf{C}_i = i\}$'s. Thus, the ensuing maximization w.r.t. θ^g can fortunately be performed based on constant expressions for these expectations. If this was not the case, maximization in closed form would easily become a hopeless task.

Typically, we obtain maximizers of (33) through differentiating and zeroing. In this case, we require that the η_ℓ 's and A are (partially) differentiable w.r.t. θ^g . The posterior expectation of $T_\ell(\mathbf{X}_i)\mathbb{1}\{\mathbf{C}_i = i\}$ for $\ell = 1, \dots, u$ and $i = 1, \dots, n$ needs to be feasible. This feasibility relies on the ability to compute all the $\mathcal{H}_{ji}^{T_\ell}$'s (see Equation (27)) as the following lemma shows.

Lemma 24. *For $\ell = 1, \dots, u$ and $0 < i \leq n$ we get*

$$\mathbb{E}(T_\ell(\mathbf{X}_i)\mathbb{1}\{\mathbf{C}_i = i\} \mid \mathbf{Y}_{1:n} = y_{1:n}) = \sum_{j=i}^{n-1} \mathcal{H}_{ij}^{T_\ell} \tilde{\mathbf{c}}_{ij} \tilde{\mathbf{q}}_{j+1} + \mathcal{H}_{in}^{T_\ell} \tilde{\mathbf{c}}_{in}$$

Proof. This follows from

$$\begin{aligned} & \mathbb{P}(\mathbf{X}_i \in \cdot, \mathbf{C}_i = i \mid \mathbf{Y}_{1:n} = y_{1:n}) \\ &= \sum_{j=i}^{n-1} \mathbb{P}(\mathbf{X}_j \in \cdot, \mathbf{C}_{j+1} = j+1, \mathbf{C}_j = i \mid \mathbf{Y}_{1:n} = y_{1:n}) + \mathbb{P}(\mathbf{X}_n \in \cdot, \mathbf{C}_n = i \mid \mathbf{Y}_{1:n} = y_{1:n}) \\ &= \sum_{j=i}^{n-1} \mathcal{H}_{ij} \tilde{\mathbf{c}}_{ij} \tilde{\mathbf{q}}_{j+1} + \mathcal{H}_{in} \tilde{\mathbf{c}}_{in} \end{aligned}$$

and a subsequent integration. \square

Lemma 24 touches each particle only once. This paves the way for an EM algorithm for segment height where each step has a time and space complexity of \mathcal{O} (number of particles). Lemma 24 and Theorem 3 clarify that the main bottle-neck here is the feasibility and complexity of computing $\tilde{\mathbf{c}}_{ji}$ and $\mathcal{H}_{ij}^{T_\ell}$ for $0 < i \leq n$, $0 \leq j \leq i$ and $\ell = 1, \dots, u$.

Distribution and \mathbf{g}	T , η and A
Gaussian $\mathbf{g}(x; \mu, \sigma) = -\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2 - \frac{1}{2}\ln(2\pi\sigma^2)$	$T(x) = -\frac{1}{2}(x^2, x, 1)^t$ $\eta(\mu, \sigma) = (1/\sigma^2, -2\mu/\sigma^2, \mu^2/\sigma^2)^t$ $A(\mu, \sigma) = -\frac{1}{2}\ln(2\pi\sigma^2)$
Laplacian $\mathbf{g}(x; \sigma) = -\frac{ x-\mu }{\sigma} - \ln(2\sigma)$	$T(x) = - x - \mu $ $\eta(\sigma) = 1/\sigma$ $A(\sigma) = -\ln(2\sigma)$
Inverse gamma $\mathbf{g}(x; \alpha, \beta) = \alpha \ln(\beta) - \ln(\Gamma(\alpha))$ $-(\alpha + 1) \ln(x) - \beta/x$	$T(x) = -(\ln(x), 1/x)^t$ $\eta(\alpha, \beta) = (1 + \alpha, -\beta)^t$ $A(\alpha, \beta) = \alpha \ln(\beta) - \ln(\Gamma(\alpha))$
Beta $\mathbf{g}(x; \alpha, \beta) = (\alpha - 1) \ln(x)$ $+(\beta - 1) \ln(1 - x) - \ln(\text{B}(\alpha, \beta))$	$T(x) = (\ln(x), \ln(1 - x))^t$ $\eta(\alpha, \beta) = (\alpha - 1, \beta - 1)^t$ $A(\alpha, \beta) = -\ln(\text{B}(\alpha, \beta))$

Table 2: A brief collection of univariate distributions for segment height and the corresponding realizations of T , η and A .

Table 2 lists some common univariate distributions for segment height and the corresponding form of T , A , and η in Lemma 23. While the normal distribution is a conjugate prior in the Gaussian change in mean model, the Laplace distribution (Section 5.2.3) is suitable for the Laplacian change in location model. The inverse gamma distribution comes into place in the Gaussian and Laplacian change in scale model and the beta distribution is involved in all sorts of change in success rate models.

The Gaussian change in mean model is completely benign. It just requires us to compute the first and second moment, and some simple derivatives w.r.t. σ and μ . In the Laplacian change in location model, where σ is to be estimated, we need to compute the first absolute moments of \mathcal{H}_{ji} . This is elaborated in Section 5.5.5.

In the Gaussian change in variance model where the inverse gamma distribution is used, the \mathcal{H}_{ji} 's are again of the inverse gamma type. Computing expectations w.r.t. these kind of simple distributions doesn't pose a problem and could be performed numerically. In turn,

the maximization involves the gamma function, which might foil a closed form maximization. However, we may again resort to a numeric method to achieve this. The same conclusions hold for the beta distribution, where the \mathcal{H}_{j_i} 's are typically beta distributions again.

The next lemma concerns the EM algorithm for observation with regards to a particular class of distributions.

Lemma 25. *Assume that $\mathbf{f}(y \mid x; \theta^f) = \eta'(\theta^f, y)^t T'(x) + A'(\theta^f)$ with functions $\eta' : \Theta^f \times \mathbb{Y} \rightarrow \mathbb{R}^v$, $A' : \Theta^f \rightarrow \mathbb{R}$ and measurable $T' : \mathbb{X} \rightarrow \mathbb{R}^v$. Any value of θ^f that maximizes*

$$nA'(\theta^f) + \sum_{\ell=1}^v \sum_{i=1}^n \eta'_\ell(\theta^f, y_i) \mathbb{E}(T'_\ell(\mathbf{X}_i) \mid \mathbf{Y}_{1:n} = y_{1:n}) \quad (34)$$

is a candidate for θ_{new}^f .

Proof. This follows directly from the particular form of \mathbf{f} . □

In the context of Lemma 25, we may again be able to achieve a space and time complexity of $\mathcal{O}(\text{number of particles})$ for each EM step. Theorem 3 and Corollary 3 show that the main bottle-neck here is the feasibility and complexity of computing $\tilde{\mathbf{c}}_{j_i}$ and $\mathcal{H}_{i_j}^{T'_\ell}$ for $0 < i \leq n$, $0 \leq j \leq i$ and $\ell = 1, \dots, v$.

However, the requirement of Lemma 25 is relatively strict. It does, for example, not hold in the Laplacian case where we want to estimate σ , since $|y - x|$ does not decompose into two separate factors. Therefore, we would need to compute $\mathbb{E}(|y_i - \mathbf{X}_i| \mid \mathbf{Y}_{1:n} = y_{1:n})$ for all $i = 1, \dots, n$, which might increase the overall complexity by a factor of n or more. We may thus, be more successful by other means as the next lemma shows.

Lemma 26. *Assume that $\mathbf{f}(y \mid x; \theta^f) = \eta'(\theta^f)^t T'(x, y) + A'(\theta^f)$ with functions $\eta' : \Theta^f \rightarrow \mathbb{R}^v$, $A' : \Theta^f \rightarrow \mathbb{R}$ and measurable $T' : \mathbb{Y} \times \mathbb{X} \rightarrow \mathbb{R}^v$. For $0 \leq j \leq i$, $0 < i \leq n$ and $\ell = 1, \dots, v$ let*

$$\mathbf{o}_{j_i}^\ell = \int \sum_{k=\max(j,1)}^i T_\ell(x, y_k) \mathcal{H}_{j_i}(dx)$$

Any value of θ^f that maximizes

$$nA'(\theta^f) + \sum_{\ell=1}^v \eta'_\ell(\theta^f) \left(\sum_{i=1}^{n-1} \sum_{j=0}^i \tilde{c}_{ji} \tilde{\mathbf{q}}_{i+1} \mathbf{o}_{ji}^\ell + \sum_{j=0}^n \tilde{c}_{jn} \mathbf{o}_{jn}^\ell \right) \quad (35)$$

is a candidate for θ_{new}^f .

Proof. The proof follows from

$$\mathbb{E} \left(\sum_{i=1}^n T'_\ell(\mathbf{X}_i, y_i) \mid \mathbf{Y}_{1:n} = y_{1:n} \right) = \sum_{i=1}^{n-1} \sum_{j=0}^i \tilde{c}_{ji} \tilde{\mathbf{q}}_{i+1} \mathbf{o}_{ji}^\ell + \sum_{j=0}^n \tilde{c}_{jn} \mathbf{o}_{jn}^\ell$$

This, in turn, can be seen as follows

$$\begin{aligned} & \mathbb{E} \left(\sum_{i=1}^n T'_\ell(\mathbf{X}_i, y_i) \mid \mathbf{Y}_{1:n} = y_{1:n} \right) \\ &= \sum_{\mathbf{c}_{1:n} \in \mathbb{N}^n} \mathbb{P}(\mathbf{C}_{1:n} = \mathbf{c}_{1:n} \mid \mathbf{Y}_{1:n} = y_{1:n}) \sum_{\substack{j \leq i \text{ with } c_i = j \text{ and} \\ (i = n \text{ or } c_{i+1} = i + 1)}} \mathbf{o}_{ji}^\ell \\ &= \sum_{i=1}^n \sum_{j=0}^i \mathbf{o}_{ji}^\ell \sum_{\substack{\mathbf{c}_{1:n} \in \mathbb{N}^n \text{ with } c_i = j \text{ and} \\ (i = n \text{ or } c_{i+1} = i + 1)}} \mathbb{P}(\mathbf{C}_{1:n} = \mathbf{c}_{1:n} \mid \mathbf{Y}_{1:n} = y_{1:n}) \\ &= \sum_{i=1}^{n-1} \sum_{j=0}^i \tilde{c}_{ji} \tilde{\mathbf{q}}_{i+1} \mathbf{o}_{ji}^\ell + \sum_{j=0}^n \tilde{c}_{jn} \mathbf{o}_{jn}^\ell \end{aligned}$$

□

The advantage of Equation (35) in combination with Lemma 26 is that we just need to compute one expectation per particle. Thus, the overall complexity of each EM step is $\mathcal{O}(\text{number of particles})$.

Table 3 provides a brief list of univariate observational distributions for certain change-point problems. We can also read the corresponding form of T' , η' and A' , and the lemma that should be applied in order to develop an EM algorithm.

As before, if we exploit conjugacy, we will obtain simple posterior distributions. Thus, computing logarithmic or inverse moments as in the Gaussian change in variance and Laplacian change in scale examples can be done numerically.

Changepoint problem and \mathbf{f}	T' , η' and A'	Apply Lemma
Gaussian change in mean $\mathbf{f}(y \mid x; \sigma) =$ $-\frac{1}{2}\left(\frac{y-x}{\sigma}\right)^2 - \frac{1}{2}\ln(2\pi\sigma^2)$	$T'(x) = -\frac{1}{2}(x^2, x, 1)^t$ $\eta'(\sigma, y) = (1/\sigma^2, -2y/\sigma^2, y^2/\sigma^2)^t$ $A'(\sigma) = -\frac{1}{2}\ln(2\pi\sigma^2)$	25
Gaussian change in variance $\mathbf{f}(y \mid x; \mu) =$ $-\frac{1}{2}\frac{(y-\mu)^2}{x} - \frac{1}{2}\ln(2\pi x)$	$T'(x) = -\frac{1}{2}(1/x, 1/x, 1/x, \ln(2\pi x))^t$ $\eta'(\mu, y) = (y^2, -2y\mu, \mu^2, 1)^t$ $A'(\mu) = 0$	25
Laplacian change in location $\mathbf{f}(y \mid x; \sigma) =$ $-\frac{ y-x }{\sigma} - \ln(2\sigma)$	$T'(x, y) = - y - x $ $\eta'(\sigma) = 1/\sigma$ $A'(\sigma) = -\ln(2\sigma)$	26
Laplacian change in scale $\mathbf{f}(y \mid x; \mu) =$ $-\frac{ y-\mu }{x} - \ln(2x)$	$T'(x) = -(1/x, \ln(2x))^t$ $\eta'(\mu, y) = (y - \mu , 1)^t$ $A'(\mu) = 0$	25

Table 3: A brief collection of univariate observational distributions for specific changepoint problems. It states the functions T' , η' and A' , and the lemma that should be applied in order to build an EM algorithm for the unknown parameters.

5.5.5 EM algorithm for the scales in the well-log example

In the following, we want to apply the results of Section 5.5.4 to the well-log example (Sections 5.2.3, 5.2.4, 5.4.1 and 5.5.3). Our task is to develop an EM algorithm for τ and σ . μ does not meet the requirements of Lemma 23 since $|x - \mu|$ does not decompose into two separate factors.

In order to build an EM algorithm for τ , we apply Lemma 23. Therefore, we need to compute

$$\mathbf{m}_0 = \sum_{i=1}^n \mathbb{E} \left(|\mathbf{X}_i - \mu| \cdot \mathbf{1} \{ \mathbf{C}_i = i \} \mid \mathbf{Y}_{1:n} = y_{1:n} \right)$$

According to Lemma 24, this entails the computation of \mathcal{H}_{ji}^T with $T(x) = |x - \mu|$ for all

$0 < i \leq n$ and $0 \leq j \leq i$. Finally, we obtain $\tau_{new} = \mathbf{m}_0 / (\tilde{\mathbf{q}}_1 + \dots + \tilde{\mathbf{q}}_n)$ through maximizing $-(\tilde{\mathbf{q}}_1 + \dots + \tilde{\mathbf{q}}_n) \ln(2\tau) - \mathbf{m}_0/\tau$ w.r.t τ .

Hence, it remains to show how to compute $\mathcal{H}_{ji}^T = \int |x - \mu| \mathcal{H}_{ji}(dx)$. Therefore, consider that $\mathcal{H}_{ji}^T \cdot \mathbf{Z}_{ji}^0 = \int |x - \mu| \exp\left(-\frac{|x-\mu|}{\tau} - \sum_{\ell=j}^i \frac{|x-y_\ell|}{\sigma}\right) dx$. This elementary integration is explained in the following corollary.

Corollary 6. *In the context of Corollary 2 the following holds*

$$\int |x - \mu| \exp\left(-\sum_{\ell=1}^r \frac{|x - z_\ell|}{\sigma_\ell}\right) dx = \sum_{\ell=1}^r \exp(\mathbf{d}_\ell z_\ell + \mathbf{e}_\ell) (\mathbf{b}_{\ell-1}^\ell - \mathbf{b}_\ell^\ell) \quad (36)$$

with

$$\mathbf{b}_j^k = \text{sgn}(z_k - \mu) \cdot \begin{cases} \frac{z_k - \mu}{\mathbf{d}_j} - \frac{1}{\mathbf{d}_j^2} & \mathbf{d}_j \neq 0 \\ z_k^2/2 - \mu z_k & \mathbf{d}_j = 0 \end{cases}$$

for $j = 0, \dots, r$ and $k = 1, \dots, r$.

