

Exact Bayesian Regression of Piecewise Constant Functions

Marcus Hutter*

Abstract. ¹We derive an exact and efficient Bayesian regression algorithm for piecewise constant functions of unknown segment number, boundary locations, and levels. The derivation works for any noise and segment level prior, e.g. Cauchy which can handle outliers. We derive simple but good estimates for the in-segment variance. We also propose a Bayesian regression curve as a better way of smoothing data without blurring boundaries. The Bayesian approach also allows straightforward determination of the evidence, break probabilities and error estimates, useful for model selection and significance and robustness studies. We discuss the performance on synthetic and real-world examples. Many possible extensions are discussed.

Keywords: Bayesian regression, exact polynomial algorithm, non-parametric inference, piecewise constant function, dynamic programming, change point problem.

1 Introduction

“In science the primary duty of ideas is to be useful and interesting, even more than to be true.”

— Wilfred Trotter (1941)

We consider the problem of fitting a piecewise constant function through noisy one-dimensional data, as e.g. in Figure 1, where the segment number, boundaries and levels are unknown. Regression with piecewise constant (PC) functions, a special case of change point detection, has many applications, e.g. in seismology, tomography, biology, and econometric modeling. Determining DNA copy numbers in cancer cells from microarray data is a very recent application.

Bayesian piecewise constant regression (BPCR). We provide a full Bayesian analysis of PC-regression. For a fixed number of segments we choose some prior over all possible segment boundary locations. Some prior on the segment levels and data noise within each segment is assumed. Finally a prior over the number of segments is chosen. From this we obtain the posterior segmentation probability distribution (Section 2). In practice we need summaries of this complicated distribution. A simple maximum (MAP) approximation or mean does not work here. The right way is to proceed in stages from determining the most critical segment number, to the boundary location, and finally to the then trivial segment levels. We also extract the evidence, the boundary probability

*RSISE @ ANU and SML @ NICTA, Canberra, Australia, <http://www.hutter1.net>

¹Submitted 26.Aug'05. Revised 20.Sep'06. A shorter version will appear in the proceedings of the Valencia 2006 meeting (Hut07).

distribution, and an interesting non-PC mean regression curve including error estimate (Section 3). We derive an exact polynomial-time dynamic-programming-type algorithm for all quantities of interest (Sections 5, 6 and 9). Our algorithm is applicable for any noise and level prior. We consider more closely the Gaussian “standard” prior and heavy-tailed robust-to-outliers distributions like the Cauchy, and briefly discuss the non-parametric case (Sections 4 and 7). Finally, some hyper-parameters like the global data average and variability and local within-level noise have to be determined. We introduce and discuss efficient semi-principled estimators, thereby avoiding problematic or expensive numerical EM or Monte-Carlo estimates (Section 8). We test our method on some synthetic examples (Section 10) and some real-world data sets (Section 11). The simulations show that our method handles difficult data with high noise and outliers well. Our basic algorithm can (easily) be modified in a variety of ways: For discrete segment levels, segment dependent variance, piecewise linear and non-linear regression, non-parametric noise prior, etc. (Section 12).

Comparison to other work. Sen and Srivastava (SS75) developed a frequentist solution to the problem of detecting a single (the most prominent) segment boundary, called change or break point. Olshen et al. (OVLW04) generalize this method to detect pairs of break points, which improves recognition of short segments. Both methods are then (heuristically) used to recursively determine further change points. Another approach is penalized Maximum Likelihood (ML). For a fixed number of segments, ML chooses the boundary locations that maximize the data likelihood, which is equivalent to minimizing the mean square data deviation in case of Gaussian noise. Jong et al. (Jon03) use a population based algorithm as minimizer, while Picard et al. (Pic05) use dynamic programming, which is structurally very close to our core recursion, to find the exact solution in polynomial time. An additional penalty term has to be added to the likelihood in order to determine the correct number of segments. The most principled penalty is the Bayesian Information Criterion (Sch78; KW95). Since it can be biased towards too simple (Wea99) or too complex (Pic05) models, in practice often a heuristic penalty is used. An interesting heuristic, based on the curvature of the log-likelihood as a function of the number of segments, has been used in (Pic05). Our Bayesian regressor is a natural response to penalized ML. Another related work to ours is Bayesian bin density estimation by Endres and Földiák (EF05), who also average over all boundary locations, but in the context of density estimation. Another related series of papers, developing exact Bayesian algorithms for PC regression is (Yao84; BH92; Fea06), which we discuss in more detail in Section 13. All three papers consider renewal processes (which have a highly informed prior for the number of segments k) and derive recursions over the number of data points n , while we consider a noninformative uniform prior over k (which is not a renewal process) and derive a recursion over k . See also (Fea05) for a similar Bayesian/MAP approach. Many other approximate, heuristic, sampling, or non-Bayesian approaches to PC regression exist; too many to list them all.

Advantages of Bayesian regression. A full Bayesian approach, when computationally feasible, has various advantages over others: A generic advantage is that it is more principled and hence involves fewer heuristic design choices. This is particularly

important for estimating the number of segments. Another generic advantage is that it can be easily embedded in a larger framework. For instance, one can decide among competing models solely based on the Bayesian evidence. Finally, Bayes often works well in practice, and provably so if the model assumptions are valid.² We can also easily extract other information, like probability estimates and variances for the various quantities of interest. Particularly interesting is the expected level (and variance) of each data point. This leads to a regression curve, which is very flat, i.e. smoothes the data, in long and clear segments, wiggles in less clear segments, follows trends, and jumps at the segment boundaries. It thus behaves somewhat between local smoothing and rigid PC-segmentation.

2 The General Model

Setup. We are given a sequence $\mathbf{y}=(y_1,\dots,y_n)$, e.g. times-series data or measurements of some function at locations $1\dots n$, where each $y_i\in\mathbb{R}$ resulted from a noisy “measurement”, i.e. we assume that the y_i are independently (e.g. Gaussian) distributed with means μ'_i and³ variances σ'_i ². The data likelihood is therefore⁴

$$\text{likelihood: } P(\mathbf{y}|\boldsymbol{\mu}',\boldsymbol{\sigma}') := \prod_{i=1}^n P(y_i|\mu'_i,\sigma'_i) \tag{1}$$

The estimation of the true underlying function $f=(f_1,\dots,f_n)$ is called regression. We assume or model f as piecewise constant. Consider k segments with segment boundaries $0=t_0 < t_1 < \dots < t_{k-1} < t_k = n$, i.e. f is constant on $\{t_{q-1}+1,\dots,t_q\}$ for each $0 < q \leq k$. If the noise within each segment is the same, we have

$$\text{piecewise constant: } \mu'_i = \mu_q \text{ and } \sigma'_i = \sigma_q \text{ for } t_{q-1} < i \leq t_q \quad \forall q \tag{2}$$

We first consider the case in which the variances of all segments coincide, i.e. $\sigma_q = \sigma \forall q$. The easier case of different variance is described in Section 12. Our goal is to estimate the segment levels $\boldsymbol{\mu}=(\mu_1,\dots,\mu_k)$, boundaries $\mathbf{t}=(t_0,\dots,t_k)$, and their number k . Bayesian regression proceeds in assuming a prior for these *quantities of interest*. We model the segment levels by a broad (e.g. Gaussian) distribution with mean ν and variance ρ^2 . For the segment boundaries we take some (e.g. uniform) distribution among all segmentations into k segments. Finally we take some prior (e.g. uniform) over the segment number k . So our prior $P(\boldsymbol{\mu},\mathbf{t},k)$ is the product of

$$\text{prior: } P(\mu_q|\nu,\rho)\forall q \text{ and } P(\mathbf{t}|k) \text{ and } P(k) \tag{3}$$

²Note that we are not claiming here that BPCR works better than the other mentioned approaches. In a certain sense Bayes is optimal if the prior is ‘true’. Practical superiority likely depends on the type of application. A comparison for micro-array data is in progress. The major aim of this paper is to derive an efficient algorithm, and demonstrate the gains of BPCR beyond bare PC-regression, e.g. the (predictive) regression curve (which is better than local smoothing which wiggles more and blurs jumps).

³More generally, μ'_i and σ'_i are location and scale parameters of a symmetric distribution.

⁴For notational and verbal simplicity we will not distinguish between probabilities of discrete variables and densities of continuous variables.

We regard the global variance ρ^2 and mean ν of $\boldsymbol{\mu}$ and the in-segment variance σ^2 as fixed hyper-parameters, and notationally suppress them in the following. We will return to their determination in Section 8.

Evidence and posterior. Given the prior and likelihood we can compute the data evidence and posterior $P(\mathbf{y}|\boldsymbol{\mu}, \mathbf{t}, k)$ by Bayes' rule:

$$\text{evidence: } P(\mathbf{y}) = \sum_{k, \mathbf{t}} \int P(\mathbf{y}|\boldsymbol{\mu}, \mathbf{t}, k) P(\boldsymbol{\mu}, \mathbf{t}, k) d\boldsymbol{\mu}$$

$$\text{posterior: } P(\boldsymbol{\mu}, \mathbf{t}, k|\mathbf{y}) = \frac{P(\mathbf{y}|\boldsymbol{\mu}, \mathbf{t}, k) P(\boldsymbol{\mu}, \mathbf{t}, k)}{P(\mathbf{y})}$$

The posterior contains all information of interest, but is a complex object for practical use. So we need summaries like the maximum (MAP) or mean and variances. MAP over continuous parameters ($\boldsymbol{\mu}$) is problematic, since it is not reparametrization invariant. This is particularly dangerous if MAP is across different dimensions (k), since then even a linear transformation ($\boldsymbol{\mu} \rightsquigarrow \alpha \boldsymbol{\mu}$) scales the posterior (density) exponentially in k (by α^k). This severely influences the maximum over k , i.e. the estimated number of segments. The mean of $\boldsymbol{\mu}$ does not have this problem. On the other hand, the mean of \mathbf{t} makes only sense for fixed (e.g. MAP) k . The most natural solution is to proceed in stages similar to as the prior (3) has been formed.