Proof. The proof is similar to the proof of Corollary 2. □

An EM algorithm for σ can be derived by means of Lemma 26. To this end, we need to compute $\mathbf{o}_{ji} = \int \sum_{\ell=\max(j,1)}^i |x - y_\ell| \mathcal{H}_{ji}(dx)$ for all $0 < i \leq n$ and $0 \leq j \leq i$. With this, we further compute

$$\mathbf{m}_1 = \sum_{i=1}^{n-1} \sum_{j=0}^i \tilde{\mathbf{c}}_{ji} \tilde{\mathbf{q}}_{i+1} \mathbf{o}_{ji} + \sum_{j=0}^n \tilde{\mathbf{c}}_{jn} \mathbf{o}_{jn}$$

Finally, we obtain $\sigma_{new} = \mathbf{m}_1/n$ by maximizing $-n \ln(2\sigma) - \frac{\mathbf{m}_1}{\sigma}$ w.r.t. σ .

The integrand in \mathbf{o}_{ji} represents a cumbersome function we haven't dealt with yet. We can avoid it by considering that

$$\mathbf{o}_{ji} = \sigma \int \frac{|x - \mu|}{\tau} + \sum_{\ell=j}^i \frac{|y_\ell - x|}{\sigma} \mathcal{H}_{ji}(dx) - \frac{\sigma}{\tau} \mathcal{H}_{ji}^T$$

In this case, the integrand and the exponent within $\mathcal{H}_{ji}(dx)/\xi(dx)$ differ only in terms of their sign. The following lemma explains this more elementary integration.

Corollary 7. *In the context of Corollary 2 the following holds*

$$\int \sum_{\ell=1}^r \frac{|x - z_\ell|}{\sigma_\ell} \exp \left(- \sum_{\ell=1}^r \frac{|x - z_\ell|}{\sigma_\ell} \right) dx = \sum_{\ell=1}^r \exp(\mathbf{d}_\ell z_\ell + \mathbf{e}_\ell) (\mathbf{b}_{\ell-1}^\ell - \mathbf{b}_\ell^\ell) \quad (37)$$

with

$$\mathbf{b}_j^k = \begin{cases} z_k - \frac{1-\mathbf{e}_j}{\mathbf{d}_j} & \mathbf{d}_j \neq 0 \\ \mathbf{e}_j z_k & \mathbf{d}_j = 0 \end{cases}$$

for $j = 0, \dots, r$ and $k = 1, \dots, r$.

Proof. The proof is similar to the proof of Corollary 2. □

Equipped with this knowledge, we now conduct a brief simulation study to examine the EM algorithm for q and σ in the Laplacian change in median model that employs a negative binomial distribution for segment length. Therefore, we compare different values of r with regard to the simultaneous estimation of the two parameters.

In frequentist settings, the changepoint locations are usually not considered as random. Thus, frequentists deem the prior distribution for segment length just as a necessary evil of the Bayesian approach. In this context, the main interest lies in estimates for σ .

However, the choice of the length distribution has to be taken with care, because it impacts the estimation of σ . Once again, this can be explained by the collinearity of the segment lengths and the observational scale (see Section 5.5.3). In the following, we are particularly interested in understanding this impact.

We repeatedly sample 4050 datapoints from a Laplace distribution with scale equal to 1 and changing median. The segment heights (the medians) are sampled from a Laplace distribution with median equal to 0 and scale equal to 10. Each data set contains 12 random changepoints drawn uniformly.

Figure 14 displays the average EM estimate for σ w.r.t. to different choices of r and an increasing number of data sets. The horizontal line marks the true value of σ . Of course, convergence towards certain limits can not be derived safely hereof, but tendencies are

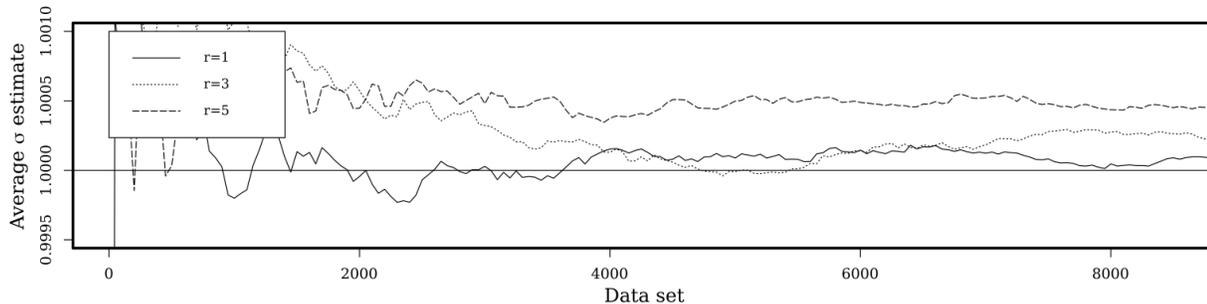


Figure 14: The graphs show how the average EM estimate for σ evolves for different values of r at an increasing number of data sets. The horizontal line indicates the true value of σ .

apparent. All graphs indicate a positive bias, which seems to increase with r . Anyway, the three biases concern at most the fourth digit of the true value of σ and are thus tiny.

The geometric distribution, i.e. $r = 1$, places the changepoints uniformly regardless of $q \in (0, 1)$ if their number is given. It is thus the most natural choice for the distribution of segment length here. This explains why $r = 1$ yields the smallest (or perhaps no) bias from the true value of σ .

On the other hand, a uniform distribution for changepoints tends to produce very short segments. Those segments combined with comparably small height changes can barely be distinguished from random perturbations in the data. Since the number of changepoints is uncertain during the estimation, segments are overlooked occasionally. This may result in an overestimation of σ . Furthermore, increasing r inhibits short segments enhancing this effect.

Table 4 displays different statistics, derived from the EM estimation of σ and q w.r.t. 10000 data sets. The mean squared error of the estimation for σ appears to be lowest for $r = 1$ and increases with r . Furthermore, the expected changepoint count seems to be a good (roughly unbiased) estimator for the true changepoint count when the length distribution is chosen appropriately. In accordance with the above reasoning, $r = 1$ is appropriate here. The MAP estimator overlooks a considerable portion of changepoints for all choices of r . This reflects its robustness towards uncertain changepoints.

r	MSE	Average expected change point count	Average MAP	Average expected segment length
1	0.000243	12.0178	11.562	336.99
3	0.000249	11.8754	11.513	334.34
5	0.000251	12.1006	11.471	327.28

Table 4: For different choices of the negative binomial distribution, the table shows several statistics derived from the EM algorithm w.r.t. to 10000 data sets.

A potentially contra intuitive observation is that the expected segment length gets smaller with increasing r . This is because the negative binomial distribution, for larger values of r , needs to be tightened more and more, in order to adjust to the entirety of possible segment lengths. At the same time, the expected change point counts exhibit an irregular and odd picture.

Therefore, we can conclude that too large values of r may provoke overfitting towards strongly peaked length distributions. This is highly undesirable and calls for a careful determination of r . On the other hand, increasing r helps to mitigate the effect of sudden outliers by inhibiting short segments (see also Section 5.5.3).

5.6 Discussion

In this section, we developed a wide range of exact inference strategies for a certain class of changepoint models. The central achievement of this work is Theorem 3. It shows that under good conditions all the considered algorithms exhibit a linear complexity w.r.t. the data size. Even though this requires an approximation through pruning, the results are still convincing enough to be referred to as virtually exact.

In order to perform exact inference over changepoint locations, it is not necessary that all involved computations are feasible in closed form. Instead, a sufficient requirement is

that certain elementary integrals, i.e.

$$\int \prod_{\ell=j}^i \mathbf{p}(\mathbf{Y}_\ell = y_\ell \mid \mathbf{X}_\ell = x) \mathcal{J}(dx) \quad (38)$$

for $0 < j \leq i \leq n$, can be computed either analytically or numerically. Therewith, the \mathbf{c}_{ji} 's and \mathcal{H}_{ji} 's can be derived in a purely arithmetic fashion.

In turn, for pointwise inference over segment heights by means of expectations w.r.t. functions \mathbf{f} , we require that

$$\int \mathbf{f}(x) \prod_{\ell=j}^i \mathbf{p}(\mathbf{Y}_\ell = y_\ell \mid \mathbf{X}_\ell = x) \mathcal{J}(dx)$$

is computable for $0 < j \leq i \leq n$.

Even though these calculations determine the desired posterior quantities uniquely, the observational densities are generally not unique, i.e. they may differ on null sets. This is related to the Borel-Kolmogorov paradox (Kolmogorov and Bharucha-Reid, 2018) and gives rise to ambiguous inference results. A common workaround is to impose topological continuity to the employed densities.

Furthermore, there might be applications where (38) is not applicable, for example, if it yields zero. This may occur when the observations depend somehow deterministically on segment heights that are continuous. Under such circumstances, we would refrain from our density based construction and resort to the measure theoretic notion of conditional distributions instead (Kolmogorov and Bharucha-Reid, 2018).

Our inference algorithms are implemented in discrete time exclusively. Even though an extension to continuous time is in principle viable, it involves dealing intensively with the backward residual times (Stewart, 2009). These quantities are trivial for exponentially distributed segment lengths, but generally of a delicate form.

The main constrained in this section is the imposed independence induced by change-points. Unfortunately, there are very basic applications that do not meet this constrained, like a continuous piecewise linear regression. The problem here is the exponential growth of the number of particles in the forward algorithm, which requires a sophisticated adaption

of the pruning algorithm to this far more demanding case. This topic lies beyond the scope of this thesis, but Fearnhead and Liu (2009) may provide a good starting point for such an algorithm. They elaborate an SMC approach (Doucet et al., 2001; Fearnhead and Clifford, 2003) for online changepoint detection.

As a general remark, applying mixtures to Lemma 11 is prone to fail since it yields mixtures for the \mathcal{H}_{ji} 's with an exponential growth in the number of components. If mixtures are involved, we should instead consider resorting to numerical integration by means of Lemma 10.

A further, rather inconspicuous restriction is that the observations are supposed to be mutually independent given the segment heights. Some applications may require instead that the observations follow a joint segmentwise and not a pointwise distribution.

This can be incorporated into our algorithms straight away if the segmentwise distributions are consistent in regards to the length of their segment. The consistency condition comes into play in the forward algorithm since it works on partial segments. If it doesn't hold, any partial segment would have to be represented explicitly as a mixture w.r.t. all possible segments containing this partial segment.

However, the forward algorithm may be sidestepped by means of a purely segmentwise approach like the sampling method for changepoint models described in Fearnhead (2006). In essence, it provides the \check{c}_{ji} 's. All our backward and pointwise algorithms may build on top of this since they rely on the \check{c}_{ji} 's directly.

A famous modeling approach, which hasn't been considered in this thesis, is to carry along the number of changepoints at each timepoint. For this purpose, we add a third index to c_{ji} , i.e. the number of changepoints that occurred hitherto.

Rigaiil et al. (2012) and Aston et al. (2012) employ this method to infer the pointwise probabilities of seeing certain changepoints. It is further necessary for the use of a prior for changepoint locations that pins the changepoint count down to a constant. Section B.1 elaborates such a uniform prior, which is ready to use within our algorithms. By assuming that k is the maximum number of changepoints considered, the space and time complexity is then increased by a factor of k regardless of pruning. Furthermore, the additional split

of the probabilities is numerically more challenging.

At several occasions, we saw that our methods share strong similarities with the popular Kalman-filter (and Kalman smoothing). The trajectory of the pointwise posterior expectations of the latent process poses a MAP for the Kalman-filter. Although this does not hold for changepoint models, Figure 11(a) showed that we may still gather valuable information from such expectations.

Moreover, Figure 11(a) reveals that the expected segment heights behave fairly steady amid segments even in the presence of strong outliers. This is because changepoint models adjust single segment heights closely to whole sequences of neighboring observations. At the same time, the expected segment heights can move very rapidly at changepoint locations.

In contrast, the volatility of the expectations in the Kalman-filter rely heavily on the variances of single observations. They can only move away abruptly if the corresponding variances are sufficiently small, i.e. the measurement is precise. Conversely, a steady behavior w.r.t. strong perturbations requires large variances.

Thus, changepoint models introduce a very specific behavioral pattern, which could perhaps be of significant practical interest. This is subject to further research.

The mathematics developed for the estimation of the success probability of the geometric distribution provides interesting links to the tremendously important topic of variable selection in statistical models (Guyon and Elisseeff, 2003). It can be stated as follows. Assume that we deal with a family of statistical models and n parameters with values in Θ . Each model is parameterized by a subsets of these n parameters. Given an observed data set, the task is to find one or more parsimonious models that represent the data well.

Therefore, we utilize their likelihoods, say $L_{I,\theta}$, where $I \subseteq \{1, \dots, n\}$ represents a subset of the parameters and $\theta \in \Theta^{\#I}$ their values. The famous Bayesian information criterion (BIC) (Schwarz et al., 1978) and the Akaike information criterion (AIC) (Akaike,

1973) are commonly used here. In this primal context, they would be defined as follows

$$\begin{aligned} \text{BIC}_I &= \ln(n) \cdot \#I - 2 \ln \left(\int L_{I,\theta} \mathcal{P}_I(d\theta) \right) \\ \text{AIC}_I &= 2 \cdot \#I - 2 \ln \left(\max_{\theta \in \Theta^{\#I}} \{L_{I,\theta}\} \right) \end{aligned}$$

whereby \mathcal{P}_I is a prior over a measurable space build from $\Theta^{\#I}$. Among all $I \subseteq \{1, \dots, n\}$, we would choose one that maximizes either BIC_I or AIC_I .

Both approaches perform a maximization that penalizes the cardinality of I . Therewith, they promote small and reasonable sets of parameters. While the BIC employs a prior in order to marginalize the parameters out, the AIC incorporates their maximum.

Unfortunately, the coefficient of $\#I$ is in both cases static and thus, these approaches appear to be a bit inflexible. This gets clearer through a natural analogy that uses the Bernoulli distribution as a prior for the elements of I , i.e. $\mathbb{P}(\mathbf{I} = I) = q^{\#I}(1 - q)^{n - \#I}$. The negative loglikelihood of this Bayesian model reads

$$-n \ln(1 - q) + \ln((1 - q)/q) \cdot \#I - \ln(L_{I,\theta}) \quad (39)$$

This yields an optimization problem similar to BIC and AIC, but with an arbitrary penalization for $\#I$.

In this context, BIC uses a prior for I that employs less than 1 parameter on average. This steers the estimation of the number of parameters to a parsimonious choice and thus focuses on specificity. However, this is inappropriate for changepoint problems and also for regression problems, where sensitivity towards individual parameters is of high value.

Hence, it might be reasonable to adjust the coefficient, say γ , of $\#I$ in (39) to the observed data in order to allow for a decent sensitivity. This could be done by means of the EM algorithm of Section 5.5.2 based on samples of \mathbf{I} (Mitchell and Beauchamp, 1988; George and McCulloch, 1993; Green, 1995). According to Lemma 19, in each EM step, we would just need to approximate the expected size of \mathbf{I} .

Having estimated γ , we may then maximize

$$\gamma \cdot \#I - \ln \left(\int L_{I,\theta} \mathcal{P}_I(d\theta) \right) \quad \text{or} \quad \gamma \cdot \#I - \ln \left(\max_{\theta} \{L_{I,\theta}\} \right)$$

w.r.t. I , a procedure which is sometimes referred to as an empirical Bayes method (Casella, 1985). This is subject to further research.

This whole section is accompanied by a Laplacian change in median example. It demonstrates the special case where the observational distribution is not part of the exponential family w.r.t. its segment height. In this case, simple recursive Bayes formulas are not available and thus, Lemma 10 combined with plain integration becomes the method of choice. Therewith, we have conveyed the ideas behind our algorithms and how to implement them in more challenging scenarios. Section 6.7 will conclude our Laplacian change in median example by a comparison with the model proposed by Fearnhead (2006).

6 Simultaneous credible regions for multiple changepoint locations

6.1 Outline

In this last big section, we present a novel approach for Bayesian changepoint models, that is based on changepoint samples. It facilitates the examination of the distribution of changepoints through a novel set estimator. For a given level α , we aim at smallest sets that cover all changepoints with a probability of at least $1 - \alpha$.

These so-called smallest simultaneous credible regions, computed for certain values of α , provide parsimonious representations of the possible changepoint locations. In addition, combining them for a range of different α 's enables very informative yet condensed visualizations. Therewith we allow for the evaluation of model choices and the analysis of changepoint data to an unprecedented degree.

This approach exhibits superior sensitivity, specificity and interpretability in comparison with highest density regions, marginal inclusion probabilities and confidence intervals inferred by `stepR`. Whilst their direct construction is usually intractable, asymptotically correct solutions can be derived from posterior samples. This leads to a novel NP-complete problem. Through reformulations into an Integer Linear Program we show empirically that a fast greedy heuristic computes virtually exact solutions.

6.2 Introduction

Detecting changepoints in a time series usually comes down to deciding on a set of changepoint locations. Thus, Bayesian frameworks aim to infer a set valued random variable that gives a reasonable representation of this decision (Fearnhead, 2006; Adams and MacKay, 2007; Lai and Xing, 2011). The non-deterministic nature of these so-called *posterior random changepoints* expresses the uncertainty of their location.

Rigaill et al. (2012) illustrates this uncertainty by means of a Bayesian model with exactly two changepoints. They plot for all possible pairs of timepoints the posterior

probability of being these changepoints. The results indicate both that posterior random changepoints are highly dependent and that generally more than one combination is likely. Unfortunately, this approach is not suitable to monitor the distribution of more than two changepoints. It is a crucial fact that the space of possible changepoint locations is very high-dimensional even for time series of moderate size. Thus, an extensive exploration is a nontrivial task.

In Bayesian research, summaries of changepoint locations, uncertainty measurements or model selection criteria are often provided by means of marginal changepoint probabilities. Rigai et al. (2012) gives a general consideration of this approach, but it has always enjoyed great popularity in the changepoint community. See, for example, Perreault et al. (2000); Lavielle and Lebarbier (2001); Tourneret et al. (2003); Fearnhead (2006); Hannart and Naveau (2009); Fearnhead and Liu (2009); Lai and Xing (2011); Aston et al. (2012); Nam et al. (2012). Marginal changepoint probabilities, as shown in changepoint histograms, simply consist of the probabilities for a changepoint at each timepoint and thus, they are pointwise statements. However, due to the uncertainty of their location, (even single) changepoints cannot be explored comprehensively by pointwise statements. On these grounds, we present a novel approach that incorporates all possible changepoint locations simultaneously.

Let $y = (y_1, \dots, y_n)$ be a time series and let $\mathbf{C} \subseteq \{1, \dots, n\}$ be the posterior random changepoints. A region that contains all changepoints simultaneously with a probability of at least $1 - \alpha$ is an $A \subseteq \{1, \dots, n\}$ with $\mathbb{P}(\mathbf{C} \subseteq A) \geq 1 - \alpha$. We call such an A a *simultaneous α level credible region* or simply *credible region* if there is no risk of confusion. It provides a sensitive assessment of the whole set of possible changepoint locations. However, in order to obtain specific assessments as well, we seek a smallest such region. This means we seek an element of

$$\mathfrak{R}(\alpha, \mathbf{C}) = \operatorname{argmin}_{A \subseteq \{1, \dots, n\}} \left\{ \#A \mid \mathbb{P}(\mathbf{C} \subseteq A) \geq 1 - \alpha \right\} \quad (40)$$

whereby $\#A$ is the cardinality of A .

Figure 15 demonstrates smallest credible regions by means of an example. The data

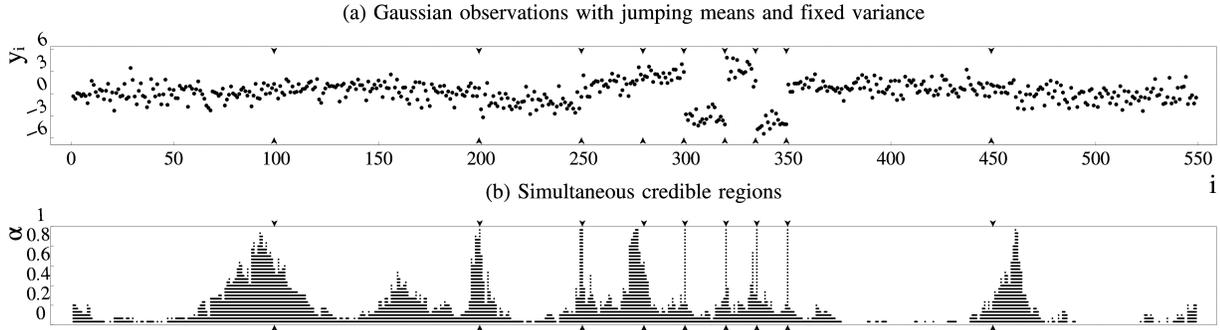


Figure 15: (a) A time series. (b) 29 approximated smallest simultaneous α level credible regions, each of them drawn as a broken horizontal line.

points in (a) were drawn independently from a normal distribution having a constant variance of 1 and mean values that are subject to successive changes. The true changepoints are marked by small vertical arrows. To build an exemplary Bayesian model here, we choose a prior for the changepoint locations and mean values. The time from one changepoint to the next is geometrically distributed with success probability $\frac{3}{550}$ and the mean values are distributed according to $\mathcal{N}(0, 25)$.

Smallest credible regions are visualized in (b). For each $\alpha \in \left\{ \frac{1}{30}, \dots, \frac{29}{30} \right\}$ the plot shows one such (approximated) region as a collection of horizontal lines. The value of $\alpha \in [0, 1]$ controls the size of the regions and thus, governs the trade of between specificity and sensitivity.

The broadness of the credible regions around a true changepoint expresses the uncertainty of its locations. At the same time, a pointed shape reveals that the model is in favor of certain changepoint locations. Existing visualization techniques, like changepoint histograms, are unable to go beyond these two estimates. However, it is also crucial to get an impression of the sensitivity of the model towards a true changepoint. We can examine this by looking at the height of the peak that relates to the true changepoint. The higher the peak, the higher the model's sensitivity. We refer to this height as the importance of the true changepoint. The importances of the true changepoints in Figure 15 are always

larger than 0.9.

Of course, importance (as well as broadness and shape) can also be used for changepoint data without true changepoints. There, we look at the peaks that belong to the features of interest. Examples for features, which may occur in practice, are changes in mean, variance, slope or any other change in distribution. Furthermore, anomalies like outliers are of concern as well although they are usually supposed to be skipped by the model. Figure 15 shows a nuisance feature in form of a small irregularity at around 170 with an importance of around 0.4.

By means of importance we can conveniently evaluate if the model is sensitive towards the desired features but skips the nuisance ones. Conversely, we can also detect the relevant features in changepoint data on the basis of a given model. Most notably, this does not require any previous knowledge about the number of features or their positions. This novel concept is one of the main outcomes of the present work. It allows for a much more detailed analysis of changepoint models and changepoint data than hitherto possible. Figure 15(b) demonstrates this. It shows that the changepoint model in use is sensitive towards the desired true changepoints and specific towards random distortions in the data.

The outline of this work is as follows. In Section 6.3 we consider the above problem from a general statistical and algorithmic viewpoint. We examine alternative approaches in Section 6.4. Section 6.5 starts with an overview over sampling strategies in changepoint models and deals with importance in a broader and more formal way. Afterwards, we compare our results with the existing approaches. We examine Dow Jones returns and demonstrate how credible regions can be used in order to perform model selection in Section 6.6. Finally, we discuss our results in Section 6.8.

6.3 The Sample Based Problem

In this section we investigate the computational and mathematical foundations of our approach. We show how to approximate credible regions in an asymptotic manner, examine the complexity and elaborate suitable algorithms.

We are given an arbitrary random set $\mathbf{C} \subseteq \{1, \dots, n\}$. Let \mathcal{P} be the distribution of \mathbf{C} , i.e. $\mathcal{P}(A) = \mathbb{P}(\mathbf{C} = A)$ for all $A \subseteq \{1, \dots, n\}$. In this section, we are mainly concerned with the task of finding an element of $\mathfrak{R}(\alpha, \mathbf{C})$ (see Equation (40)), where α can take any value in $[0, 1]$.

Deriving $\mathbb{P}(\mathbf{C} \subseteq A)$ is generally intractable since it arises from a summation of $2^{\#A}$ probabilities. Likewise, finding an element of $\mathfrak{R}(\alpha, \mathbf{C})$ requires a search over the 2^n subsets of $\{1, \dots, n\}$. Hence, we address this problem in an approximate manner by using relative frequencies instead of probabilities.

Definition 2. For $A \subseteq \{1, \dots, n\}$ and $s_1, \dots, s_m \subseteq \{1, \dots, n\}$, let

$$\mathbf{r}(A, s_{1:m}) = \frac{1}{m} \sum_{i=1}^m \mathbb{1}\{s_i \subseteq A\}$$

where $\mathbb{1}\{\dots\} \in \{0, 1\}$ is the indicator function, which is equal to 1 iff the bracketed condition is true. $\mathbf{r}(A, s_{1:m})$ corresponds to the relative frequency by which the subsets are covered by A .

In the following theorem, we will show that having independent samples $s_1, \dots, s_m \subseteq \{1, \dots, n\}$ from \mathcal{P} , we can approximate an element of $\mathfrak{R}(\alpha, \mathbf{C})$ by finding an element of

$$\mathfrak{S}(\alpha, s_{1:m}) = \operatorname{argmin}_{A \subseteq \{1, \dots, n\}} \left\{ \#A \mid \mathbf{r}(A, s_{1:m}) \geq 1 - \alpha \right\}$$

We denote this problem as the *Sample Based Problem (SBP)*.

Theorem 4. Let $\mathbf{S}_1, \mathbf{S}_2, \dots \subseteq \{1, \dots, n\}$ be independent random sets distributed according to \mathcal{P} and $\alpha \in [0, 1]$. If there exists an $A_\infty \in \mathfrak{R}(\alpha, \mathbf{C})$ with $\mathbb{P}(\mathbf{C} \subseteq A_\infty) > 1 - \alpha$, then $\mathfrak{S}(\alpha, S_{1:m}) \subseteq \mathfrak{R}(\alpha, \mathbf{C})$ eventually almost surely.

Proof. Let $A \subseteq \{1, \dots, n\}$, since the $\mathbb{1}\{\mathbf{S}_i \subseteq A\}, i = 1, 2, \dots$ are i.i.d. with finite expectations, the strong law of large numbers states that $\lim_{m \rightarrow \infty} \mathbf{r}(A, \mathbf{S}_{1:m}) = \mathbb{P}(\mathbf{C} \subseteq A)$ almost surely. Since $2^{\{1, \dots, n\}}$ is finite, $\mathbf{r}(A, \mathbf{S}_{1:m})$ even converges to $\mathbb{P}(\mathbf{C} \subseteq A)$ for all $A \subseteq \{1, \dots, n\}$ almost surely.

Let $s_1, s_2, \dots \subseteq \{1, \dots, n\}$ be an arbitrary but fixed sequence with $\lim_{m \rightarrow \infty} \mathbf{r}(A, s_{1:m}) = \mathbb{P}(\mathbf{C} \subseteq A)$ for all $A \subseteq \{1, \dots, n\}$. For $A \subseteq \{1, \dots, n\}$ with $\mathbb{P}(\mathbf{C} \subseteq A) < 1 - \alpha$ we can pick an $m_A \in \mathbb{N}$ such that $\mathbf{r}(A, s_{1:m}) < 1 - \alpha$ for all $m > m_A$. Hence, $\mathfrak{S}(\alpha, s_{1:m}) \subseteq \{A \subseteq \{1, \dots, n\} \mid \mathbb{P}(\mathbf{C} \subseteq A) \geq 1 - \alpha\}$ holds for all $m > m_1 = \max\{m_A \mid \mathbb{P}(\mathbf{C} \subseteq A) < 1 - \alpha\}$. Furthermore, since $\mathbb{P}(\mathbf{C} \subseteq A_\infty) > 1 - \alpha$, we can pick an $m_0 \in \mathbb{N}$ with $\mathbf{r}(A_\infty, s_{1:m}) > 1 - \alpha$ for all $m > m_0$. Thus, by choosing $\ell = \max\{m_0, m_1\}$, we obtain $\mathfrak{S}(\alpha, s_{1:m}) \subseteq \mathfrak{R}(\alpha, \mathbf{C})$ for all $m > \ell$. \square

Since the exact value of α is usually irrelevant, the condition that there exists a certain solution can in most cases be neglected.

6.3.1 Reformulation of the Sample Based Problem and its complexity

In this section we consider the complexity of SBP, formulate it as an *Integer Linear Program* (ILP) and introduce a fast and fairly accurate approximation by a greedy method. The next theorem shows that there is no hope to find a polynomial time algorithm to solve SBP unless $P = NP$ (Garey and Johnson, 1979). The corresponding proof is provided in the appendix (Section B.2).

Theorem 5 (Marc Hellmuth). *The (decision version of) SBP is NP-complete.*

SBP is NP-hard and thus, there is no polynomial-time algorithm to optimally solve a given SBP-instance. Nevertheless, SBP can be formulated as an ILP and thus, SBP becomes accessible to highly efficient ILP solvers (Meindl and Templ, 2013). Such ILP solvers can be employed to optimally solve at least moderately-sized SBP-instances. The ILP formulation is as follows.

ILP (Marc Hellmuth). *Following the notion of Section 6.3, we are given a set of samples $s_1, \dots, s_m \subseteq \{1, \dots, n\}$ and an $\alpha \in [0, 1]$. We declare binary variables $U_i, F_j \in \{0, 1\}$ for all $i \in \{1, \dots, n\}$ and $j \in \{1, \dots, m\}$. For $i \in \{1, \dots, n\}$ let $\mathcal{D}(s_{1:m}, i) = \{j \mid i \in s_j\}$*

denote the set of samples that contain i . Now define the following constraints

$$\text{(I)} \quad \sum_{j=1}^m F_j \geq m \cdot (1 - \alpha)$$

$$\text{(II)} \quad \forall i \in \{1, \dots, n\} : \sum_{j \in \mathcal{D}(s_{1:m}, i)} (1 - F_j) \geq \#\mathcal{D}(s_{1:m}, i) \cdot (1 - U_i)$$

under which the objective function $\sum_{i=1}^n U_i$ needs to be minimized.

Having computed an optimum, $A = \{i \mid U_i = 1\} \in \mathfrak{S}(\alpha, s_{1:m})$, i.e. this set is a solution of SBP.

The binary variable F_j represents the j -th sample. If $F_j = 1$ then $s_j \subseteq A$. Thus, the constraint **(I)** states that A covers at least $m \cdot (1 - \alpha)$ samples, which is equivalent to $\mathbf{r}(A, s_{1:m}) \geq 1 - \alpha$ in the definition of $\mathfrak{S}(\alpha, s_{1:m})$. The constraint **(II)** states that timepoint i can only be dropped from A , if $i \in s_j$ implies that $F_j = 0$. From the perspective of SBP these constraints are obviously necessary. The proof of sufficiency is provided in the appendix (Section B.3).

A benchmark of several ILP solvers can be found in Meindl and Templ (2013). Following this advice, we use CPLEX V12.6.3 for Linux x86-64 IBM (2016) on a Lenovo Yoga 2 Pro (8GB Ram, 4 x 1.8GHZ Intel CPU) to solve our ILP instances.

6.3.2 A greedy heuristic

To provide an alternative way to address SBP, we now resort to a simple greedy heuristic. This approach starts with the whole set of timepoints and greedily removes all timepoints iteratively. The greedy rule used here chooses a timepoint that is contained in the smallest number of samples. In the subsequent steps, these samples will be ignored.

Greedy. Let $A_0 = \{1, \dots, n\}$. Compute $A_{\ell+1} = A_\ell \setminus \{k_{\ell+1}\}$ iteratively where $k_{\ell+1} \in \operatorname{argmax}_{i \in A_\ell} \{\mathbf{r}(A_\ell \setminus \{i\}, s_{1:m})\}$ until $A_{\ell+1} = \emptyset$. Use A_ℓ as a solution proposal for SBP for any α with $\mathbf{r}(A_{\ell+1}, s_{1:m}) < 1 - \alpha \leq \mathbf{r}(A_\ell, s_{1:m})$.