3 Quantities of Interest

We now define estimators for all quantities of interest in stages as suggested in Section 2. Our first quantities are the posterior of the number of segments and the MAP segment number:

$$\# \text{ segments: } P(k|\mathbf{y}) \quad \text{and} \quad \hat{k} = \arg \max_k P(k|\mathbf{y})$$

Second, for each boundary t_q its posterior and MAP, given the MAP estimate of k :

$$\text{boundaries: } P(t_q|\mathbf{y}, \hat{k}) \quad \text{and} \quad \hat{t}_q = \arg \max_{t_q} P(t_q|\mathbf{y}, \hat{k})$$

Different estimates of t_q (e.g. the mean or MAP based on the joint \mathbf{t} posterior) will be discussed elsewhere. Finally we want the segment level means for the MAP segmentation:

$$\text{segment level: } P(\mu_q|\mathbf{y}, \hat{\mathbf{t}}, \hat{k}) \quad \text{and} \quad \hat{\mu}_q = \int P(\mu_q|\mathbf{y}, \hat{\mathbf{t}}, \hat{k}) \mu_q d\mu_q$$

The estimate $(\hat{\boldsymbol{\mu}}, \hat{\mathbf{t}}, \hat{k})$ uniquely defines the piecewise constant (PC) function \hat{f} , which is our estimate of f . A quite different quantity is to Bayes-average over all piecewise constant functions and to ask for the mean at location i as an estimate for f_i :

$$\text{regression curve: } P(\mu'_i|\mathbf{y}) \quad \text{and} \quad \hat{\mu}'_i = \int P(\mu'_i|\mathbf{y}) \mu'_i d\mu'_i$$

We will see that $\boldsymbol{\mu}'$ behaves similar to a local smoothing of \mathbf{y} , but without blurring true jumps. Standard deviations of all estimates may also be reported.

4 Specific Models

We now complete the specification by discussing some (purely exemplary) choices for the data noise and prior.

Gaussian model. The standard assumption on the noise is independent Gauss:

$$\text{Gaussian noise: } P(y_i|\mu'_i, \sigma'_i) = \frac{1}{\sqrt{2\pi}\sigma'_i} e^{-\frac{(y_i-\mu'_i)^2}{2\sigma'^2_i}} \quad (4)$$

The corresponding standard conjugate prior on the means μ_q for each segment q is also Gauss:

$$\text{Gaussian prior: } P(\mu_q|\nu, \rho) = \frac{1}{\sqrt{2\pi}\rho} e^{-\frac{(\mu_q-\nu)^2}{2\rho^2}} \quad (5)$$

Cauchy model. The standard problem with Gauss is that it does not handle outliers well. If we do not want to or cannot remove outliers by hand, we have to properly model them by a prior with heavier tails. This can be achieved by a mixture of Gaussians or by a Cauchy distribution:

$$\text{Cauchy noise: } P(y_i|\mu'_i, \sigma'_i) = \frac{1}{\pi} \frac{\sigma'_i}{\sigma'^2_i + (y_i - \mu'_i)^2} \quad (6)$$

Note that μ'_i and σ'_i determine the location and scale of Cauchy but are not its mean and variance (which do not exist). The prior on the levels μ_q may as well be modeled as Cauchy:

$$\text{Cauchy prior: } P(\mu_q|\nu, \rho) = \frac{1}{\pi} \frac{\rho}{\rho^2 + (\mu_q - \nu)^2} \quad (7)$$

Actually, the Gaussian noise model may well be combined with a non-Gaussian prior and vice versa if appropriate.

Segment boundaries. The most natural assumption is a uniform (non-informative) prior over all segmentations into k segments. Since there are $\binom{n-1}{k-1}$ ways of placing the $k-1$ inner boundaries (ordered and without repetition) on $(1, \dots, n-1)$, we have:

$$\text{uniform boundary prior: } P(\mathbf{t}|k) = \binom{n-1}{k-1}^{-1} \quad (8)$$

We discuss more general factorizable \mathbf{t} -priors, including general renewal processes, later.

Number of segments. Finally, consider the number of segments k , which is an integer between 1 and n . Sure, if we have prior knowledge on the [minimal,maximal] number of segments $[k_{min}, k_{max}]$ we could/should set $P(k)=0$ outside this interval. Otherwise, any non-extreme choice of $P(k)$ has little influence on the final results, since it gets swamped by the (implicit) strong (exponential) dependence of the likelihood on k . So we suggest a uniform prior

$$P(k) = \frac{1}{k_{max}} \quad \text{for } 1 \leq k \leq k_{max} \quad \text{and } 0 \quad \text{otherwise}$$

with $k_{max}=n$ as default (or $k_{max} < n$ discussed later).

5 The Core Recursion

We now derive expressions for all quantities of interest, which need time $O(k_{max}n^2)$ and space $O(n^2)$. The key fact we exploit is that, fixing a boundary t_p , all probabilities $P(\cdot|t_p, k)$ factorize into a P depending only on quantities left of the boundary times another P depending only on quantities right of the boundary. This leads to a recursion that terminates with an evidence matrix and moments A for single segments. The algorithm naturally breaks into two parts. (i) determination of A from data y for specific noise distribution and segment level prior as described in Section 7, and (ii) the core recursion for specific boundary prior described in this section. In the next section we show how the core recursion can be used to compute all quantities of interest. The dynamic program (ii) depends only on A , i.e. can be applied for any noise and level model, as long as A can be computed.

Notation. Throughout this and the next two sections we use the following notation: k is the total number of segments, t some data index, q some segment index, $0 \leq i < h < j \leq n$ are data item indices of segment boundaries $t_0 \leq t_l < t_p < t_m \leq t_k$, i.e. $t_0 = 0$, $t_l = i$, $t_p = h$, $t_m = j$, $t_k = n$. Further, $y_{ij} = (y_{i+1}, \dots, y_j)$ is data with segment boundaries $t_{lm} = (t_l, \dots, t_m)$ and segment levels $\mu_{lm} = (\mu_{l+1}, \dots, \mu_m)$. In particular $y_{0n} = \mathbf{y}$, $t_{0k} = \mathbf{t}$, and $\mu_{0k} = \boldsymbol{\mu}$. All introduced matrices below (capital symbols with indices) will be important in our algorithm.

General recursion. For $m = l + 1$, y_{ij} is data from a single segment with mean μ_m whose joint distribution (given segment boundaries and $m = l + 1$) is

$$\text{single segment: } P(y_{ij}, \mu_m | t_{m-1, m}) = P(\mu_m) \prod_{t=i+1}^j P(y_t | \mu_m) \quad (9)$$

by the model assumptions (1) and (2). Here and in the following, if the last argument of P is an integer, it denotes the number of segments in the considered data range $i \dots j$. The probabilities for a general but fixed segmentation are independent, i.e. for $m - l$ segments from i to j with breaks at t_{lm} the joint data y_{ij} and segment mean μ_{lm} distribution factorizes as:

$$\begin{aligned} P(y_{ij}, \mu_{lm} | t_{lm}) &= \prod_{p=l+1}^m \left[P(\mu_p) \prod_{t=t_{p-1}+1}^{t_p} P(y_t | \mu_p) \right] \\ &= P(y_{ih}, \mu_{lp} | t_{lp}) P(y_{hj}, \mu_{pm} | t_{pm}) \quad (\text{any } p) \end{aligned} \quad (10)$$

This is our key recursion. We now marginalize over the inner boundaries $t_{l+1} \dots t_{m-1}$. If we first keep break t_p fixed at h , then the probability of $y_{ij} \mu_{lm}$ factorizes into $y_{ih} \mu_{lp}$ and $y_{hj} \mu_{pm}$. We can then average over location t_p . It is more convenient to define the

scaled quantity Q below and combine both steps:

$$Q(y_{ij}, \mu_{lm} | m-l) := \binom{j-i-1}{m-l-1} P(y_{ij}, \mu_{lm} | t_l, t_m) \quad (11)$$

$$\stackrel{(a)}{=} \binom{j-i-1}{m-l-1} \sum_{t_{lm}: i=t_l < \dots < t_m=j} P(y_{ij}, \mu_{lm} | t_{lm}) P(t_{lm} | t_l, t_m) \quad (12)$$

$$\stackrel{(b)}{=} \sum_{t_{lm}: i=t_l < \dots < t_m=j} P(y_{ij}, \mu_{lm} | t_{lm})$$

$$\stackrel{(c)}{=} \sum_{t_p=i+p-l}^{j+p-m} \sum_{t_{lp}: i=t_l < \dots < t_p=h} P(y_{ih}, \mu_{lp} | t_{lp}) \sum_{t_{pm}: h=t_p < \dots < t_m=j} P(y_{hj}, \mu_{pm} | t_{pm})$$

$$= \sum_{h=i+p-l}^{j+p-m} Q(y_{ih}, \mu_{lp} | p-l) Q(y_{hj}, \mu_{pm} | m-p) \quad (13)$$

(a) is just an instance of the law of total probability $P(A|B) = \sum_i P(A|H_i, B) P(H_i|B)$ for a partitioning (H_i) of the sample space, here $H_i \hat{=} t_{l+1} \dots t_{m-1}$. In (b) we exploited uniformity (8), which implies that $P(t_{lm} | t_l, t_m) = \binom{j-i-1}{m-l-1}^{-1}$ is also uniform, i.e. independent from the concrete segmentation $t_{l+1, m-1}$. In (c) we insert (10), fix segment boundary t_p , sum over the left and right segmentations, and finally over t_p . The sum starts at $h=i+p-l$, since y_{ih} must contain at least $p-l$ points in order to allow for $p-l$ segments, and ends at $h=j+p-m$, since y_{hj} must allow for $m-p$ segments.