Pseudocode 4 provides a possible implementation of **Greedy**. The runtime of this algorithm is $\mathcal{O}(n^2 m \ln(m))$. To see this, observe that the first for-loop runs n times and

Pseudocode 4 Greedy Algorithm for SBP

```
1: Input: List  $s[j]$  that represents the sample  $j$ ,  $1 \leq j \leq m$ 
2:         Sorted lists  $\mathcal{C}[i]$  that represents the samples that contain timepoint  $i$ ,  $1 \leq i \leq n$ 
3: Output:  $A_1, \dots, A_n$  ▷  $A_{\ell+1} = A_\ell \setminus \{k_{\ell+1}\}$  with  $k_{\ell+1} \in \operatorname{argmax}_{i \in A_\ell} \{r(A_\ell \setminus \{i\}, s_{1:m})\}$ 
4:  $A_0 = \{1, \dots, n\}$ 
5: for  $\ell = 0, \dots, n - 1$  do ▷ Discard all timepoints iteratively
6:      $i = \operatorname{argmin} \{\mathcal{C}[k].\text{length} \mid k \in A_\ell\}$  ▷ Complexity  $\mathcal{O}(n)$ 
7:      $A_{\ell+1} = A_\ell \setminus \{i\}$ 
8:     for  $j = 1, \dots, \mathcal{C}[i].\text{length}$  and  $k = 1, \dots, s[j].\text{length}$  do ▷ Remove all samples that contain  $i$  from  $\mathcal{C}$ 
9:          $\mathcal{C}[s[j][k]].\text{remove}(j)$  ▷ Removing has a complexity of  $\mathcal{O}(\ln(m))$ 
10:    end for
11: end for
```

the assignment in Line 6 needs $\mathcal{O}(n)$ time. The second for-loop runs $\mathcal{O}(mn)$ times and the removal of an element in Line 9 can be done in $\mathcal{O}(\ln(m))$ time. Thus, the overall time-complexity is $\mathcal{O}(n^2 + n^2m \ln(m)) = \mathcal{O}(n^2m \ln(m))$.

Although, the worst case runtime-complexity is more or less cubic, Algorithm 4 can be well applied in practice for the following reasons. Assume that every sample has exactly k elements uniformly distributed over $\{1, \dots, n\}$. In this case, the number of samples containing i is binomially distributed with parameters $\frac{k}{n}$ and m . Thus, their average number is $m \cdot \frac{k}{n}$. Finally, the average complexity reduces to $\mathcal{O}(k^2m)$ if we use a hash table to store elements of the $\mathcal{C}[i]$'s, a sorted list to store their lengths and assume that $n \ln(n) < k^2m$. In turn, if the elements of the samples are non-uniformly distributed and k only represents their maximal cardinality the runtime will be shortened even further.

Greedy provides credible regions for all α values at once. However, in contrast to solutions to the ILP, the resulting credible regions will be nested w.r.t. increasing α values.

6.4 Alternative approaches

In this section we explain two alternative Bayesian approaches that can be derived from **C** as well. In the first approach we join highest density regions in order to infer credible regions. Highest density regions can be considered as a straightforward simultaneous tool to

explore distributions of interest (Held, 2004). The second approach, marginal changepoint probabilities, is a straightforward pointwise tool in the context of changepoints.

Besides this, it should be noted that Guédon (2015) addresses uncertainty of changepoint locations through the entropy of \mathbf{C} .

6.4.1 Highest density regions

A highest density region (HDR) is a certain subset of a probability space with elements having a higher density value than elements outside of it. Such a subset can be utilized to characterize and visualize the support of the corresponding probability distribution (Hyndman, 1996). In a Bayesian context, $(1 - \alpha)$ -HDR's are often used as simultaneous α level credible regions (Held, 2004). In this section, we examine HDR's in general and for the case of random subsets of $\{1, \dots, n\}$.

Let \mathbf{X} be a random variable with a density \mathbf{p} . For $\alpha \in [0, 1]$, let q_α be an α -quantile of $p(\mathbf{X})$, i.e. $\mathbb{P}(\mathbf{p}(\mathbf{X}) \leq q_\alpha) \geq \alpha$ and $\mathbb{P}(\mathbf{p}(\mathbf{X}) \geq q_\alpha) \geq 1 - \alpha$.

Definition 3. *The set $\{x \mid \mathbf{p}(x) \geq q_\alpha\}$ is referred to as the $(1 - \alpha)$ -HDR of \mathbf{X} .*

The $(1 - \alpha)$ -HDR is a smallest subset of the state space with a probability of at least $1 - \alpha$ (Box and Tiao, 1973).

Now we consider $\mathbf{X} = \mathbf{C}$. Our credible regions are subsets of $\{1, \dots, n\}$, whereas the $(1 - \alpha)$ -HDR of \mathbf{C} would be a subset of $2^{\{1, \dots, n\}}$. For the purpose of comparison, we join all successes in the HDR:

Definition 4. *Let $s_1, \dots, s_\ell \subseteq \{1, \dots, n\}$ be the $(1 - \alpha)$ -HDR of \mathbf{C} . We refer to $\bigcup_{i=1}^\ell s_i$ as the joined $(1 - \alpha)$ -HDR of \mathbf{C} .*

Let A be the joined $(1 - \alpha)$ -HDR and \mathcal{P} the distribution of \mathbf{C} . Then

$$\mathbb{P}(\mathbf{C} \subseteq A) = \mathcal{P}(2^A) \geq \sum_{i=1}^{\ell} \mathcal{P}(s_i) \geq 1 - \alpha$$

Even though, s_1, \dots, s_ℓ have high probabilities, $2^{\#A}$ may be much larger than ℓ . Likewise, $\mathbb{P}(\mathbf{C} \subseteq A)$ might be considerably larger than $1 - \alpha$. Thus, A might be substantially larger

than elements of $\mathfrak{R}(\alpha, \mathbf{C})$. In Section 6.5 we will see that this happens, especially for small α .

Unfortunately, in many cases we cannot compute HDR's directly and therefore, we use an approximation scheme (Held, 2004). Let now $s_1, \dots, s_m \subseteq \{1, \dots, n\}$ be independent samples from the distribution of \mathbf{C} that are arranged in descending order according to their values under \mathcal{P} , i.e. $i > j \Rightarrow \mathcal{P}(s_i) \leq \mathcal{P}(s_j)$. We use $\bigcup_{i=1}^{\ell} s_i$ with $\ell = \lceil m \cdot (1 - \alpha) \rceil$ as an approximate joined $(1 - \alpha)$ -HDR.

6.4.2 Marginal changepoint probabilities

Here we consider the marginal changepoint probabilities, i.e. $\mathbb{P}(i \in \mathbf{C})$ for $i = 1, \dots, n$. They can be used to derive subsets of $\{1, \dots, n\}$ that are closely related to credible regions.

Lemma 27. $\{i \mid \mathbb{P}(i \in \mathbf{C}) > \alpha\}$ is a subset of all elements of $\mathfrak{R}(\alpha, \mathbf{C})$.

Proof. Since $\mathbb{P}(i \in \mathbf{C}) > \alpha$ implies $\mathbb{P}(\mathbf{C} \subseteq \{1, \dots, n\} \setminus i) < 1 - \alpha$, i has to be part of any α level credible region. \square

The Bonferroni correction (Dunnnett, 1955) can be applied to construct a credible region. Let $\mathfrak{B}(\alpha, \mathbf{C}) = \left\{i \mid \mathbb{P}(i \in \mathbf{C}) > \frac{\alpha}{n}\right\}$.

Lemma 28. $\mathbb{P}(\mathbf{C} \subseteq \mathfrak{B}(\alpha, \mathbf{C})) \geq 1 - \alpha$

Proof. Let $A = \mathfrak{B}(\alpha, \mathbf{C})^c$. We conclude that $1 - \mathbb{P}(\mathbf{C} \subseteq A^c) = \mathbb{P}(\mathbf{C} \cap A \neq \emptyset) \leq \sum_{i \in A} \mathbb{P}(i \in \mathbf{C}) \leq \alpha$. Thus, $\mathbb{P}(\mathbf{C} \subseteq \mathfrak{B}(\alpha, \mathbf{C})) = \mathbb{P}(\mathbf{C} \subseteq A^c) \geq 1 - \alpha$ applies. \square

Pointwise statements suffer from their inability to reflect dependencies. To see this, we assume that $\mathbb{P}(\#\mathbf{C} = 1) = 1$ and $\mathbb{P}(i \in \mathbf{C}) = \frac{1}{n}$ for all $i = 1, \dots, n$. Single timepoints are strongly dependent, e.g. $i \in \mathbf{C}$ implies $j \notin \mathbf{C}$ for all $j \neq i$. Clearly, $\mathfrak{B}(\alpha, \mathbf{C}) = \{1, \dots, n\}$ for all $\alpha < 1$. Moreover, $\{i \mid \mathbb{P}(i \in \mathbf{C}) > \alpha\} = \emptyset$ for $\alpha \geq \frac{1}{n}$. In contrast to this, elements of $\mathfrak{R}(\alpha, \mathbf{C})$ become smaller if α becomes larger. More precisely, $\#A = \lceil (1 - \alpha)n \rceil$ for $A \in \mathfrak{R}(\alpha, \mathbf{C})$. This is due to the fact that pointwise statements ignore dependencies, whereas simultaneous statements incorporate them. Hence, in practice $\mathfrak{B}(\alpha, \mathbf{C})$ may be

much broader than elements of $\mathfrak{R}(\alpha, \mathbf{C})$ and $\{i \mid \mathbb{P}(i \in \mathbf{C}) > \alpha\}$ may be much smaller. In Section 6.5 we will give an example that underpins this claim.

6.5 Multiple changepoints and credible regions in more detail

In this section we discuss sampling strategies in changepoint models and we consider the concept of importance in a more formal way. Afterwards we reconsider the aforementioned example in order to compare existing approaches with the credible region approach. Finally, we investigate the performance of **Greedy** and the convergence speed of SBP empirically.

6.5.1 Importance revisited

Characterizing the features in the data that should or should not trigger a changepoint enables the formulation of changepoint problems in the first place. An intrinsic property of virtually every changepoint problem is that the location of the changepoint, placed as a consequence of a feature, is uncertain. Hence, a feature in the data is usually related to a set of locations instead of a single location. In order to deal with this fact, we require a simultaneous approach that considers all changepoint locations that belong to the same feature in a coherent fashion.

To this end, as intimated already, importance provides an estimate of the sensitivity of changepoint models towards features in changepoint data. However, it remains to substantiate this claim mathematically. We start with our definition of sensitivity. The sensitivity of a changepoint model towards a feature in the data is defined as the probability by which posterior random changepoints occur as a consequence of this feature. Thus, if $A \subseteq \{1, \dots, n\}$ represents the set of timepoints which are related to the feature, the sensitivity equates to $\mathbb{P}(\mathbf{C} \not\subseteq A^c)$. In general, determining the set A is a fuzzy task. However, since changepoint data usually exhibits the relevant features only occasionally, in most cases we can easily specify these sets to a sufficient extent.

In the introduction the importance of a feature is roughly defined as the height of the peak that relates to this feature. This means we derive importance from a given range of

credible regions. Assume that we are given a whole range of credible regions $R_\alpha \in \mathfrak{R}(\alpha, \mathbf{C})$ for all $\alpha \in [0, 1]$. We now define the importance of a feature as $\hat{\alpha} = \inf\{\alpha \mid R_\alpha \subseteq A^c\}$, where A is again the set of locations that belongs to the feature. With this the following lemma applies.

Lemma 29. *The importance of a feature is an upper bound for the sensitivity of the changepoint model towards this feature.*

Proof. If there exists an $\epsilon > 0$ with $R_\alpha \subseteq A^c$ for all $\alpha \in]\hat{\alpha}, \hat{\alpha} + \epsilon[$, we conclude

$$\begin{aligned} 1 - \alpha &\leq \mathbb{P}(\mathbf{C} \subseteq R_\alpha) \leq \mathbb{P}(\mathbf{C} \subseteq A^c) \text{ for all } \alpha \in]\hat{\alpha}, \hat{\alpha} + \epsilon[\\ \Leftrightarrow 1 - \hat{\alpha} &\leq \mathbb{P}(\mathbf{C} \subseteq A^c) \quad \Leftrightarrow \quad \hat{\alpha} \geq \mathbb{P}(\mathbf{C} \not\subseteq A^c) \end{aligned}$$

In the case where $R_{\hat{\alpha}} \subseteq A^c$, we may apply the second part of the above equation directly. \square

Since the credible regions are not necessarily nested, in some rare cases, the importance read by means of the above definition may not correspond to a global peak in the credible regions. Furthermore, if credible regions are only drawn for an incomplete range of α 's, we may slightly underestimate the heights of the peaks. In exchange for these negligible inaccuracies, we obtain an intuitive, visual estimate that provides great insights into the distribution of changepoints and is thus of high practical value.

6.5.2 Exemplary comparison

Figure 16(a) displays the same dataset as Figure 15(a). There are 6 obvious true changepoints at 200, 250, 300, 320, 335 and 350 having a large jump height and two true changepoints at 280 and 450 with smaller jump heights of 1 and 0.8, respectively. Besides, there is a true changepoint at 100 with an even smaller jump from 0 to 0.5.

We now want to compare the credible regions approach with the highest density regions approach, marginal jump probabilities and confidence intervals inferred by the R package `stepR`. We use the same model as in the introduction. Figure 16(b) displays the credible regions with respect to $\alpha = \frac{1}{30}, \dots, \frac{29}{30}$ inferred by solving the ILP.

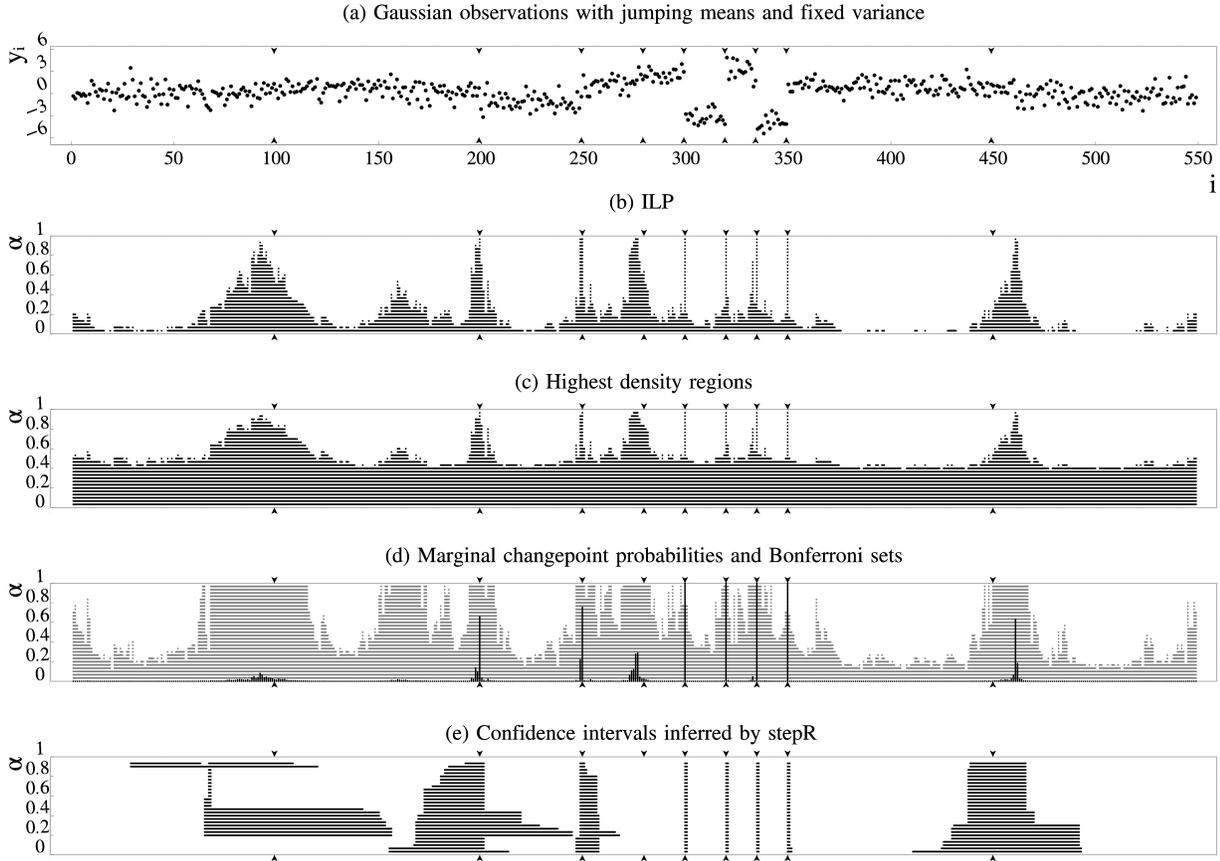


Figure 16: (a) Simulated independent Gaussian observations with variance 1 and successive changes in mean. The true changepoints are marked by vertical arrows. (b) Solutions derived by solving the ILP using 10^5 samples. (c) Joined $(1 - \alpha)$ -HDR's using 10^7 samples. (d) Marginal changepoint probabilities in black and $\mathfrak{B}(\alpha, \mathbf{C})$ in gray. (e) Joined confidence intervals inferred by `stepR`.