Left and right recursions. If we integrate (11) over μ_{lm} , the integral factorizes and we get a recursion in (a quantity that is proportional to) the evidence of y_{ij} . Let us define more generally r^{th} “Q-moments” of μ'_t . Provided they exist, for $i < t \leq j$ we have

$$\begin{aligned} Q_t^r(y_{ij} | m-l) &:= \int Q(y_{ij}, \mu_{lm} | m-l) \mu_t'^r d\mu_{lm} \quad (14) \\ &= \sum_{h=i+p-l}^{t-1} Q^0(y_{ih} | p-l) Q_t^r(y_{hj} | m-p) + \sum_{h=t}^{j+p-m} Q_t^r(y_{ih} | p-l) Q^0(y_{hj} | m-p) \end{aligned}$$

Depending on whether $h < t$ or $h \geq t$, the $\mu_t'^r$ term combines with the right or left Q in recursion (13) to Q_t^r , while the other Q simply gets integrated to $Q_t^0 = Q^0$ independent of t . The recursion terminates with

$$A_{ij}^r := Q_t^r(y_{ij} | 1) = \int P(\mu_m) \prod_{t=i+1}^j P(y_t | \mu_m) \mu_m^r d\mu_m, \quad (0 \leq i < j \leq n) \quad (15)$$

Note $A_{ij}^0 = P(y_{ij} | t_{m-1, m})$ is the evidence and $A_{ij}^r / A_{ij}^0 = \mathbf{E}[\mu_m^r | y_{ij}, t_{m-1, m}]$ the r^{th} moment of $\mu'_t = \mu_m$ in case y_{ij} is modeled by a single segment. It is convenient to formally start the recursion with $Q^0(y_{ij} | 0) = \delta_{ij} = \begin{cases} 1 & \text{if } i=j \\ 0 & \text{else} \end{cases}$ (consistent with the recursion) with interpretation that only an empty data set ($i=j$) can have 0 segments. Since p was an arbitrary split number, we can choose it conveniently. We need a left recursion for $r=0$, $i=0$, $p-l=k$, and $m-p=1$:

$$L_{k+1, j} := Q^0(y_{0j} | k+1) = \sum_{h=k}^{j-1} Q^0(y_{0h} | k) Q^0(y_{hj} | 1) = \sum_{h=k}^{j-1} L_{kh} A_{hj}^0$$

That is, (apart from binomial factors) the evidence of y_{0j} with $k+1$ segments equals the evidence of y_{0h} with k segments times the single-segment evidence of y_{hj} , summed over all locations h of boundary k . The recursion starts with $L_{1j} = A_{0j}^0$, or more conveniently with $L_{0j} = \delta_{j0}$. A related recursion has been derived in (Fea06). We also need a right recursion for $r=0$, $j=n$, $p-l=1$, $m-p=k$:

$$R_{k+1,i} := Q^0(y_{in}|k+1) = \sum_{h=i+1}^{n-k} Q^0(y_{ih}|1)Q^0(y_{hn}|k) = \sum_{h=i+1}^{n-k} A_{ih}^0 R_{kh}$$

The recursion starts with $R_{1i} = A_{in}^0$, or more conveniently with $R_{0i} = \delta_{in}$.

6 Efficient Solution for Quantities of Interest

We now use the core recursion of the last section to derive expressions for all quantities of interest. We denote them by capital symbols (with indices).

Data evidence. Note that

$$L_{kn} = R_{k0} = Q^0(\mathbf{y}|k) = \binom{n-1}{k-1} P(\mathbf{y}|k)$$

are proportional to the data evidence for fixed k . So the data evidence can be computed as

$$E := P(\mathbf{y}) = \sum_{k=1}^n P(\mathbf{y}|k)P(k) = \frac{1}{k_{max}} \sum_{k=1}^{k_{max}} \frac{L_{kn}}{\binom{n-1}{k-1}} \quad (16)$$

Number of segments. The posterior of k and its MAP estimate are

$$C_k := P(k|\mathbf{y}) = \frac{P(\mathbf{y}|k)P(k)}{P(\mathbf{y})} = \frac{L_{kn}}{\binom{n-1}{k-1} k_{max} E} \quad \text{and} \quad \hat{k} = \arg \max_{k=1..k_{max}} C_k \quad (17)$$

Segment boundaries. We now determine the segment boundaries. Consider recursion (11) for $i=l=0$, $m=k$, $j=n$, but keep $t_p=h$ fixed, i.e. do not sum over it. Then (12) and (13) reduce to the l.h.s. and r.h.s. of

$$\binom{n-1}{k-1} P(\mathbf{y}, \boldsymbol{\mu}, t_p|k) = Q(y_{0h}, \mu_{0p}|p) Q(y_{hn}, \mu_{pk}|k-p) \quad (18)$$

Integration over $\boldsymbol{\mu}$ gives

$$\binom{n-1}{k-1} P(\mathbf{y}, t_p|k) = Q^0(y_{0h}|p) Q^0(y_{hn}|k-p)$$

Hence the posterior probability that boundary p is located at $t_p=h$, given \hat{k} , is

$$B_{ph} := P(t_p = h|\mathbf{y}, \hat{k}) = \frac{\binom{n-1}{\hat{k}-1} P(\mathbf{y}, t_p|\hat{k})}{\binom{n-1}{\hat{k}-1} P(\mathbf{y}|\hat{k})} = \frac{L_{ph} R_{\hat{k}-p,h}}{L_{\hat{k}n}} \quad (19)$$

So our estimate for segment boundary p is

$$\hat{t}_p := \arg \max_h P(t_p = h | \mathbf{y}, \hat{k}) = \arg \max_h \{B_{ph}\} = \arg \max_h \{L_{ph} R_{\hat{k}-p, h}\} \quad (20)$$

Segment levels. Finally we need the segment levels, given the segment number \hat{k} and boundaries $\hat{\mathbf{t}}$. The r^{th} moment of segment m with boundaries $i = \hat{t}_{m-1}$ and $j = \hat{t}_m$ is

$$\widehat{\mu}_m^r = \mathbf{E}[\mu_m^r | \mathbf{y}, \hat{\mathbf{t}}, \hat{k}] = \mathbf{E}[\mu_m^r | y_{ij}, \hat{t}_{m-1, m}] = \frac{\int P(y_{ij}, \mu_m | \hat{t}_{m-1, m}) \mu_m^r d\mu_m}{\int P(y_{ij}, \mu_m | \hat{t}_{m-1, m}) d\mu_m} = \frac{A_{ij}^r}{A_{ij}^0} \quad (21)$$

Note that this expression is independent of other segment boundaries and their number, as it should.

Regression curve. ⁵We reduce probabilities of μ'_t to probabilities of μ_m : We exploit the fact that in every segmentation, μ'_t lies in some segment. Let this (unique) segment be m with (unique) boundaries $i = t_{m-1} < t \leq t_m = j$. Then $\mu'_t = \mu_m$. Summing now over all such segments we get

$$P(\mu'_t | \mathbf{y}, k) = \sum_{m=1}^k \sum_{i=0}^{t-1} \sum_{j=t}^n P(\mu_m, t_{m-1} = i, t_m = j | \mathbf{y}, k) \quad (22)$$

By fixing t_p in (12) we arrived at (18). Similarly, dividing the data into three parts and fixing t_l and t_m we can derive

$$\binom{n-1}{k-1} P(\mathbf{y}, \boldsymbol{\mu}, t_l, t_m | k) = Q(y_{0i}, \mu_{0l} | l) Q(y_{ij}, \mu_{lm} | m - l) Q(y_{jn}, \mu_{mk} | k - m)$$

Setting $l = m - 1$, integrating over μ_{0l} and μ_{mk} , dividing by $\binom{n-1}{k-1} P(\mathbf{y} | k)$, and inserting into (22), we get

$$P(\mu'_t | \mathbf{y}, k) = \frac{1}{L_{kn}} \sum_{m=1}^k \sum_{i < t \leq j} L_{m-1, i} Q(y_{ij}, \mu_m | 1) R_{k-m, j}$$

The posterior moments of μ'_t , given \hat{k} , can hence be computed by

$$\widehat{\mu}'_t{}^r = \sum_{i < t \leq j} F_{ij}^r \quad \text{with} \quad F_{ij}^r := \frac{1}{L_{\hat{k}n}} \sum_{m=1}^{\hat{k}} L_{m-1, i} A_{ij}^r R_{\hat{k}-m, j} \quad (23)$$

While segment boundaries and values make sense only for fixed k (we chose \hat{k}), the regression curve $\hat{\mu}'_t$ could actually be averaged over all k instead of fixing $k = \hat{k}$.