In the same fashion as before, Figure 16(c) displays several approximated joined HDR's derived from 10^7 samples. This confirms that joining the elements of an HDR leads to larger subsets of $\{1, \dots, n\}$ compared to elements of $\mathfrak{R}(\alpha, \mathbf{C})$. At an α level less than 0.4, the joined HDR already covers most of the timepoints and thus, holds no special information about the changepoint locations anymore.

Figure 16(d) displays the marginal changepoint probabilities, i.e. $\mathbb{P}(i \in \mathbf{C})$ highlighted

in black. Unfortunately, they don't reflect the possible set of changepoints very well since they appear to be too sparse. Furthermore, due to a dispersal of the probabilities they are not able to express the sensitivity towards the true changepoints (see, for example, the changepoints at 100 and 280). This leads to the conclusion that this approach suffers from a lack of interpretability and is less sensitive than the credible regions approach.

Additionally, Figure 16(d) shows several credible regions corresponding to $\mathfrak{B}(\alpha, \mathbf{C})$. As prognosticated in Section 6.4.2 these regions are very broad compared to smallest credible regions and thus, less specific.

Figure 16(e) displays several joined confidence intervals inferred by `stepR` (Hotz and Sieling, 2016). `stepR` first estimates the number of changepoints and produces one confidence interval for each changepoint. The plot shows the union of these confidence intervals. Although confidence sets and credible regions are different by definition, they intend to make similar statements.

Unfortunately, `stepR` does not forecast a confidence interval for the true changepoint at 280. Furthermore, the disappearance of certain changepoint locations at decreasing α values seems somewhat confusing and hard to make sense of. So, the shapes, broadnesses and importances read from the confidence sets are if at all only of little informative value. Therefore, the authors of this work consider the confidence sets inferred by `stepR` as unsuitable for this purpose, because they suffer from a grave lack of interpretability.

The supplement contains a collection of pictures similar to Figure 16. There, the data was repeatedly generated using the same changepoint locations and mean values. It becomes apparent that these kind of plots, produced with `stepR`, are almost consistently of poor quality.

In Section 6.5.1 we show that the importance of a feature in the data is an upper bound for the sensitivity of the model towards this feature. We can examine the differences between these two estimates empirically by means of the above example. Sensitivity and importance match for all the true changepoints except the first one. The importance of the first changepoint is approximately 0.94, whereas its sensitivity with regard to the interval from 0 to 130 is approximately 0.72. That's a deviation of around 0.2. The small

irregularity at around 170 results in an even bigger deviation of around 0.3.

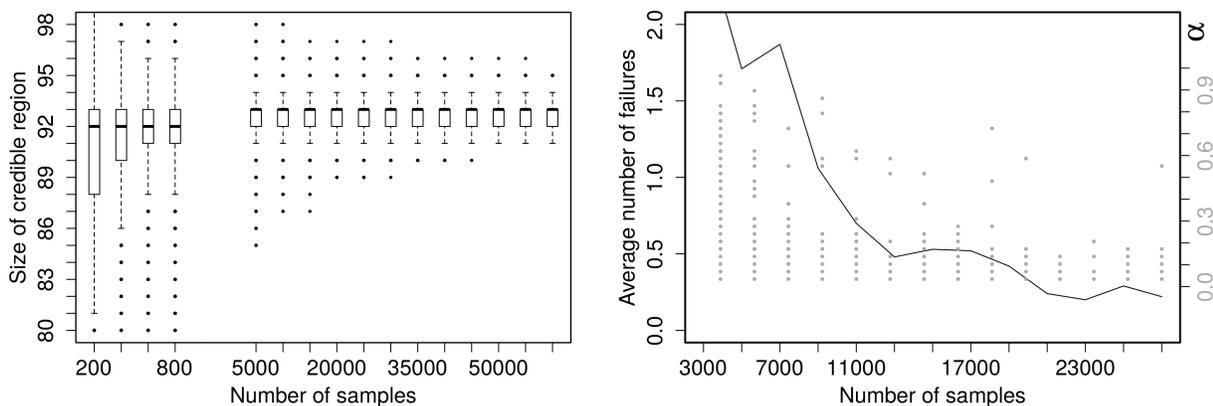


Figure 17: (left) An empirical convergence proof of SBP. Boxplots over the sizes of the elements of $\mathfrak{S}(0.3, s_{1:m})$ where m is varying according to the x-axis. (right) Illustration of the accuracy of solutions inferred by **Greedy**. The black graph illustrates the average number of failures within 29 different credible regions (left y-axis). The gray dots mark the α values (right y-axis) where failures occurred.

6.5.3 Empirical proof of convergence and the accuracy of **Greedy**

Figure 17 (left) demonstrates how solutions to SBP evolve with increasing sample size. For several sample sizes m it shows boxplots (using 1000 repetitions) over the size of elements of $\mathfrak{S}(0.3, s_{1:m})$. It can be observed that the sizes of the credible regions increase with an increasing sample size. This is due to the fact that we need a certain number of samples to cover the possible changepoint locations satisfactorily. However, at the same time, we observe an increasing preciseness, which is a result of Theorem 4.

For different sample sizes we compared for each $\alpha \in \left\{ \frac{1}{30}, \dots, \frac{29}{30} \right\}$ the solutions provided by solving the ILP with those provided by **Greedy**. In Figure 17 (right), the black graph illustrates the average number of α 's (using 100 repetitions) where **Greedy** failed to provide an optimal solution. At higher sample counts, there are less than 0.3 of the 29 credible

regions wrong. The gray points represent the α values (right axis) where the ILP was able to compute smaller regions than **Greedy**.

Hence, **Greedy** performs virtually exact on this changepoint problem. However, at smaller sample counts it gets more frequently outwitted by random. Fortunately, this shows that if **Greedy** does not compute an ideal region for a certain α , it will compute correct ones later again. This is due to the fact that credible regions for changepoint locations are roughly nested (like the solutions provided by **Greedy**). Only in some cases it is possible through solving the ILP to remove certain timepoints a little bit earlier than **Greedy**. Furthermore, because **Greedy** removes samples having a high number of changepoints fast, it works well together with changepoint models, since they prefer to explain the data through a small number of changepoints. Thus, we can conjecture that **Greedy** will perform just as well in most changepoint scenarios.

6.6 An example of use for model selection

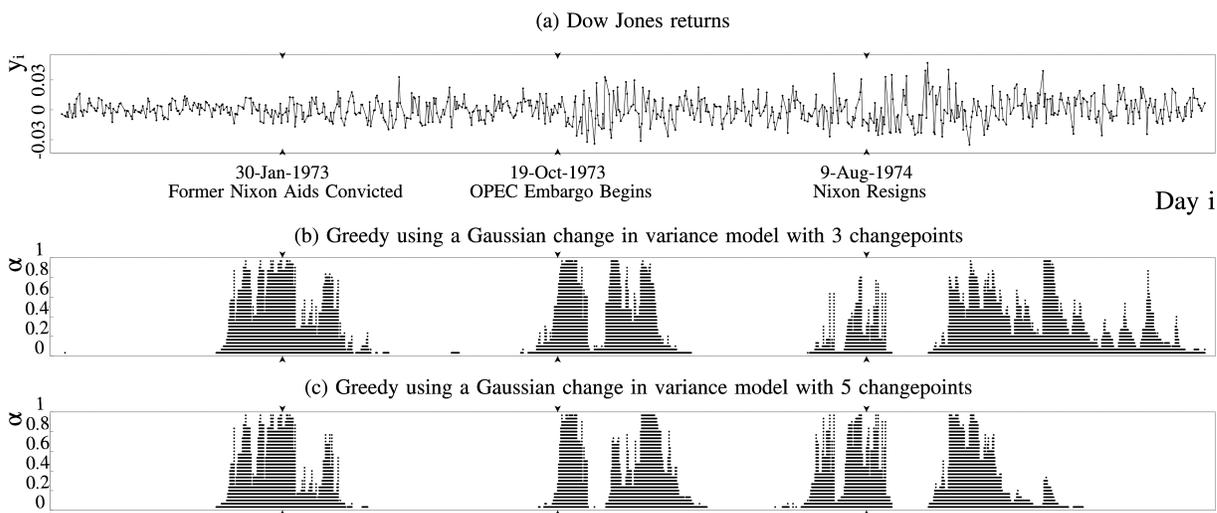


Figure 18: (a) Dow Jones returns. (b) and (c) Different regions provided by **Greedy**.

Now we examine Dow Jones returns observed between 1972 and 1975 (Adams and MacKay, 2007), see Figure 18(a). There are three documented events highlighted on

January 1973, October 1973 and August 1974. The data is modeled as normally distributed with constant mean equal to 0 and jumping variances. The variances are distributed according to an inverse gamma distribution with parameters 1 and 10^{-4} (see also Adams and MacKay (2007)). Instead of allowing a random number of changepoints, this time we predetermine the number of changepoints to three respectively five and our aim is to compare these two model choices by means of credible regions.

As we can see in Figure 18(b), the regions become fairly broad especially in the last third of the picture giving rise to additional, nonsensical changepoint locations. The reason for this is that assuming only one changepoint after the second event, yields to a misjudgment of the third or fourth variance. Thus, the third changepoint becomes superfluous and its location highly uncertain.

In contrast, in the case of five changepoints in (c) the illustration turns out to be much more differentiated. The regions stay fairly narrow even for very small values of α . Just a quick look at the credible regions immediately reveals that fixing the number of changepoints may impair the inference dramatically. Thus, whenever possible, we should use a model that allows the number of changepoints to be determined at random.

6.7 Simultaneous credible regions in the well-log example

One last time, we consider the well-log example of Section 5.2.4 and the changepoint model developed there. Figure 19 (a) depicts the data. The regions inferred by **Greedy** from 10^7 samples are shown in (b).

The credible regions clearly reflect that the model was trained for the sake of changepoint inference via MAP estimation and not uncertainty. In Section 5.4.1 we have already seen that the expected number of changepoints equates to 17.8. Thus, within the samples we can find plenty of unnecessary changepoints. The plot shows that these nuisance changepoints are mainly located in the first and last part of the data, but also around 1300 and 2200.

We may now try to reduce the uncertainty by increasing the segmental variance in order to obtain a model with a better fit. However, several experiments have shown that

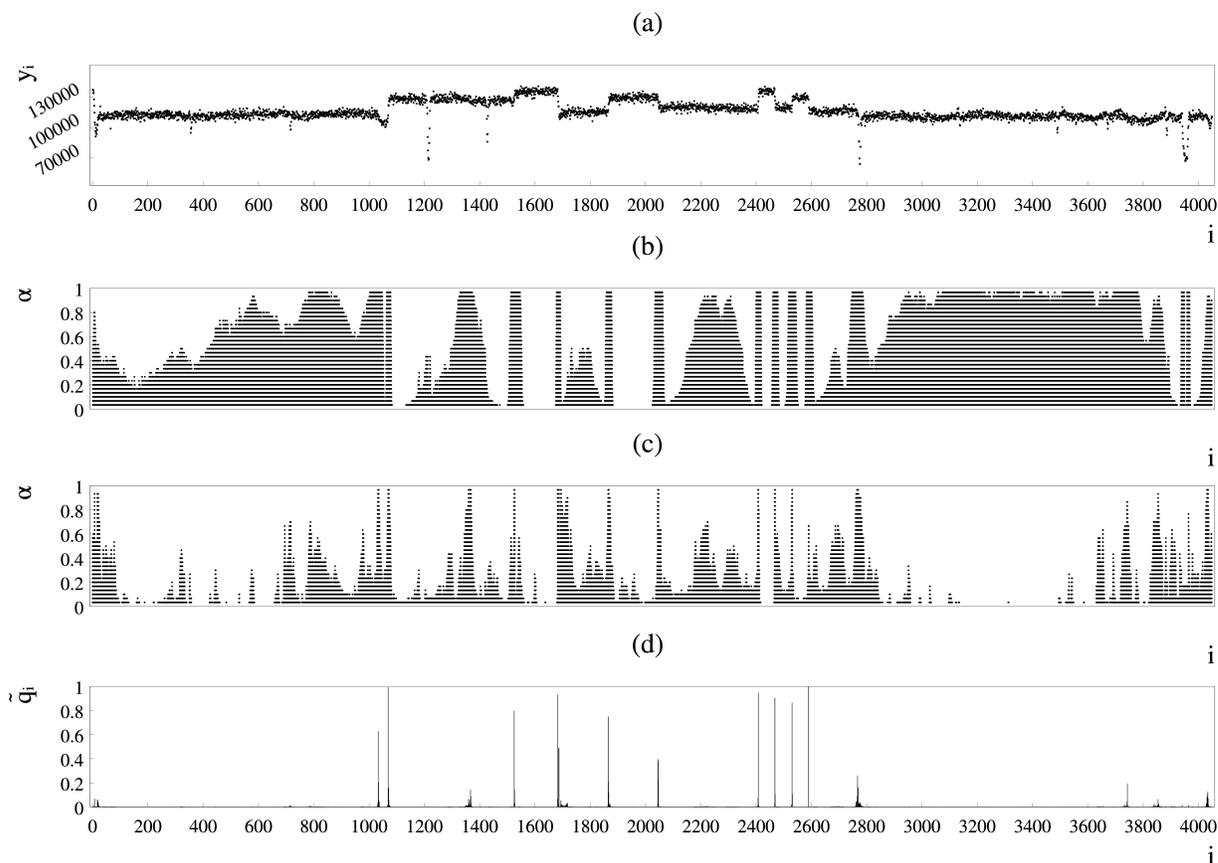


Figure 19: (a) Well-Log data. (b) and (c) credible regions inferred by **Greedy** for the Laplacian change in median model and the model of Fearnhead (2006), respectively. (d) marginal changepoint probabilities for the model of Fearnhead (2006).

an appropriate reduction of uncertainty gives less than twelve changepoints when we use the MAP.

This suggests that a fundamentally different model should be employed in case the focus is more on uncertainty or authenticity. Fearnhead (2006) proposes a changepoint model that employs a piecewise Kalman-filter marginalized over the variance of the latent process. It is obviously more naturally adapted to the dip at around 3800 and the small shifts in mean. However, the outliers are challenging since the Kalman-filter is based on the normal distribution. Therefore, Fearnhead (2006) eliminates the outliers in a preprocessing

step, which leaves gaps in the data.

19 (c) shows the corresponding regions inferred by **Greedy** from 10^7 samples and (d) the marginal changepoint probabilities. The uncertainty appears to be much lower here. Especially on the first and last third of the data, there is only a very low tendency to place changepoints. However, at around 1300 there is a high importance to place a nuisance changepoint.

In fact, 65% of the samples exhibit a changepoint between 1100 and 1400. Our changepoint model, in turn, only places a changepoint for 36% of the samples there. Conversely, our model places a changepoint between 2900 and 3900 in 98% of the samples, whereas the model of Fearnhead (2006) in only 38%.

As we have already acknowledged, the marginal changepoint probabilities (Figure 19 (a) and 11 (b)) yield a significant underestimation of sensitivity. Hence, particularly in such difficult and rather experimental application scenarios, the user should refrain from using marginal changepoint probabilities alone to verify changepoint models.

Unfortunately, the importances expressed by credible regions tend to overestimate the sensitivity here. This is because the nuisance changepoints spread quite independently over larger sets of timepoints without explaining essential changes in segment height. As a result, a considerable fraction of the samples contain a mix of the desired and also haphazard changepoints. This is challenging for our credible regions, since their computation is based on a sophisticated rejection of samples.

Even though overestimated, the importances still point to areas of substantial uncertainty, which indicates that both models are not perfectly adapted to the data and application. To fix this, the model of Fearnhead (2006) might just need an adjustment of the parameters in order to reduce the sensitivity towards the nuisance changepoints at around 1300. However, since Fearnhead (2006) used marginal changepoint probabilities in order to justify the parameter choices, this small imperfection remained unnoticed.

6.8 Discussion

In this section we develop a novel set estimator in the context of Bayesian changepoint analysis. It enables a new visualization technique that provides very detailed insights into the distribution of changepoints. The resulting plots can be analyzed manually just by considering the concepts of broadness to assess the uncertainty of the changepoint locations in regards to a feature, shape to explore the changepoint locations the model is in favor of to explain a feature, and importance to get an idea about the sensitivity of the model towards a feature. This greatly facilitates the evaluation and the adjustment of changepoint models on the basis of a given dataset. However, by means of these three concepts, we are also able to conveniently analyze changepoint datasets on the basis of a predetermined changepoint model.

We say that a credible region points to a feature in the data if the region contains at least one of the changepoint locations that are triggered by the model as a consequence of this feature. An α level credible region will point to all features having a sensitivity larger than α .

Against this backdrop, we may derive a single α level credible region from a given model in order to analyze the features in a dataset. Thus, we need to choose α in such a way that the corresponding credible region points to the features of interest, but skips the uninteresting ones. A good changepoint model will represent the degree of interest through its sensitivity. On that basis, we can choose α such that the requirements regarding accuracy and parsimony are met. Frick et al. (2014) recommend to choose $\alpha = 0.4$ for their method. This is a reasonable choice. However, it becomes apparent in Section 6.5.2 that the difference between sensitivity and importance can be around 0.3 in difficult cases. Thus, we recommend to choose $\alpha = 0.7$ instead.

An essential quality of the changepoint model is its willingness to jump since it highly affects the model's sensitivity towards features. In the supplement you can find several video files illustrating this through the success probability for the distribution of the sojourn time between successive changepoints. In this context, if changepoint samples

are available but the necessary modifications of the model are not feasible anymore, a sloppy but perhaps effective way to obtain a representative credible region is to increase or decrease α until a reasonable coarseness comes about. However, systematic defects as in the model depicted in Figure 18(b) cannot be tackled in this way.

Credible regions provide valuable knowledge about groups of changepoint locations that explain single features jointly. Besides, they reveal a little bit about combinations of changepoints with respect to more than one feature. The shapes of the credible regions shown in Figure 16(b) and 18(c) suggest that the data can be explained very well through combinations of nine respectively five changepoints, which are limited to very few locations. On the other hand, the medium importance of the feature around 170 in Figure 16(b) shows that there is a notable alternate representation, which could be quite different to that of the true changepoints.

Our theory so far is built on random sets that represent the posterior random changepoints. However, since we can specify a bijective function between subsets of $\{1, \dots, n\}$ and binary sequences in $\{0, 1\}^n$, all the theory in this section applies equally to sequences of binary random variables of equal length. Thus, we may apply credible regions to “Spike and Slab” regression (Mitchell and Beauchamp, 1988; George and McCulloch, 1993; Ishwaran and Rao, 2005) as well. There, each covariate is assigned to a binary variable that determines if the covariate is relevant for explaining the responses. However, covariates can generally not be described in terms of a time series and this could hinder a graphical analysis through credible regions.

Even though the construction of credible regions evoke acceptance regions from statistical testing, we do not intent to create a method in this direction. As we can see in the example in Section 6.5.2, credible regions can be fairly broad for say $\alpha = 0.05$. While this would interfere with statistical testing purposes, it provides valuable insights into changepoint models.

The ILP can be improved by introducing a constraint for each single changepoint in the samples, instead of having constraints for larger sets of samples that share the same changepoint location (Constraint II). This is because for an ILP solver, many small

constraints are easier to handle than a big one involving many variables. However, due to high runtimes we do not recommend using the ILP. While we are able to compute solutions to the Gaussian change in mean example, it is not possible to compute all 29 solutions to the Dow Jones example within a week.

To conclude, the authors of this work highly advice the use of **Greedy**'s credible regions to evaluate and justify changepoint models. To the same extent, we recommend to use them for the analysis of changepoint data. To this end, the R Package *SimCredRegR* provides a fast implementation of **Greedy** and plotting routines that can be applied to changepoint samples without further ado.

7 Final Remarks

In this thesis, we elaborated intensively upon three very important topics that crop up in changepoint research: sampling, inference and uncertainty. While changepoint samples derived from a bare likelihood open up one way to conduct changepoint inference, the computation of certain quantities derived from a changepoint model unbar another. The user of a particular changepoint model has to decide which of the two ways or if even a mix of both is used. This thesis provides the right tools to make and implement these decisions.

Numerous existing and novel inference approaches that go deep into Bayesian changepoint models and their characteristics are discussed. Therewith, the tremendous complexity that arises from random segmentations applied to observed data are demonstrated. This complexity is mainly owed by the simultaneous discrete and continuous nature of changepoint problems.

The resulting mixed state spaces also pose the source of big confusions within the statistics community. Hence, a sense of measure theory is required to gain a comprehensive understanding here. However, with such a sophisticated notational framework at hand, until then seemingly difficult issues become readily manageable and straightforward again.

Nevertheless, runtime challenges like polynomial growth and NP-completeness limits the practical feasibility. Consequently, heuristic approaches like pruning and greedy are of utter importance. Their skillful use allows for the computation of virtually accurate solutions to otherwise unfeasible problems.

Throughout my time as a PhD student, I worked intensively on these problems. To this end, I utilized state of the art computing technologies to develop and run R and C++ programs, and I applied sophisticated mathematical tools like measure theory. With this thesis, I intend to convey my main findings and to help others in solving their own problems.

I finally reckon that spike and slab regression (Mitchell and Beauchamp, 1988; Ishwaran and Rao, 2005) is a promising related topic that may benefit considerably from my research.

8 Related own papers and significant contributions of other researchers to this thesis

The content of Section 3 is part of a paper that was accepted in *The Mathematical Gazette* and will be available in the issue July 2020 (see Siems (2019) for a preprint). Section 4 accounts for a paper that is currently available as a preprint (Siems and Koepfel, 2018). Section 6 reflects the content of Siems et al. (2019).

Professor Marc Hellmuth laid the foundations for Theorem 5 and came up with the corresponding proof idea. He also pointed the way for the ILP of Section 6.3.1. Lisa Köpkel indicated that Metropolis-within-Gibbs might be an expedient method to approach post-hoc translations (see Section 4.3.2).

A Supplementary material

The Sections 4 and 5 are each accompanied by a C++/Qt project providing the relevant source code that was used to compute the presented results. Both projects contain a README file, which explain how to compile, run and browse through the programs on Ubuntu.

The R Package *SimCredRegR* that comes with the supplementary material of Section 6, provides all the applied datasets and sampling algorithms. It further enables the computation of credible regions according to **Greedy**, joined highest density regions, marginal changepoint probabilities, Bonferroni sets and **stepR**'s joined confidence intervals. Besides, you can find the R Package *SimCredRegILPR* to compute credible regions according to the ILP. This package can only be installed if IBM's ILP solver CPLEX (IBM, 2016) is in place. Several video (".mp4") files, which demonstrate how credible regions evolve at different parameter choices, can be found. By means of different realizations of the data in Figure 16, we also provide a comparison of our credible regions and **stepR**'s confidence intervals. This can be found in the file "collection_of_different_data_simulations.pdf".

B Appendix

B.1 How to draw a fixed number of changepoints uniformly

Within a set of timepoints $1, \dots, n$, we want to uniformly place exactly k changepoints, where $0 < k \leq n$. The following lemma shows how to achieve this in an iterative fashion by deciding for every timepoint separately if a changepoint is placed or not.

Lemma 30. *By iterating through $i = 1, \dots, n$ and applying the following rule, we obtain k changepoints in an uniform manner. At timepoint i , with $\ell < k$ previous changepoints, place a changepoint with probability $\frac{k-\ell}{n-i+1}$ or stop if $\ell = k$.*

Proof. For $0 \leq i \leq n$, let $\mathbf{K}_i \in \{0, \dots, i\}$ be a random number with

$$\mathbb{P}(\mathbf{K}_i = \ell \mid \mathbf{K}_{i-1} = \ell') = \begin{cases} \frac{1}{2} & \ell = \ell' \\ \frac{1}{2} & \ell = \ell' + 1 \end{cases}$$

for $0 < i \leq n$. If we look at the random sets $\{i \mid \mathbf{K}_i \neq \mathbf{K}_{i-1}\}$, we see that any two such sets with the same number of elements exhibit the same probability. Thus, by conditioning all \mathbf{K}_i on $\mathbf{K}_n = k$, we obtain k uniformly distributed timepoints.

For $0 \leq k - \ell \leq n - i$ and $0 < i \leq n$, we get $\mathbb{P}(\mathbf{K}_i = \ell, \mathbf{K}_n = k \mid \mathbf{K}_{i-1} = \ell') \propto \binom{n-i}{k-\ell}$ since there are $\binom{n-i}{k-\ell}$ options to pick the $k - \ell$ remaining changepoints from $\{i + 1, \dots, n\}$. Therefore, $\mathbb{P}(\mathbf{K}_i = \ell \mid \mathbf{K}_n = k, \mathbf{K}_{i-1} = \ell) = \frac{\binom{n-i}{k-\ell}}{\binom{n-i}{k-\ell} + \binom{n-i}{k-\ell-1}} = 1 - \frac{k-\ell}{n-i+1}$. \square

B.2 Hypergraphs and the complexity of the SBP

In this section we introduce the k minimum edge union problem (k -minEU), which is shown to be equivalent to the SBP. We show the NP-completeness of the k -minEU. In order to establish the theory, we introduce hypergraphs and some of its basic notions. Indeed, there is a close connection between hypergraphs and the statistical theory established in this research. As we shall see, a hypergraph can be defined in terms of a family of subsets of a finite set and vice versa. We further provide an Integer Linear Program (ILP) formulation to k -minEU that allows to compute exact solutions and we proof its correctness.

B.2.1 Hypergraphs

Here, we briefly discuss (multi-)hypergraphs and their structure and refer to Berge (1984); Voloshin (2009) for the interested reader. Before we start with the formal definitions, we recall that multisets are a natural generalization of usual sets (Calude et al., 2001). Whereas a usual set contains each element only once, a multiset can contain each element arbitrary often. Therefore, a multiset over a set A is defined in terms of a mapping from A to \mathbb{N} that assigns to each $a \in A$ the number of occurrences of a in the multiset.

Definition 5. A hypergraph \mathcal{H} is a pair (V, h) where V is a finite nonempty set and h is a multiset over 2^V , i.e. $h : 2^V \rightarrow \mathbb{N}$.

Definition 6. We say that $e \in 2^V$ is an edge of the hypergraph $\mathcal{H} = (V, h)$ iff $h(e) > 0$ and write $e \in \mathcal{H}$. Moreover, the cardinality $\#\mathcal{H}$ of \mathcal{H} is given as the number of edges $\sum_{e \in \mathcal{H}} h(e)$. The edge union of a hypergraph \mathcal{H} is defined as the set $\mathsf{U}(\mathcal{H}) = \bigcup_{e \in \mathcal{H}} e$. Removing an edge e of \mathcal{H} is achieved by setting $h(e)$ to $\max\{0, h(e) - 1\}$, whereas completely removing an edge e from \mathcal{H} means to set $h(e) = 0$.

Therefore, a set of changepoint samples $s_1, \dots, s_m \subseteq \{1, \dots, n\}$ represents a practical example of a hypergraph. There, the hypergraph is constructed through $V = \{1, \dots, n\}$ and $h(e) = \#\{i \mid e = s_i\}$. Since we can also build a hypergraph from a set of subsets of $\{1, \dots, n\}$, hypergraphs and families of subsets of a finite set are equal.

Definition 7. A hypergraph $\mathcal{G} = (V, g)$ is a sub-hypergraph of a hypergraph $\mathcal{H} = (V, h)$ iff $g(e) \leq h(e)$ for all $e \in 2^V$ and we write $\mathcal{G} \subseteq \mathcal{H}$. Moreover, for a given vertex $x \in V$ and a hypergraph $\mathcal{H} = (V, h)$ we define the sub-hypergraph $\mathcal{D}(\mathcal{H}, x) = (V, g)$ through

$$g(e) = \begin{cases} h(e) & , \text{if } x \in e \\ 0 & , \text{otherwise} \end{cases}$$

$\mathcal{D}(\mathcal{H}, x)$ represents the sub-hypergraph of \mathcal{H} that consists of all edges that contain vertex x .

B.2.2 Computational complexity of the SBP

For a given hypergraph \mathcal{H} and an integer k with $0 \leq k \leq \#\mathcal{H}$, we want to solve the problem of finding a sub-hypergraph $\mathcal{G} \subseteq \mathcal{H}$ that has at least k edges but an edge union of minimum cardinality. Equivalently, we want to find at least one element of the set $\operatorname{argmin}_{\mathcal{G} \subseteq \mathcal{H}} \left\{ \#\mathsf{U}(\mathcal{G}) \mid \#\mathcal{G} \geq k \right\}$. We refer to this task as the k *minimum edge union (optimization) problem* (k -minEU).

Remark 4. *Provided that a solution $\mathcal{G} = (V, g)$ to a k -minEU instance (\mathcal{H}, k) with $\mathcal{H} = (V, h)$ is known, one can easily determine further solutions $\mathcal{G}' = (V, g')$ by choosing for all edges $e \in \mathcal{H}$ an arbitrary integer ℓ_e with $g(e) \leq \ell_e \leq h(e)$ and setting*

$$g'(e) = \begin{cases} \ell_e & , \text{ if } g(e) > 0 \\ 0 & , \text{ otherwise} \end{cases}$$

Clearly, $\#U(\mathcal{G}) = \#U(\mathcal{G}')$ and $\#\mathcal{G}' \geq \#\mathcal{G} \geq k$.

Theorem 6. *The problems k -minEU and SBP are equivalent.*

Before proving this theorem, we need to consider the following lemma.

Lemma 31. *Given a hypergraph $\mathcal{H} = (\{1, \dots, n\}, h)$ and a sequence $s_1, \dots, s_m \subseteq \{1, \dots, n\}$ with $m = \#\mathcal{H}$ and $h(e) = \#\{i \mid e = s_i\}$ for all $e \subseteq \{1, \dots, n\}$, the following applies*

$$\min_{A \subseteq \{1, \dots, n\}} \left\{ \#A \mid \sum_{i=1}^m \mathbf{1} \{s_i \subseteq A\} \geq k \right\} = \min_{\mathcal{G} \subseteq \mathcal{H}} \left\{ \#U(\mathcal{G}) \mid \#\mathcal{G} \geq k \right\} \quad (41)$$

Proof of Lemma 31. For $A \subseteq \{1, \dots, n\}$ with $\sum_{i=1}^m \mathbf{1} \{s_i \subseteq A\} \geq k$, we construct the hypergraph $\mathcal{G} = (\{1, \dots, n\}, g)$ with $g(e) = \#\{i \mid e = s_i, s_i \subseteq A\}$. Since $\#\mathcal{G} \geq k$, $\mathcal{G} \subseteq \mathcal{H}$ and $\#U(\mathcal{G}) \leq \#A$ we conclude that \geq holds in Equation (41).