Relative log-likelihood. Another quantity of interest is how likely it is that \mathbf{y} is sampled from \hat{f} . The log-likelihood of \mathbf{y} is

$$ll := \log P(\mathbf{y} | \hat{f}) = \log P(\mathbf{y} | \hat{\boldsymbol{\mu}}, \hat{\mathbf{t}}, \hat{k}) = \sum_{i=1}^n \log P(y_i | \hat{\mu}'_i, \sigma)$$

⁵Recursion (14) allows in principle to compute the regression curve $\mathbf{E}[\mu'_t | \mathbf{y}]$ by defining $(L_t^{r=1})_{kj}$ and $(R_t^{r=1})_{ki}$ analogous to L_{kj} and R_{ki} , but this procedure needs $O(n^3)$ space and $O(k_{max} n^3)$ time, one $O(n)$ worse than our target performance.

Like for the evidence, the number itself is hard to interpret. We need to know how many standard deviations it is away from its mean(=entropy). Since noise (1) is i.i.d., mean and variance of ll are just n times the mean and variance of the log-noise distribution of a single data item. For Gaussian and Cauchy noise we get

$$\begin{aligned} \text{Gauss:} \quad \mathbf{E}[ll|\hat{f}] &= \frac{n}{2} \log(2\pi e\sigma^2), & \text{Var}[ll|\hat{f}] &= \frac{n}{2} \\ \text{Cauchy:} \quad \mathbf{E}[ll|\hat{f}] &= n \log(4\pi\sigma), & \text{Var}[ll|\hat{f}] &= \frac{n}{3}\pi^2 \end{aligned}$$

7 Computing the Single Segment Distribution

We now determine (at least in the Gaussian case efficient) expressions for the moments (15) of the distribution (9) of a single segment.

Gaussian model. For Gaussian noise (4) and prior (5) we get

$$A_{ij}^r = \left(\frac{1}{\sqrt{2\pi}\sigma} \right)^d \frac{1}{\sqrt{2\pi}\rho} \int_{-\infty}^{\infty} e^{-\frac{1}{2\sigma^2} \sum_{t=i+1}^j (y_t - \mu_m)^2 - \frac{1}{2\rho^2} (\mu_m - \nu)^2} \mu_m^r d\mu_m$$

where $d=j-i$. This is an unnormalized Gaussian integral with the following normalization, mean, and variance (Bol04, Sec.10.2):

$$P(y_{ij}|t_{m-1,m}) = A_{ij}^0 = \frac{\exp\left\{\frac{1}{2\sigma^2} \left[\frac{(\sum_t (y_t - \nu))^2}{d + \sigma^2/\rho^2} - \sum_t (y_t - \nu)^2 \right]\right\}}{(2\pi\sigma^2)^{d/2} (1 + d\rho^2/\sigma^2)^{1/2}} \quad (24)$$

$$\mathbf{E}[\mu_m|y_{ij}, t_{m-1,m}] = \frac{A_{ij}^1}{A_{ij}^0} = \frac{\rho^2 (\sum_t y_t) + \sigma^2 \nu}{d\rho^2 + \sigma^2} \approx \frac{1}{d} \sum_{t=i+1}^j y_t \quad (25)$$

$$\text{Var}[\mu_m|y_{ij}, t_{m-1,m}] = \frac{A_{ij}^2}{A_{ij}^0} - \left(\frac{A_{ij}^1}{A_{ij}^0} \right)^2 = \left[\frac{d}{\sigma^2} + \frac{1}{\rho^2} \right]^{-1} \approx \frac{\sigma^2}{d} \quad (26)$$

where \sum_t runs from $i+1$ to j . The mean/variance is just the weighted average of the mean/variance of y_{ij} and μ_m . One may prefer to use the segment prior only for determining A_{ij}^0 , but use the unbiased estimators (\approx) for the moments. Higher moments A_{ij}^r can also be computed from the central moments

$$\mathbf{E}[(\mu_m - A_{ij}^1/A_{ij}^0)^r | y_{ij}, t_{m-1,m}] = \frac{1 \cdot 3 \cdot \dots \cdot (r-1)}{[d\sigma^{-2} + \rho^{-2}]^{r/2}} \approx 1 \cdot 3 \cdot \dots \cdot (r-1) \cdot \left(\frac{\sigma^2}{d} \right)^{r/2}$$

for even r , and 0 for odd r .

Other models. Analytic expressions for A_{ij}^r are possible for all distributions in the exponential family. For others like Cauchy we need to perform integral (15) numerically. A very simple approximation is to replace the integral by a sum on a uniform grid: The stepsize/range of the grid should be some fraction/multiple of the typical scale of the integrand, and the center of the grid should be around the mean. A crude estimate of the mean and scale can be obtained from the Gaussian model (25) and (26). Or even

simpler, use the estimated global mean and variance (27), and in-segment variance (28) for determining the range (e.g. $[\hat{\nu} - 25\hat{\rho}, \dots, \hat{\nu} + 25\hat{\rho}]$) and stepsize (e.g. $\hat{\sigma}/10$) of one grid used for all A_{ij}^r . Note that if y_{ij} really stem from one segment, the integrand is typically unimodal and the above estimates for stepsize and range are reasonable, hence the approximation will be good. If y_{ij} ranges over different segments, the discretization may be crude, but since in this case, A_{ij}^r is (very) small, crude estimates are sufficient. Note also that even for the heavy-tailed Cauchy distribution, the first and second moments A_{ij}^1 and A_{ij}^2 exist, since the integrand is a product of at least two Cauchy distributions, one prior and one noise for each y_t . Preferably, standard numerical integration routines (which are faster, more robust and more accurate) should be used.

8 Determination of the Hyper-Parameters

Hyper-Bayes and Hyper-ML. The developed regression model still contains three (hyper)parameters, the global variance ρ^2 and mean ν of $\boldsymbol{\mu}$, and the in-segment variance σ^2 . If they are not known, a proper Bayesian treatment would be to assume a hyper-prior over them and integrate them out. Since we do not expect a significant influence of the hyper-prior (as long as chosen reasonable) on the quantities of interest, one could more easily proceed in an empirical Bayesian way and choose the parameters such that the evidence $P(\mathbf{y}|\sigma, \nu, \rho)$ is maximized (“hyper-ML” or “ML-IP”). (We restored the till now omitted dependency on the hyper-parameters).

Exhaustive (grid) search for the hyper-ML parameters is expensive. For data which is indeed noisy piecewise constant, $P(\mathbf{y}|\sigma, \nu, \rho)$ is typically unimodal⁶ in (σ, ν, ρ) and the global maximum can be found more efficiently by greedy hill-climbing, but even this may cost a factor of 10 to 1000 in efficiency. Below we present a very simple and excellent heuristic for choosing (σ, ν, ρ) .

Estimate of global mean and variance ν and ρ . A reasonable choice for the level mean and variance ν and ρ are the empirical global mean and variance of the data \mathbf{y} .

$$\hat{\nu} \approx \frac{1}{n} \sum_{t=1}^n y_t \quad \text{and} \quad \hat{\rho}^2 \approx \frac{1}{n-1} \sum_{t=1}^n (y_t - \hat{\nu})^2 \tag{27}$$

This overestimates the variance ρ^2 of the segment levels, since the expression also includes the in-segment variance σ^2 , which one may want to subtract from this expression.

Estimate of in-segment variance σ^2 . At first there seems little hope of estimating the in-segment variance σ^2 from \mathbf{y} without knowing the segmentation, but actually we can use a simple trick. If \mathbf{y} would belong to a single segment, i.e. the y_t were i.i.d. with

⁶A little care is necessary with the in-segment variance σ^2 . If we set it (extremely close) to zero, all segments will consist of a single data point y_t with (close to) infinite evidence (see e.g. (24)). Assuming $k_{max} < n$ eliminates this unwished maximum. Greedy hill-climbing with proper initialization will also not be fooled.

variance σ^2 , then the following expressions for σ^2 would hold:

$$\mathbf{E}\left[\frac{1}{n}\sum_{t=1}^n(y_t - \mu_1)^2\right] = \sigma^2 = \frac{1}{2(n-1)}\mathbf{E}\left[\sum_{t=1}^{n-1}(y_{t+1} - y_t)^2\right]$$

i.e. instead of estimating σ^2 by the squared deviation of the y_t from their mean, we can also estimate σ^2 from the average squared difference of successive y_t . This remains true even for multiple segments if we exclude the segment boundaries in the sum. On the other hand, if the number of segment boundaries is small, the error from including the boundaries will be small, i.e. the second expression remains approximately valid. More precisely, we have within a segment and at the boundaries

$$\begin{aligned}\mathbf{E}\left[\sum_{t=t_{m-1}+1}^{t_m-1}(y_{t+1} - y_t)^2\right] &= 2(t_m - t_{m-1} - 1)\sigma^2 \quad \text{and} \\ \mathbf{E}[(y_{t_{m+1}} - y_{t_m})^2] &= 2\sigma^2 + (\mu_{m+1} - \mu_m)^2\end{aligned}$$