Given a hypergraph $\mathcal{G} \subseteq \mathcal{H}$ with $\#\mathcal{G} \geq k$, we can choose an $I \subseteq \{1, \dots, m\}$ with $\#I \geq k$ and $\bigcup_{i \in I} s_i \subseteq U(\mathcal{G})$. Since $\sum_{i \in I} \mathbf{1} \{s_i \subseteq U(\mathcal{G})\} \geq k$ we can also conclude that \leq holds in Equation (41). \square

Proof of Theorem 6. Let (\mathcal{H}, k) be a k -minEU instance with $\mathcal{H} = (\{1, \dots, n\}, h)$. We construct an equivalent SBP instance with $s_1, \dots, s_m \subseteq \{1, \dots, n\}$ and $\alpha \in [0, 1]$ such that a solution to this SBP instance provides a solution to the k -minEU instance. Therefore, choose $m = \#\mathcal{H}$, $\alpha = 1 - \frac{k}{m}$ and s_1, \dots, s_m so that $h(e) = \#\{i \mid e = s_i\}$. If A is a solution to this SBP instance, then the hypergraph $\mathcal{G} = (\{1, \dots, n\}, g)$ with $g(e) = \#\{i \mid e = s_i, s_i \subseteq A\}$ is a solution to the k -minEU instance (\mathcal{H}, k) . Indeed, since $\#U(\mathcal{G}) = \#A$ and $\#\mathcal{G} \geq k$, Lemma 31 implies that \mathcal{G} is a solution to the k -minEU instance.

Conversely, given an SBP instance with $\alpha \in [0, 1]$ and $s_1, \dots, s_m \subseteq \{1, \dots, n\}$, we construct an equivalent k -minEU instance (\mathcal{H}, k) such that a solution to this k -minEU instance provides a solution to the SBP instance. Therefore, let $k = \lceil m \cdot (1 - \alpha) \rceil$ and $\mathcal{H} = (\{1, \dots, n\}, h)$ with $h(e) = \#\{i \mid e = s_i\}$. If \mathcal{G} is a solution to this k -minEU instance, then $\mathsf{U}(\mathcal{G})$ is a solution to the SBP instance. Indeed, since $\sum_{i=1}^m \mathbb{1}\{s_i \subseteq \mathsf{U}(\mathcal{G})\} \geq m \cdot (1 - \alpha)$, Lemma 31 implies that $\mathsf{U}(\mathcal{G})$ is a solution to the SBP instance. \square

We now show that (the decision version of) k -minEU (and hence of SBP) is NP-complete (Garey and Johnson, 1979). Thus, there is no polynomial time algorithm to solve this problem, unless $P = NP$.

The decision version of k -minEU is as follows:

Problem (Decision Version of k -minEU).

Input: Hypergraph $\mathcal{H} = (V, h)$ and $k, l \in \mathbb{N}$ with $0 \leq k \leq \#\mathcal{H}$ and $0 < l \leq \#\mathsf{U}(\mathcal{H})$.

Question: Is there a $\mathcal{G} = (V, g) \subseteq \mathcal{H}$ such that $\#\mathcal{G} \geq k$ and $\#\mathsf{U}(\mathcal{G}) \leq l$?

In order to prove the NP-completeness of k -minEU, we use the well-known NP-complete Knapsack-problem (Karp, 1972; Garey and Johnson, 1979).

Problem (Knapsack).

Input: A finite set U and for each $u \in U$ a weight $w(u) \in \mathbb{N}$, a value $v(u) \in \mathbb{N}$ and positive integers a and b .

Query: Is there a subset $U' \subseteq U$ such that $\sum_{u \in U'} w(u) \leq b$ and $\sum_{u \in U'} v(u) \geq a$?

Theorem 7. k -minEU is NP-complete.

Proof. We begin with showing that k -minEU \in NP. To this end, it suffices to demonstrate that a candidate solution to k -minEU can be verified in polynomial time. However, this is easy to see, since we only need to check whether for a possible solution $\mathcal{G} \subseteq \mathcal{H}$ it holds that $\#\mathsf{U}(\mathcal{G}) \leq l$ and $\#\mathcal{G} \geq k$. Both tasks can be done in linear time in the number of edges of \mathcal{G} .

We proceed to show by reduction from Knapsack that k -minEU is NP-hard. Thus, let us assume we are given an arbitrary instance of Knapsack, that is, a finite set U , for each

$u \in U$ the weight $w(u) \in \mathbb{N}$ and the value $v(u) \in \mathbb{N}$, as well as positive integers b and a . Now, we construct an instance of k -minEU as follows. For each $u \in U$ we set an edge $e_u = \{(u, 1), \dots, (u, v(u))\}$ and $h(e_u) = w(u)$. The vertex set of the hypergraph $\mathcal{H} = (V, h)$ is then $V = \cup_{u \in U} e_u$. Note, the edges in \mathcal{H} are pairwise disjoint. Clearly, this reduction can be done in polynomial time in the number of elements in U and the values $v(u)$.

In what follows, we show that **Knapsack** has a solution for given integers b, a if and only if k -minEU has a solution with $k = \#\mathcal{H} - b$ and $l = \#\mathbf{U}(\mathcal{H}) - a$.

Let $U' = \{u_1, \dots, u_n\} \subseteq U$ such that $\sum_{i=1}^n w(u_i) \leq b$ and $\sum_{i=1}^n v(u_i) \geq a$. Completely remove all corresponding edges e_{u_i} , $1 \leq i \leq n$ from \mathcal{H} to obtain the sub-hypergraph \mathcal{G} . Hence, $\#\mathcal{G} = \#\mathcal{H} - \sum_{i=1}^n w(u_i) \geq \#\mathcal{H} - b = k$ and $\#\mathbf{U}(\mathcal{G}) = \#\mathbf{U}(\mathcal{H}) - \sum_{i=1}^n v(u_i) \leq \#\mathbf{U}(\mathcal{H}) - a = l$.

Conversely, assume that $\mathcal{G} = (V, g)$ is a valid solution for the hypergraph $\mathcal{H} = (V, h)$ (as constructed above) and given integers $k \geq \#\mathcal{H}$ and $l \leq \#\mathbf{U}(\mathcal{H})$. Thus, we can write $k = \#\mathcal{H} - b$ and $l = \#\mathbf{U}(\mathcal{H}) - a$. Hence, $\#\mathcal{G} \geq \#\mathcal{H} - b$ and $\#\mathbf{U}(\mathcal{G}) \leq \#\mathbf{U}(\mathcal{H}) - a$. Therefore, at least b edges must have been removed from \mathcal{H} resulting in \mathcal{G} where the edge union of \mathcal{G} has at least a fewer vertices than $\mathbf{U}(\mathcal{H})$. Note, to obtain fewer vertices in $\mathbf{U}(\mathcal{H})$ one needs to completely remove edges from \mathcal{H} . Let $E_0 = \{e \in \mathcal{H} \mid g(e) = 0\}$ be the set of all edges that have been completely removed from \mathcal{H} . By construction, each edge e_u is uniquely identified with an element $u \in U$. We show that $U' = \{u \in U \mid e_u \in E_0\}$ provides a solution for SBP. To this end, observe that $\#\mathcal{H} - \sum_{e \in E_0} h(e) \geq \#\mathcal{G} \geq \#\mathcal{H} - b$, which implies that $\sum_{e \in E_0} h(e) = \sum_{u \in U'} w(u) \leq b$, as desired. Moreover, by construction we have $\#\mathbf{U}(\mathcal{G}) = \#\mathbf{U}(\mathcal{H}) - \sum_{e \in E_0} \#e = \#\mathbf{U}(\mathcal{H}) - \sum_{u \in U'} v(u) \leq \#\mathbf{U}(\mathcal{H}) - a$. Thus, $\sum_{u \in U'} v(u) \geq a$, which completes the proof. \square

Combining Theorem 6 and 7 we obtain the following

Corollary 8. *The decision version of SBP is an NP-complete problem.*

B.3 Proof of correctness of the ILP

We introduce for a hypergraph $\mathcal{H} = (V, h)$ and an integer k the following binary variables $U_x, F_e \in \{0, 1\}$:

$U_x = 1$ if and only if vertex x of \mathcal{H} is contained in the edge union $U(\mathcal{G})$ of $\mathcal{G} \subseteq \mathcal{H}$.

$F_e = 1$ if and only if the edge $e \in \mathcal{H}$ is contained in \mathcal{G} .

Moreover, the number of edges that contain a vertex x is given by the constants $D_x = \#\mathcal{D}(\mathcal{H}, x)$ for all $x \in V$.

To find a solution for the k -minEU problem, we need to minimize the number of vertices in the edge union of $\mathcal{G} = (V, g) \subseteq \mathcal{H}$, which is achieved by minimizing the objective function

$$\sum_{x \in U(\mathcal{H})} U_x \quad (42)$$

By Remark 4 it is always possible to find a sub-hypergraph \mathcal{G} with minimum edge union such that $g(e) = h(e)$ for all $e \in \mathcal{G}$. We will construct such a sub-hypergraph.

To ensure that $\#\mathcal{G} \geq k$ we add the constraint

$$\sum_{e \in \mathcal{H}} h(e) \cdot F_e \geq k \quad (43)$$

Note, e is contained in \mathcal{G} if and only if $F_e = 1$ and by construction, $\#\mathcal{G} = \sum_{e \in \mathcal{H}} (h(e) \cdot F_e)$. Thus, constraint (43) is satisfied if and only if $\#\mathcal{G} \geq k$.

Finally, we have to ensure that $U_x = 1$ if and only if there is an edge in \mathcal{G} that contains x . To this end, we add for all $x \in U(\mathcal{H})$ the constraint

$$\sum_{e \in \mathcal{D}(\mathcal{H}, x)} h(e) \cdot (1 - F_e) \geq D_x \cdot (1 - U_x) \quad (44)$$

Now, if there is no edge containing x in \mathcal{G} and thus, $F_e = 0$ for all $e \in \mathcal{D}(\mathcal{H}, x)$, then Constraint (44) implies that $\sum_{e \in \mathcal{D}(\mathcal{H}, x)} h(e) \geq D_x(1 - U_x)$. Since $\sum_{e \in \mathcal{D}(\mathcal{H}, x)} h(e) = D_x$, we have two choices for $U_x \in \{0, 1\}$. However, the optimization function ensures that U_x is

set to 0. Conversely, assume that there is an edge e that contains x and hence, $F_e = 1$. Thus, $\sum_{e \in \mathcal{D}(\mathcal{H}, x)} h(e) \cdot (1 - F_e) < D_x$. The only way to satisfy Constraint (44) is achieved by setting $U_x = 1$.

Taken together the latter arguments we can infer the following result.

Theorem 8. *The ILP formulation in Eqs. (42) - (44) correctly solves the k -minEU problem.*

List of Symbols

Sets

\mathbb{N}	Natural numbers including 0
\mathbb{N}_+	Natural numbers excluding 0
\mathbb{R}	Real numbers
$\mathbb{R}_{\geq 0}$	Non-negative real numbers
$\mathbb{R}_{> 0}$	Positive real numbers
$\#$	Cardinality

Measures

\mathbb{P}	The all-embracing probability measure
\otimes	Product operator for either two measures or a measure and a Markov kernel
$/$	π/μ is the density of π w.r.t. μ
δ_x	The Dirac or point measure around x
$\#$	Counting measure

Random Variables

$\stackrel{d}{\sim}$	$\mathbf{X} \stackrel{d}{\sim} \mathbf{Y}$ means that the random variables \mathbf{X} and \mathbf{Y} follow the same distribution
\sim	$\mathbf{X} \sim \pi$ means that π is the distribution of \mathbf{X}
$\perp\!\!\!\perp$	$\mathbf{X} \perp\!\!\!\perp \mathbf{Y}$ means that the random variables \mathbf{X} and \mathbf{Y} are independent
$ $	Opens a list of conditions
$;$	Opens a list of (non-random) parameters
\mathbb{E}	Expectation

Functions

\exp	Exponential function in standard form
\ln	Natural logarithm
$\mathbb{1}$	Indicator function
$\phi(x; \mu, \sigma^2)$	Density of normal distribution with expectation μ , variance σ^2 evaluated at x

Others

- \mathcal{O} Big O notation
- Universal placeholder that represents all possible arguments

Distributions

- $\text{NB}(\ell; q, r)$ Probability of ℓ failures under the negative binomial distribution with success probability q and number of successes r
- $\mathcal{N}(\mu, \sigma^2)$ Normal distribution with expectation μ and variance σ^2

List of Figures

1	Excerpt from the counting world	1
2	Ion channel data	2
3	Well-log data	3
4	Dow Jones data	4
5	Coal mine data	4
6	Ad-Hoc and post-hoc transition graphs	28
7	Gaussian change in mean example dataset	37
8	Graphical representation of the changepoint model	54
9	Graph of energy and density of generalized Laplace distribution	67
10	Well-log data, run-lengths and number of particles after pruning	68
11	Well-log data with posterior expectations, variances, and changepoint histogram	78
12	Loglikelihood with MAP and Q-Q plot	89
13	EM estimate for q with MAP and loglikelihood w.r.t. q	92
14	Average EM estimates for variances	101
15	Gaussian data and credible regions	111
16	Gaussian data and credible regions from ILP, HDR, Bonferroni sets, stepR	121
17	SBP convergence and accuracy of Greedy	123
18	Dow Jones data and credible regions	124

19	Well-log data and credible regions	126
----	--	-----

List of Tables

1	Acceptance rates for particular moves	42
2	EM estimates for segment height	95
3	EM for data distributions	98
4	Several statistics derived from EM estimates	102

List of Pseudocodes

1	Forward inference	58
2	Forward inference with pruning	61
3	Sampling	73
4	Greedy Algorithm for SBP	116

List of Acronyms

MAP	Maximum A-Posteriori
EM	Expectation Maximization
BIC	Bayesian Information Criterion
AIC	Akaike Information Criterion
MCMC	Markov Chain Monte Carlo
SDT	Semi Deterministic Translation
SBP	Sampling Based Problem
ILP	Integer Linear Programming
HDR	Highest Density Region

List of References

- Adams, R. P. and D. J. MacKay (2007). Bayesian online changepoint detection. ArXiv e-prints – <https://arxiv.org/abs/0710.3742>.
- Akaike, H. (1973). Information Theory and an Extension of the Maximum Likelihood Principle. In *Selected Papers of Hirotugu Akaike*. Springer New York.
- Andrieu, C., P. Djurić, and A. Doucet (2001). Model selection by MCMC computation. *Signal Processing* 81(1), 19–37.
- Aston, J. A., J.-Y. Peng, and D. E. Martin (2012). Implied distributions in multiple change point problems. *Statistics and Computing* 22(4), 981–993.
- Bansal, N. K., H. Du, and G. Hamedani (2008). An application of EM algorithm to a change-point problem. *Communications in Statistics—Theory and Methods* 37(13), 2010–2021.
- Bayes, T. (1763). An essay towards solving a problem in the doctrine of chances. *Philosophical transactions of the Royal Society of London* (53), 370–418.
- Bedard, M. (2008). Optimal acceptance rates for Metropolis algorithms: Moving beyond 0.234. *Stochastic Processes and their Applications* 118(12), 2198–2222.
- Bellman, R. (1966). Dynamic programming. *Science* 153(3731), 34–37.
- Berge, C. (1984). *Hypergraphs: Combinatorics of Finite Sets*. North-Holland Mathematical Library. Elsevier Science.
- Besag, J. (2001). Markov chain Monte Carlo for statistical inference. *Center for Statistics and the Social Sciences* 9, 24–25.
- Bishop, C. M. and T. M. Mitchell (2014). *Pattern Recognition and Machine Learning*. Springer.
- Boutilier, C., N. Friedman, M. Goldszmidt, and D. Koller (1996). Context-specific independence in Bayesian networks. In *Proceedings of the Twelfth international conference on Uncertainty in artificial intelligence*, pp. 115–123. Morgan Kaufmann Publishers Inc.
- Box, G. E. P. and G. C. Tiao (1973). *Bayesian inference in statistical analysis*. Addison-Wesley Series in Behavioral Science: Quantitative Methods. Addison-Wesley Pub. Co.