Summing over all k segments and $k-1$ boundaries and solving w.r.t. σ^2 we get

$$\begin{aligned}\sigma^2 &= \frac{1}{2(n-1)}\left\{\mathbf{E}\left[\sum_{t=1}^{n-1}(y_{t+1} - y_t)^2\right] - \sum_{m=1}^{k-1}(\mu_{m+1} - \mu_m)^2\right\} \\ &= \frac{1}{2(n-1)}\mathbf{E}\left[\sum_{t=1}^{n-1}(y_{t+1} - y_t)^2\right] \cdot \left[1 - O\left(\frac{k}{n}\frac{\rho^2}{\sigma^2}\right)\right]\end{aligned}$$

The last expression holds, since there are k boundaries in n data items, and the ratio between the variance of $\boldsymbol{\mu}$ to the in-segment variance is ρ^2/σ^2 . Hence we may estimate σ^2 by the upper bound

$$\hat{\sigma}^2 \approx \frac{1}{2(n-1)}\sum_{t=1}^{n-1}(y_{t+1} - y_t)^2 \quad (28)$$

If there are not too many segments ($k \ll n$) and the regression problem is hard (high noise $\rho \lesssim \sigma$), this is a very good estimate. In case of low noise ($\rho \gg \sigma$), regression is very easy, and a crude estimate of σ^2 is sufficient. If there are many segments, $\hat{\sigma}^2$ tends to overestimate σ^2 , resulting in a (marginal) bias towards estimating fewer segments (which is then often welcome).

If the estimate is really not sufficient, one may use (28) as an initial estimate for determining an initial segmentation \hat{t} , which then can be used to compute an improved estimate of $\hat{\sigma}^2$, and possibly iterate.

Hyper-ML estimates. Expressions (27) are the standard estimates of mean and variance of a distribution. They are particularly suitable for (close to) Gaussian distributions, but also for others, as long as ν and ρ parameterize mean and variance. If mean and variance do not exist or the distribution is quite heavy-tailed, we need other

estimates. The “ideal” hyper-ML estimates may be approximated as follows. If we assume that each data point lies in its own segment, we get

$$\begin{aligned}
 (\hat{\nu}, \hat{\rho}) &\approx \arg \max_{(\nu, \rho)} \prod_{t=1}^n P(y_t | \hat{\sigma}, \nu, \rho) \quad \text{with} \\
 P(y_t | \sigma, \nu, \rho) &= \int P(y_t | \mu, \sigma) P(\mu | \nu, \rho) d\mu
 \end{aligned}
 \tag{29}$$

The in-segment variance $\hat{\sigma}^2$ can be estimated similarly to the last paragraph considering data differences and ignoring segment boundaries:

$$\begin{aligned}
 \hat{\sigma} &\approx \arg \max_{\sigma} \prod_{t=1}^{n-1} P(y_{t+1} - y_t | \sigma) \quad \text{with} \\
 P(y_{t+1} - y_t = \Delta | \sigma) &\approx \int_{-\infty}^{\infty} P(y_{t+1} = a + \Delta | \mu, \sigma) P(y_t = a | \mu, \sigma) da
 \end{aligned}
 \tag{30}$$

Note that the last expression is independent of the segment level (this was the whole reason for considering data differences) and exact iff y_t and y_{t+1} belong to the same segment. In general (beyond the exponential family) $(\hat{\nu}, \hat{\rho}, \hat{\sigma})$ can only be determined numerically.

Using median and quartile. We present some simpler estimates based on median and quartiles. Let $[\mathbf{y}]$ be the data vector \mathbf{y} , but sorted in ascending order. Then, item $[\mathbf{y}]_{\alpha n}$ (where the index is assumed to be rounded up to the next integer) is the α -quantile of empirical distribution \mathbf{y} . In particular $[\mathbf{y}]_{n/2}$ is the median of \mathbf{y} . It is a consistent and robust-to-outliers estimator of the mean segment level

$$\hat{\nu} \approx [\mathbf{y}]_{n/2}
 \tag{31}$$

if noise and segment levels have symmetric distributions. Further, half of the data points lie in the interval $[a, b]$, where $a := [\mathbf{y}]_{n/4}$ is the first and $b := [\mathbf{y}]_{3n/4}$ is the last quartile of \mathbf{y} . So, using (29), $\hat{\rho}$ should (!) be estimated such that

$$P(a \leq y_t \leq b | \sigma, \hat{\nu}, \hat{\rho}) \stackrel{!}{\approx} \frac{1}{2}$$

Ignoring data noise (assuming $\sigma \approx 0$), we get

$$\hat{\rho} \approx \frac{[\mathbf{y}]_{3n/4} - [\mathbf{y}]_{n/4}}{2\alpha} \quad \text{with } \alpha = 1 \text{ for Cauchy and } \alpha \doteq 0.6744 \text{ for Gauss,}
 \tag{32}$$

where α is the quartile of the standard Cauchy/Gauss/other segment prior. For the data noise σ we again consider the differences $\Delta_t := y_{t+1} - y_t$. Using (30), $\hat{\sigma}$ should be estimated such that

$$P(a' \leq y_{t+1} - y_t \leq b' | \hat{\sigma}) \stackrel{!}{\approx} \frac{1}{2}$$

where $a' = [\Delta]_{n/4}$ and $b' = [\Delta]_{3n/4} \approx -a'$. One can show that

$$\hat{\sigma} \approx \frac{[\Delta]_{3n/4} - [\Delta]_{n/4}}{2\beta} \quad \text{with } \beta = 2 \text{ for Cauchy and } \beta \doteq 0.6744\sqrt{2} \text{ for Gauss,}
 \tag{33}$$

where β is the quartile of the one time with itself convolved standard Cauchy/Gauss/other (noise) distribution. Use of quartiles for estimating σ is very robust to the “outliers” caused by the segment boundaries, so yields better estimates than (28) if noise is low. Again, if the estimates are really not sufficient, one may iteratively improve them.

9 The Algorithm

The computation of A , L , R , E , C , B , \hat{t}_p , $\hat{\mu}_m^r$, F , and $\hat{\mu}_t^r$ by the formulas/recursions derived in Sections 5–7, are straightforward. In (15) one should compute the product, or in (24), (25), (26) the sum, incrementally from $j \rightsquigarrow j+1$. Similarly $\hat{\mu}_t^r$ should be computed incrementally by

$$\hat{\mu}_{t+1}^r = \hat{\mu}_t^r - \sum_{i=0}^{t-1} F_{it}^r + \sum_{j=t+1}^n F_{tj}^r$$

Typically $r=0,1,2$. In this way, all quantities can be computed in time $O(k_{max}n^2)$ and space $O(n^2)$, as clear from Algorithm 1. Space can be reduced to $O(k_{max}n)$ by computing A and F on-the-fly in the various expressions at the cost of a slowdown by a constant factor. Table 1 contains the algorithm in pseudo-C code. The complete code including examples and data is available at (Hut05a). Since A^0 , L , R , and E can be exponentially large or small in n , i.e. huge or tiny, actually their logarithm has to be computed and stored. In the expressions, the logarithm is pulled in by $\log(x \cdot y) = \log(x) + \log(y)$ and $\log(x + y) = \log(x) + \log(1 + \exp(\log(y) - \log(x)))$ for $x > y$ and similarly for $x < y$. Instead of A_{ij}^r we have to compute A_{ij}^r/A_{ij}^0 by pulling the denominator into the integral.

10 Synthetic Examples

Description. In order to test our algorithm we created various synthetic data sets. We considered piecewise constant functions with noisy observations. The considered function was defined -1 in its first quarter, $+1$ in its second quarter, and 0 in the last half. So the function consists of two small and one large segments, with a large jump at the first and a small jump at the second boundary. For n we chose 100, i.e. $f_1 \dots f_{25} = -1$, $f_{26} \dots f_{50} = +1$, and $f_{51} \dots f_{100} = 0$. Data y_t was obtained by adding independent Gaussian/Cauchy noise of same scale σ for all t . We considered low $\sigma=0.1$, medium $\sigma=0.32$, and high $\sigma=1$ noise, resulting in an easy, medium, and hard regression problem (Figures 1-14). We applied our regression algorithm to these 6 data sets (named GL,GM,GH,CL,CM,CH), where we modeled noise and prior as Gaussian or Cauchy with hyper-parameters also estimated by the algorithms in Table 1. Table 2 contains these and other scalar summaries, like the evidence, likelihood, MAP segment number \hat{k} and their probability.

Three segment Gaussian with low noise. Regression for low Gaussian noise ($\sigma =$

Table 1: Regression algorithm in pseudo C code

EstGauss(\mathbf{y}, n) and **EstGeneral**($\mathbf{y}, n, \alpha, \beta$) compute from data (y_1, \dots, y_n) , estimates for ν, ρ, σ (hat ‘^’ omitted), and from that the evidence A_{ij}^0 of a single segment ranging from $i+1$ to j , and corresponding first and second moments A_{ij}^1 and A_{ij}^2 . The expressions (27), (28), (24), (25), (26) are used in EstGauss() for Gaussian noise and prior, and (31), (32), (33) and numerical integration on a uniform Grid in EstGeneral() for arbitrary noise and prior P , e.g. Cauchy. $[\mathbf{y}]$ denotes the sorted \mathbf{y} array, Grid is the uniform integration grid, += and *= are additive/multiplicative updates, and [] denotes arrays.

EstGauss($\mathbf{y}_{[]}, n$)

```
[ ν = 1/n ∑_{t=1}^n y_t;
  ρ² = 1/(n-1) ∑_{t=1}^n (y_t - ν)²;
  σ² = 1/(2(n-1)) ∑_{t=1}^{n-1} (y_{t+1} - y_t)²;
  for (i=0..n)
    [ m=0; s=0;
      for (j=i+1..n)
        [ d=j-i; m+=y_j-ν; s+=(y_j-ν)²;
          A_{ij}^0 = exp{ 1/(2σ²) [ d + σ²/ρ² - s ] } /
            (2πσ²)^{d/2} (1+dρ²/σ²)^{1/2};
          A_{ij}^1 = A_{ij}^0 (ν+m/d);
          [ [ A_{ij}^2 = A_{ij}^0 ((A_{ij}^1/A_{ij}^0)² + σ²/d);
        ]
      ]
    ]
  [ return (A_{[] []}^0, ν, ρ, σ);
```

Regression(\mathbf{A}, n, k_{max}) takes \mathbf{A} , n , and an upper bound on the number of segments k_{max} , and computes the evidence $E = P(\mathbf{y})$ (16), the probability $C_k = P(k|\mathbf{y})$ of k segments and its MAP estimate \hat{k} (17), the probability $B_i = P(\exists p : t_p = i | \mathbf{y}, \hat{k})$ that a boundary is at i (19) and the MAP location \hat{t}_p of the p^{th} boundary (20), the first and second segment level moments μ_p and μ_p^2 of all segments p (21), and the Bayesian regression curve μ'_t and its second moment $\mu_t'^2$ (23).

EstGeneral($\mathbf{y}_{[]}, n, \alpha, \beta$)

```
[ ν = [y]_{n/2};
  ρ = ([y]_{3n/4} - [y]_{n/4}) / 2α;
  for (t=1..n-1) Δ_t = y_{t+1} - y_t;
  σ = ([Δ]_{3n/4} - [Δ]_{n/4}) / 2β;
  Grid = (σ/10 Z) ∩ [ν - 25ρ, ν + 25ρ];
  for (i=0..n)
    [ for (μ ∈ Grid) R_μ = P(μ | ν, ρ);
      for (j=i+1..n)
        [ for (μ ∈ Grid) R_μ * = P(y_j | μ, σ);
        ]
      [ [ A_{ij}^r = σ/10 ∑_{μ ∈ Grid} R_μ μ^r; (r=0,1,2)
    ]
  ]
  [ return (A_{[] []}^0, ν, ρ, σ);
```

Regression($\mathbf{A}_{[] []}^0, n, k_{max}$)

```
[ for (i=0..n) { L_{0i} = δ_{i0}; R_{0i} = δ_{in}; }
  for (k=0..k_{max}-1)
    [ for (i=0..n) L_{k+1,i} = ∑_{h=k}^{i-1} L_{kh} A_{hi}^0;
      [ for (i=0..n) R_{k+1,i} = ∑_{h=i+1}^{n-k} A_{ih}^0 R_{kh};
      E = k_{max}^{-1} ∑_{k=1}^{k_{max}} L_{kn} / \binom{n-1}{k-1};
      for (k=0..k_{max}) C_k = L_{kn} / [ \binom{n-1}{k-1} k_{max} E ];
      \hat{k} = argmax_{k=1..k_{max}} { C_k };
      for (i=0..n) B_i = ∑_{p=0}^{\hat{k}} L_{pi} R_{\hat{k}-p,i} / L_{\hat{k}n};
      for (p=0..{\hat{k}}) \hat{t}_p = argmax_h { L_{ph} R_{\hat{k}-p,h} };
      for (p=1..{\hat{k}}) \widehat{\mu}_p^r = A_{\hat{t}_{p-1}\hat{t}_p}^r / A_{\hat{t}_{p-1}\hat{t}_p}^0; (r=1,2)
      for (i=0..n) for (j=i+1..n)
        [ F_{ij}^r = ∑_{m=1}^{\hat{k}} L_{m-1,i} A_{ij}^r R_{\hat{k}-m,j} / L_{\hat{k}n};
          \mu_0'^r = 0; (r=1,2)
          for (t=0..n-1)
            [ \widehat{\mu}_{t+1}'^r = \widehat{\mu}_t'^r - ∑_{i=0}^{t-1} F_{it}^r + ∑_{j=t+1}^n F_{tj}^r
          ]
        ]
      [ return (E, C_{[]}, \hat{k}, B_{[]}, \widehat{t}_{[]}, \widehat{\mu}_{[]}'^r, \widehat{\mu}_{[]}'^r);
```




Figure 1: [GL: low Gaussian noise] data (blue), PCR (black), BP (red), and variance^{1/2} (green).




Figure 2: [GM: medium Gaussian noise] data (blue), PCR (black), BP (red), and variance^{1/2} (green).




Figure 3: [GM: medium Gaussian noise] data with Bayesian regression ± 1 std.-deviation.




Figure 4: [GH: high Gaussian noise] data.

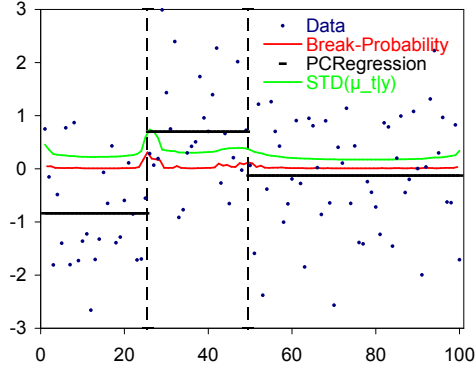


Figure 5: [GH: high Gaussian noise] data (blue), PCR (black), BP (red), and variance^{1/2} (green).




Figure 6: [GH: high Gaussian noise] data with Bayesian regression ± 1 std.-deviation.

0.1) is very easy. Figure 1 shows the data points $(1, y_1), \dots, (100, y_{100})$ together with the estimated segment boundaries and levels, i.e. the Piecewise Constant Regression (PCR) curve (black). The red BP curve (with the two spikes) is the posterior probability that a boundary (break point BP) is at t . It is defined as $B_t := \sum_{p=1}^k B_{pt}$. Our Bayesian regressor (BPCR) is virtually sure that the boundaries are at $t_1 = 25$ ($B_{25} = 100\%$) and $t_2 = 50$ ($B_{50} = 99.9994\%$). The segment levels $\hat{\mu}_1 = -0.98 \approx -1$, $\hat{\mu}_2 = 0.97 \approx 1$, $\hat{\mu}_3 = 0.01 \approx 0$ are determined with high accuracy i.e. with low standard deviation (green curve) $\sigma/\sqrt{25} = 2\%$ for the first two and $\sigma/\sqrt{50} \doteq 1.4\%$ for the last segment. The Bayesian regression curve $\hat{\mu}_t$ (not shown) is indistinguishable from the PCR.

Three segment Gaussian with medium noise. Little changes for medium Gaussian noise ($\sigma = 0.32$). Figure 2 shows that the number and location of boundaries is still correctly determined, but the posterior probability of the second boundary location (red BP curve) starts to get a little broader ($B_{50} = 87\%$). The regression curve in Figure 3 is still essentially piecewise constant. At $t = 50$ there is a small kink and the error band gets a little wider, as can better be seen in the (kink of the) green $\sqrt{\text{Var}[\mu'_t | \dots]}$ curve in Figure 2. In Figure 13 we study the sensitivity of our regression to the noise estimate $\hat{\sigma}$. Keeping everything else fixed, we varied σ from 0.1 to 1 and plotted the log-evidence $\log P(\mathbf{y}|\sigma)$ and the segment number estimate $\hat{k}(\sigma)$ as a function of σ . We see that our estimate $\hat{\sigma} \doteq 0.35$ is close to the hyper-ML value $\sigma_{\text{HML}} = \text{argmax}_{\sigma} P(\mathbf{y}|\sigma) \doteq 0.33$, which itself is close to the true $\sigma = 0.32$. The number of segments \hat{k} is correctly recovered for a wide range of σ around $\hat{\sigma}$. If σ is chosen too small (below the critical value 0.2), BPCR cannot regard typical deviations from the segment level as noise anymore and has to break segments into smaller pieces for a better fit (\hat{k} increases). For higher noise, the critical value gets closer to $\hat{\sigma}$, but also the estimate becomes (even) better. For lower noise, $\hat{\sigma}$ overestimates the true σ , but BPCR is at the same time even less sensitive to it.

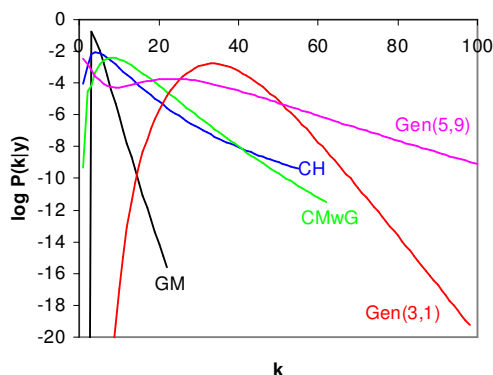


Figure 7: Posterior segment number probability $P(k|\mathbf{y})$ for medium Gaussian noise (GM, black), high Cauchy noise (CH, blue), medium Cauchy noise with Gaussian regression (CMwG, green), aberrant gene copy # of chromosome 1 (Gen(3,1), red), normal gene copy # of chromosome 9 (Gen(5,9), pink).