- Calude, C. S., G. Paun, G. Rozenberg, and A. Salomaa (2001). *Multiset processing: mathematical, computer science, and molecular computing points of view*, Volume 2235. Springer Science & Business Media.
- Camerino, D. C., D. Tricarico, and J.-F. Desaphy (2007). Ion channel pharmacology. *Neurotherapeutics* 4(2), 184–198.
- Carlin, B. P. and S. Chib (1995). Bayesian model choice via Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society: Series B (Methodological)* 57(3), 473–484.
- Casella, G. (1985). An introduction to empirical Bayes data analysis. *The American Statistician* 39(2), 83–87.
- Cauchy, A. (1847). Méthode générale pour la résolution des systemes d'équations simultanées. *Comptes rendus de l'Académie des Sciences* 25(1847), 536–538.
- Chen, M.-H., Q.-M. Shao, and J. G. Ibrahim (2012). *Monte Carlo methods in Bayesian computation*. Springer Science & Business Media.
- Dempster, A. P., N. M. Laird, and D. B. Rubin (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society: Series B* 39(1), 1–38.
- Doob, J. L. (1948). Renewal theory from the point of view of the theory of probability. *Transactions of the American Mathematical Society* 63(3), 422–438.
- Doucet, A., N. D. Freitas, and N. Gordon (2001). *An introduction to sequential Monte Carlo methods*. Springer.
- Dunnnett, C. W. (1955). A Multiple Comparison Procedure for Comparing Several Treatments with a Control. *Journal of the American Statistical Association* 50(272), 1096–1121.
- Eckley, I. A., P. Fearnhead, and R. Killick (2011). Analysis of changepoint models. *Bayesian Time Series Models*, 205–224.
- Fearnhead, P. (2006). Exact and efficient Bayesian inference for multiple changepoint problems. *Statistics and Computing* 16(2), 203–213.
- Fearnhead, P. and P. Clifford (2003). On-line inference for hidden Markov models via particle filters. *Journal of the Royal Statistical Society: Series B* 65(4), 887–899.

- Fearnhead, P. and Z. Liu (2007). On-line inference for multiple changepoint problems. *Journal of the Royal Statistical Society: Series B* 69(4), 589–605.
- Fearnhead, P. and Z. Liu (2009). Efficient Bayesian analysis of multiple changepoint models with dependence across segments. *Statistics and Computing* 21(2), 217–229.
- Fearnhead, P. and G. Rigaiil (2019). Changepoint detection in the presence of outliers. *Journal of the American Statistical Association* 114(525), 169–183.
- Frick, K., A. Munk, and H. Sieling (2014). Multiscale change point inference. *Journal of the Royal Statistical Society: Series B* 76(3), 495–580.
- Friedrich, F., A. Kempe, V. Liescher, and G. Winkler (2008). Complexity Penalized M-Estimation: Fast Computation. *Journal of Computational and Graphical Statistics* 17(1), 201–224.
- Frobenius, G. (1912). *Über Matrizen aus nicht negativen Elementen*. Reimer.
- Fronk, E.-M. and P. Giudici (2004). Markov Chain Monte Carlo model selection for DAG models. *Statistical Methods and Applications* 13(3), 259–273.
- Fryzlewicz, P. et al. (2014). Wild binary segmentation for multiple change-point detection. *The Annals of Statistics* 42(6), 2243–2281.
- Garey, M. R. and D. S. Johnson (1979). *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co.
- George, E. I. and R. E. McCulloch (1993). Variable Selection via Gibbs Sampling. *Journal of the American Statistical Association* 88(423), 881–889.
- Georgii, H.-O. (2011). *Gibbs measures and phase transitions*, Volume 9. Walter de Gruyter.
- Geyer, C. J. and J. Møller (1994). Simulation procedures and likelihood inference for spatial point processes. *Scandinavian journal of statistics*, 359–373.
- Godsill, S. J. (2001). On the relationship between Markov chain Monte Carlo methods for model uncertainty. *Journal of computational and graphical statistics* 10(2), 230–248.
- Green, P. and A. O’Hagan (1998). Model choice with MCMC on product spaces without using pseudo-priors. Technical report, Department of Mathematics, University of Nottingham.

- Green, P. J. (1995). Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika* 82(4), 711–732.
- Green, P. J. (2003). Trans-dimensional markov chain monte carlo. *Oxford Statistical Science Series*, 179–198.
- Green, P. J. and D. I. Hastie (2009). Reversible jump MCMC. *Genetics* 155(3), 1391–1403.
- Guédon, Y. (2015). Segmentation uncertainty in multiple change-point models. *Statistics and Computing* 25(2), 303–320.
- Guyon, I. and A. Elisseeff (2003). An introduction to variable and feature selection. *Journal of machine learning research* 3(Mar), 1157–1182.
- Hannart, A. and P. Naveau (2009). Bayesian multiple change points and segmentation: Application to homogenization of climatic series. *Water resources research* 45(10).
- Hastings, W. K. (1970). Monte Carlo Sampling Methods Using Markov Chains and Their Applications. *Biometrika*.
- Haynes, K., I. A. Eckley, and P. Fearnhead (2017). Computationally efficient changepoint detection for a range of penalties. *Journal of Computational and Graphical Statistics* 26(1), 134–143.
- Held, L. (2004). Simultaneous Posterior Probability Statements from Monte Carlo Output. *Journal of Computational and Graphical Statistics* 13(1), 20–35.
- Hobolth, A. and E. A. Stone (2009). Simulation from endpoint-conditioned, continuous-time Markov chains on a finite state space, with applications to molecular evolution. *The Annals of Applied Statistics* 3(3), 1204.
- Hoffman, M. D. and A. Gelman (2014). The No-U-turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research* 15(1), 1593–1623.
- Hotz, T. and H. Sieling (2016). stepR: Fitting Step-Functions. <http://CRAN.R-project.org/package=stepR>. R package version 1.0-4.
- Hyndman, R. J. (1996). Computing and Graphing Highest Density Regions. *The American Statistician* 50(2), 120–126.
- IBM (2016). IBM ILOG CPLEX C++ Optimizer.

- Ishwaran, H. and J. S. Rao (2005). Spike and slab variable selection: Frequentist and Bayesian strategies. *The Annals of Statistics* 33(2), 730–773.
- Jackson, B., J. D. Scargle, D. Barnes, S. Arabhi, A. Alt, P. Gioumouisis, E. Gwin, P. Sangtrakulcharoen, L. Tan, and T. T. Tsai (2005). An algorithm for optimal partitioning of data on an interval. *IEEE Signal Processing Letters* 12(2), 105–108.
- Jafari, A., M. Yarmohammadi, and A. Rasekhi (2016). Bayesian analysis to detect change-point in two-phase Laplace model. *Scientific Research and Essays* 11(18), 187–193.
- Jannink, J.-L. and R. L. Fernando (2004). On the Metropolis-Hastings acceptance probability to add or drop a quantitative trait locus in Markov chain Monte Carlo-based Bayesian analyses. *Genetics* 166(1), 641–643.
- Jarrett, R. (1979). A note on the intervals between coal-mining disasters. *Biometrika* 66(1), 191–193.
- Kalman, R. E. (1960). A New Approach to Linear Filtering and Prediction Problems. *Journal of Basic Engineering* 82(Series D), 35–45.
- Karp, R. M. (1972). Reducibility among combinatorial problems. In *Complexity of computer computations*, pp. 85–103. Springer.
- Keshavarz, M. and B. Huang (2014). Expectation Maximization method for multivariate change point detection in presence of unknown and changing covariance. *Computers & Chemical Engineering* 69, 128–146.
- Killick, R., P. Fearnhead, and I. A. Eckley (2012). Optimal Detection of Changepoints With a Linear Computational Cost. *Journal of the American Statistical Association* 107(500), 1590–1598.
- Koenig, W. (2005). Stochastische Prozesse I: Markovketten in diskreter und stetiger Zeit. <https://www.wias-berlin.de/people/koenig/www/StPrI.pdf>. [Online; accessed 15-October-2019].
- Kolmogorov, A. N. and A. T. Bharucha-Reid (2018). *Foundations of the theory of probability: Second English Edition*. Courier Dover Publications.
- Kotz, S., T. Kozubowski, and K. Podgorski (2012). *The Laplace distribution and generalizations: a revisit with applications to communications, economics, engineering, and finance*. Springer Science & Business Media.

- Lai, T. L. and H. Xing (2011). A simple Bayesian approach to multiple change-points. *Statistica Sinica* 21, 539–569.
- Lavielle, M. and E. Lebarbier (2001). An application of MCMC methods for the multiple change-points problem. *Signal Processing* 81(1), 39–53.
- Lawrance, A. J. (1973). Dependency of intervals between events in superposition processes. *Journal of the Royal Statistical Society: Series B* 35(2), 306–315.
- Li, S. and R. Lund (2012). Multiple changepoint detection via genetic algorithms. *Journal of Climate* 25(2), 674–686.
- Liu, Y.-Y., S. Li, F. Li, L. Song, and J. M. Rehg (2015). Efficient learning of continuous-time hidden markov models for disease progression. In *Advances in neural information processing systems*, pp. 3600–3608.
- Meindl, B. and M. Templ (2013). Analysis of Commercial and Free and Open Source Solvers for the Cell Suppression Problem. *Transactions on Data Privacy* 6(2), 147–159.
- Metropolis, N., A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller (1953). Equation of state calculations by fast computing machines. *The journal of chemical physics* 21(6), 1087–1092.
- Mitchell, T. J. and J. J. Beauchamp (1988). Bayesian Variable Selection in Linear Regression. *Journal of the American Statistical Association* 83(404), 1023–1032.
- Nam, C. F., J. A. Aston, and A. M. Johansen (2012). Quantifying the uncertainty in change points. *Journal of Time Series Analysis* 33(5), 807–823.
- Nikodym, O. (1930). Sur une généralisation des intégrales de MJ Radon. *Fundamenta Mathematicae* 15(1), 131–179.
- Pearl, J. (1998). *Graphical models for probabilistic and causal reasoning*. Springer.
- Perreault, L., E. Parent, J. Bernier, B. Bobee, and M. Slivitzky (2000). Retrospective multivariate Bayesian change-point analysis: a simultaneous single change in the mean of several hydrological sequences. *Stochastic Environmental Research and Risk Assessment* 14(4), 243–261.
- Peskun, P. H. (1973). Optimum monte-carlo sampling using markov chains. *Biometrika* 60(3), 607–612.
- Rabiner, L. R. and B.-H. Juang (1986). An introduction to hidden Markov models. *IEEE ASSP Magazine* 3(1), 4–16.

- Raiffa, H. and R. Schlaifer (1961). *Applied statistical decision theory*. Division of Research, Graduate School of Business Administration, Harvard.
- Rauch, H. E., C. T. Striebel, and F. Tung (1965). Maximum likelihood estimates of linear dynamic systems. *AIAA Journal* 3, 1445–1450.
- Rigaill, G., E. Lebarbier, and S. Robin (2012). Exact posterior distributions and model selection criteria for multiple change-point detection problems. *Statistics and Computing* 22(4), 917–929.
- Roberts, W. J. and Y. Ephraim (2008). An EM algorithm for ion-channel current estimation. *IEEE Transactions on Signal Processing* 56(1), 26–33.
- Roodaki, A., J. Bect, and G. Fleury (2011). Note on the computation of the Metropolis-Hastings ratio for Birth-or-Death moves in trans-dimensional MCMC algorithms for signal decomposition problems. ArXiv e-prints – <https://arxiv.org/abs/1111.6245>.
- Sakmann, B. and E. Neher (1984). Patch clamp techniques for studying ionic channels in excitable membranes. *Annual Review of Physiology* 46(1), 455–472.
- Sambridge, M., K. Gallagher, A. Jackson, and P. Rickwood (2006). Trans-dimensional inverse problems, model comparison and the evidence. *Geophysical Journal International* 167(2), 528–542.
- Särkkä, S. (2013). *Bayesian filtering and smoothing*, Volume 3. Cambridge University Press.
- Schwarz, G. et al. (1978). Estimating the dimension of a model. *The Annals of Statistics* 6(2), 461–464.
- Siekmann, I., M. Fackrell, E. J. Crampin, and P. Taylor (2016). Modelling modal gating of ion channels with hierarchical Markov models. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 472(2192), 20160122.
- Siekmann, I., J. Sneyd, and E. J. Crampin (2012). MCMC can detect nonidentifiable models. *Biophysical journal* 103(11), 2275–2286.
- Siekmann, I., J. Sneyd, and E. J. Crampin (2014). Statistical analysis of modal gating in ion channels. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 470(2166), 20140030.
- Siekmann, I., L. E. Wagner II, D. Yule, C. Fox, D. Bryant, E. J. Crampin, and J. Sneyd (2011). MCMC estimation of markov models for ion channels. *Biophysical Journal* 100(8), 1919–1929.

- Siems, T. (2019). Markov Chain Monte Carlo on Finite State Spaces. ArXiv e-prints – <https://arxiv.org/abs/1903.09019>.
- Siems, T., M. Hellmuth, and V. Liebscher (2019). Simultaneous credible regions for multiple changepoint locations. *Journal of Computational and Graphical Statistics*, 1–9.
- Siems, T. and L. Koepfel (2018). A note on the Metropolis-Hastings acceptance probabilities for mixture spaces. ArXiv e-prints – <https://arxiv.org/abs/1808.00789>.
- Sisson, S. A. (2005). Transdimensional Markov chains: A decade of progress and future perspectives. *Journal of the American Statistical Association* 100(471), 1077–1089.
- Stewart, W. J. (2009). *Probability, Markov Chains, Queues, and Simulation: The Mathematical Basis of Performance Modeling*. Princeton University Press.
- Tierney, L. (1994). Markov chains for exploring posterior distributions. *The Annals of Statistics*, 1701–1728.
- Tierney, L. (1998). A note on Metropolis-Hastings kernels for general state spaces. *The Annals of Applied Probability* 8(1), 1–9.
- Tourneret, J.-Y., M. Doisy, and M. Lavielle (2003). Bayesian off-line detection of multiple change-points corrupted by multiplicative noise: application to SAR image edge detection. *Signal Processing* 83(9), 1871–1887.
- Turner, R. D., S. Bottone, M. Deisenroth, Y. Saatçi, C. Rasmussen, F. Huszár, D. Roy, P. Orbanz, and J. M. Hernández (2010). *Gaussian Processes for State Space Models and Change Point Detection*. Ph. D. thesis, University of Cambridge.
- Viterbi, A. (1967). Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *IEEE transactions on Information Theory* 13(2), 260–269.
- Voloshin, V. I. (2009). *Introduction to Graph and Hypergraph Theory*. Nova Science Publishers.
- Waagepetersen, R. and D. Sorensen (2001). A Tutorial on Reversible Jump MCMC with a View toward Applications in QTL-mapping. *International Statistical Review* 69(1), 49–61.
- Weinmann, A., M. Storath, and L. Demaret (2015). The L1-Potts Functional for Robust Jump-Sparse Reconstruction. *SIAM Journal on Numerical Analysis* 53(1), 644–673.
- Wikipedia Contributors (2018). Conjugate prior — Wikipedia, The Free Encyclopedia. https://en.wikipedia.org/wiki/Conjugate_prior. [Online; accessed 10-September-2018].

- Yildirim, S., S. S. Singh, and A. Doucet (2013). An online expectation–maximization algorithm for changepoint models. *Journal of Computational and Graphical Statistics* 22(4), 906–926.
- Yu, S. (2010). Hidden semi-Markov models. *Artificial Intelligence* 174(2), 215–243. Special Review Issue.