Figure 8: [CM: medium Cauchy noise] data (blue), PCR (black), BP (red), and variance^{1/2} (green).

Three segment Gaussian with high noise. Figure 4 shows the data with Gaussian noise of the same order as the jump of levels ($\sigma=1$). One can imagine some up-trend in the first quarter, but one can hardly see any segments. Nevertheless, BPCR still finds the correct boundary number and location of the first boundary (Figure 5). The second boundary is one off to the left, since y_{50} was accidentally close to zero, hence got assigned to the last segment. The (red) boundary probability (BP) curve is significantly blurred, in particular at the smaller second jump with quite small $B_{49} = 12\%$ and $B_{50} = 10\%$. The levels themselves are within expected accuracy $\sigma/\sqrt{25} = 20\%$ and $\sigma/\sqrt{50} \doteq 14\%$, respectively, yielding still a PCR close to the true function. The Bayesian regression (and error) curve (Figure 6), though, changed shape completely. It resembles more a local data smoothing, following trends in the data (more on this in the next section). The variance (green curve in Figure 5) has a visible bump at $t=25$, but only a broad slight elevation around $t=50$.

Three segment Cauchy. The qualitative results for the Cauchy with low noise ($\sigma=0.1$) are the same as for Gauss, perfect recovery of the underlying function, and is hence not shown. Worth mentioning is that the estimate $\hat{\sigma}$ based on quartiles is excellent(ly close to hyper-ML) even for this low noise (and of course higher noise), i.e. is very robust against the segment boundaries.

Also for medium Cauchy noise ($\sigma=0.32$, Figure 8) our BPCR does not get fooled (even) by (clusters of) “outliers” at $t=16$, $t=48,49$, and $t=86,89,90$. The second boundary is one off to the right, since y_{51} is slightly too large. Break probability B_t




Figure 9: [CH: high Cauchy noise] data.




Figure 10: [CH: high Cauchy noise] data (blue), PCR (black), BP (red), and variance^{1/2} (green).




Figure 11: [CH: high Cauchy noise] data with Bayesian regression ± 1 std.-deviation.




Figure 12: [CMwG: medium Cauchy noise] data (blue), but with Gaussian PCR (black), BP (red), and variance^{1/2} (green).




Figure 13: [GM: medium Gaussian noise] $\log P(\mathbf{y})$ (blue) and \hat{k} (green) as function of σ and our estimate $\hat{\sigma}$ of $(\arg)\max_{\sigma} P(\mathbf{y})$ and $\hat{k}(\hat{\sigma})$ (black triangles).

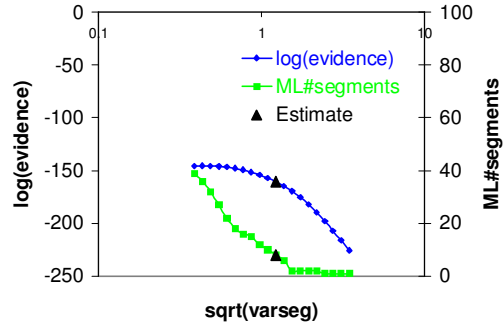


Figure 14: [CMwG: medium Cauchy noise] with Gaussian regression, $\log P(\mathbf{y})$ (blue) and \hat{k} (green) as function of σ and our estimate $\hat{\sigma}$ of $(\arg)\max_{\sigma} P(\mathbf{y})$ and $\hat{k}(\hat{\sigma})$ (black triangles).

(red) and variance $\text{Var}[\mu'_t | \mathbf{y}, \hat{k}]$ (green) are nicely peaked at $\hat{t}_1 = 25$ and $\hat{t}_2 = 51$.

For high Cauchy noise ($\sigma = 1$, Figure 9) it is nearly impossible to see any segment (levels) at all. Amazingly, BPCR still recovers three segments (Figure 10), but the first boundary is significantly displaced ($\hat{t}_1 = 14$). B_t and $\text{Var}[\mu'_t | \mathbf{y}, \hat{k}]$ contain many peaks indicating that BPCR was quite unsure where to break. The Bayesian regression in Figure 11 identifies an upward trend in the data $y_{14:35}$, explaining the difficulty/impossibility of recovering the correct location of the first boundary.

Cauchy analyzed with Gauss and vice versa. In order to test the robustness of BPCR under misspecification, we analyzed the data with Cauchy noise by Gaussian BPCR (and vice versa). Gaussian BPCR perfectly recovers the segments for low Cauchy noise. For medium noise (CMwG, Figure 12) the outlier at $t = 49$ is not tolerated and placed in its own segment, and the last segment is broken in two parts, but overall the distortion is less than possibly expected (e.g. not all outliers are in their own segments). The reason for this robustness can be attributed to the way we estimate σ . Figure 14 shows that the outliers have increased $\hat{\sigma}$ far beyond the peak of $P(\mathbf{y} | \sigma)$, which in turn leads to lower (more reasonable) number of segments. This is a nice stabilizing property of $\hat{\sigma}$. The other way round, segmentation of data with medium Gaussian noise is essentially insensitive to whether performed with Gaussian BPCR (Fig. 2 and 3) or Cauchy BPCR (GMwC, not shown), which confirms (once again) the robustness of the Cauchy model. But for high noise BPCR fails in both misspecification directions.

11 Real-World Example & More Discussion

Gene copy number data. We now turn to a real-world data set. All chromosomes (except for the sex chromosomes in males) in a healthy human cell come in pairs, but

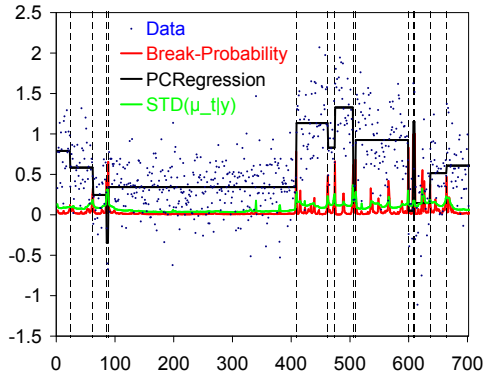


Figure 15: [Gen31: Aberrant gene copy # of chromosome 1] data (blue), PCR (black), BP (red), and variance^{1/2} (green).

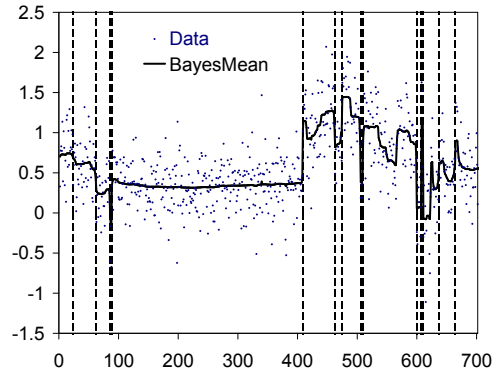


Figure 16: [Gen31: Aberrant gene copy # of chromosome 1] data with Bayesian regression ± 1 std.-deviation.

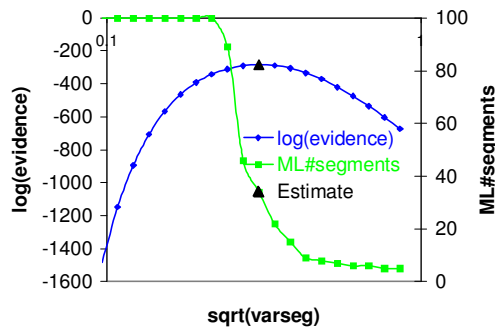


Figure 17: [Gen31: Aberrant gene copy # of chromosome 1] $\log P(\mathbf{y})$ (blue) and \hat{k} (green) as function of σ and our estimate $\hat{\sigma}$ of $(\arg)\max_{\sigma} P(\mathbf{y})$ and $\hat{k}(\hat{\sigma})$ (black triangles).




Figure 18: [Gen59: normal gene copy # of chromosome 9] with Bayesian regression.

pieces or entire chromosomes can be lost or multiplied in tumor cells. With modern micro-arrays one can measure the local copy number along a chromosome. It is important to determine the breaks, where copy-number changes. The measurements are very noisy (Pin98). Hence this is a natural application for piecewise constant regression of noisy (one-dimensional) data. An analysis with BPCR of chromosomal aberrations of real tumor samples, its biological interpretation, and comparison to other methods will be given elsewhere. Here, we only show the regression results of one aberrant and one healthy chromosome (without biological interpretation).

Most tumor samples are impure, so the copy number is unfortunately not an integer. The “log-ratios” \mathbf{y} of the copy numbers of a normal cell (and also the Δ of any cell) are very close to Gaussian distributed, so we chose Gaussian BPCR. The log-ratios \mathbf{y} of chromosome 1 of a sample known to have multiple myeloma are shown in Figure 15, together with the regression results. Visually, the segmentation is very reasonable. Long segments (e.g. $t = 89 \dots 408$) as well as very short ones around $t = 87$ and 641 of length 3 are detected. The Bayesian regression curve in Figure 16 also behaves nicely. It is very flat i.e. smoothes the data in long and clear segments, wiggles in less clear segments, and has jumps at the segment boundaries. Compare this to local smoothing techniques (Rin06), which wiggle much more within a segment and severely smooth boundaries. In this sense our Bayesian regression curve is somewhere in-between local smoothing and hard segmentation. We also see that the regression curve has a broad dip around $t = 535 \dots 565$, although $t = 510 \dots 599$ has been assigned to a single segment. This shows that other contributions breaking the segment have been mixed into the Bayesian regression curve. The PCR favor for a single segment is close to “tip over” as can be seen from the spikes in the break probability (red curve) in this segment.

The dependence of evidence and segment number on σ is shown in Figure 17. Our estimate $\hat{\sigma}$ (black triangle) perfectly maximizes $P(\mathbf{y}|\sigma)$ (blue curve). It is at a deep slope of $P(k|\mathbf{y},\sigma)$ (green curve), which means that the segmentation is sensitive to a good estimate of $\hat{\sigma}$. There is no unique (statistically) correct segmentation (number). Various segmentations within some range are supported by comparable evidence.

Figure 18 shows a healthy chromosome 9, correctly lumped into one big segment.

Posterior probability of the number of segments $P(k|\mathbf{y})$. One of the most critical steps for good segmentation is determining the right segment number, which we did by maximizing $P(k|\mathbf{y})$. The whole curves shown in Figure 7 give additional insight. A representative selection is presented.

For truly piecewise constant functions with $k_0 \ll n$ segments and low to medium noise, $\log P(k|\mathbf{y})$ typically raises rapidly with k till k_0 and thereafter decays approximately linear (black GM curve). This shows that BPCR certainly does not underestimate k_0 ($P(k < k_0|\mathbf{y}) \approx 0$). Although it also does not overestimate k_0 , only $P(k \geq k_0|\mathbf{y}) \approx 1$, but $P(k_0|\mathbf{y}) \not\approx 1$ due to the following reason: If a segment is broken into two (or more) and assigned (approximately) equal levels, the curve and hence the likelihood does not change. BPCR does not explicitly penalize this, only implicitly by the Bayesian averaging (Bayes factor phenomenon (Goo83; Jay03; Mac03)). This gives very roughly

Table 2: **(Regression summary)** The table summarizes the regression results for the considered synthetic and real-world data sets described in the main text: Data consisting of three segments with (low,medium,high) Gaussian/Cauchy noise (GL,GM,GH)/(CL,CM,CH), cross regressed (GMwC,CMwG), and the (tumor,healthy) DNA copy number data (Gen31,Gen59). The second last column contains \hat{k} with and without counting multiple breaks at the same point. The last column contains $P(\hat{k})$ and in parentheses $P(\hat{k}-1)$ and $P(\hat{k}+1)$ in percent.

Gauss, Cauchy, Low, Medium, High noise, Gene	true noise scale	data size	method	global mean estimate	global deviation estimate	in-segment deviation est.	log-evidence $\log P(\mathbf{y})$	rel. log-likelihood $\frac{U - \mathbf{E}[U]}{\text{Var}[U]^{1/2}}$	Opt.# segm.	Confidence $P(\hat{k}(-1,+1) \mathbf{y})$
Name	σ	n	P	$\hat{\nu}$	$\hat{\rho}$	$\hat{\sigma}$	$\log E$	$\frac{U - \mathbf{E}}{\sigma_U}$	\hat{k}	$C_{k(-1,+1)}$
GL	0.10	100	G	-0.01	0.69	0.18	39	4.9	3 3	74%(0 20)
GM	0.32	100	G	-0.03	0.73	0.35	-48	1.2	3 3	44%(0 29)
GH	1.00	100	G	-0.10	1.15	1.03	-156	0.3	3 4	13%(10 12)
CL	0.10	100	C	-0.02	0.58	0.09	-17	1.0	3 3	69%(0 21)
CM	0.32	100	C	-0.09	0.70	0.27	-127	0.8	3 3	38%(0 27)
CH	1.00	100	C	-0.20	0.99	0.86	-234	0.9	3 4	12%(11 11)
GMwC	0.32	100	C	0.00	0.49	0.17	-70	1.5	3 3	27%(0 26)
CMwG	0.32	100	G	0.01	1.24	1.22	-160	2.9	5 8	8%(8 8)
Gen31	-	769	G	0.55	0.45	0.30	-283	-1.5	15 34	6%(6 6)
Gen59	-	483	G	1.05	0.47	0.44	-336	-2.3	1 1	8%(0 6)

an additive term in the log-likelihood of $\frac{1}{2} \log n$ for each additional degree of freedom (segment level and boundary). This observation is the core of the Bayesian Information Criterion (BIC) (Sch78; KW95; Wea99).

With increasing noise, the acute maximum become more round (blue CH curve), i.e. as expected, BPCR becomes less sure about the correct number of segments. This uncertainty gets pronounced under misspecification (green CMwG curve), and in particular when the true number of segments is far from clear (or nonexistent) like in the genome abberation example (red Gen(3,1) curve). The pink Gen(5,9) curve shows that $\log P(k|\mathbf{y})$ is not necessarily unimodal.

Miscellaneous. Table 2 summarizes the most important quantities of the considered examples.

While using the variance of Δ as estimate for $\hat{\sigma}$ tends to overestimate σ for low noise, the quartile method does not suffer from this (non)problem.

The usefulness of quoting the evidence cannot be overestimated. While the absolute number itself is hard to comprehend, comparisons (based on this absolute(!) number)

are invaluable. Consider, for instance, the three segment medium Gaussian noise data \mathbf{y}_{GM} from Figure 2. Table 2 shows that $\log E(\text{GM}) = -48$, while $\log E(\text{GMwC}) = -70$, i.e. the odds that \mathbf{y}_{GM} has Cauchy rather than Gaussian noise is tiny $e^{48-70} < 10^{-9}$, and similarly the odds that \mathbf{y}_{CM} has Gaussian rather than Cauchy noise is $e^{127-160} < 10^{-14}$. This can be used to decide on the model to use. For instance it clearly indicates that noise in Gene31 and Gen59 is not Cauchy for which log-evidences would be -398 and -406 , respectively. The smallness of the relative log-likelihoods does not indicate any gross misspecification.

The indicated 4th segment for GH and CH is spurious, since it has length zero (two breaks at the same position). In Gene31, only 15 out of the indicated 34 segments are real. The spurious ones would be real had we estimated the breaks $\hat{\mathbf{t}}$ jointly, rather than the marginals t_p separately. They would often be single data segments at the current boundaries, since it costs only a single extra break to cut off an “outlier” at a boundary versus two breaks in the middle of a segment.

In the last column we indicated the confidence $C_{\hat{k}}$ ($C_{\hat{k}-1}, C_{\hat{k}+1}$) of BPCR in the estimate \hat{k} . For clean data (GL,GM,CL,GM) it is certain that there are at least 3 segments. We already explained the general tendency to also believe in higher number of segments.

12 Extensions & Outlook

The core Regression($\mathbf{A}, n, k_{\text{max}}$) algorithm does not care where the in-segment evidence matrix and moments \mathbf{A} come from. This allows for plenty of easy extensions of the basic idea.

Instead of the *scalar mean* μ_q of segment q , we can consider any other scalar, vector, or discrete segment property μ_q . The only relevant property is that, given the segment boundaries, the μ_q of different segments are independent. Similarly, the global parameters may differ from (σ, ν, ρ) . Some of the generalizations described below, like segment dependent variance σ_q , are of this type.

If the segment levels are known to belong to a discrete set (e.g. integer DNA copy numbers (PRLD05)), this simply corresponds to a discrete prior on μ and leads naturally (rather than by need) to a Grid sum as in EstGeneral().

If each segment can have its own (unknown) variance σ_m^2 , we can assume some prior over σ_m and average (15) (which depends on σ_m , notationally suppressed) additionally over σ_m . Possibly $P(\sigma_m|\dots)$ depends on some hyper-parameter that now has to be estimated instead of σ ; all the better if not.

Our recursion can easily be generalized to more general factorizable boundary prior

$$P(\mathbf{t}|k) = T_{t_0 t_1} \cdot \dots \cdot T_{t_{k-1} t_k} / T_{0n}^{0k} \quad \text{with norm} \quad T_{0n}^{0k} := \sum_{0 < t_1 < \dots < t_{k-1} < n} T_{t_0 t_1} \cdot \dots \cdot T_{t_{k-1} t_k} \quad (34)$$

for any $T_{ij} \geq 0$, e.g. $T_{ij} \equiv 1$ and $T_{0n}^{0k} = \binom{n-1}{k-1}$ for our uniform boundary prior (8), or $T_{ij} = P(t_{i+1}=j|t_i=i)$ for renewal processes discussed in Section 13. The only change to the algorithm in Table 1 is to replace $\binom{n-1}{k-1}$ everywhere by T_{0n}^{0k} (computable in time $O(k_{max}n^2)$) and call Regression() with argument $(A_{ij}^r T_{ij})$ instead of (A_{ij}^r) .

We assumed a constant regression function within a segment. Actually any other function could be used. We simply choose likelihood and prior for a single segment and compute its evidence A_{ij}^0 . This is all what Regression() needs to determine the segment number and boundaries. Once we have the segment boundaries it is easy to compute the in-segment quantities we are interested in, e.g. the MAP or mean regression curve.

For instance, if we consider all linear functions within a segment, we get a piecewise linear regression curve. But note that this curve is not continuous. This model is, for instance good, if the true function is essentially piecewise constant, but there is an additional underlying trend (slope) in the segments. Using non-linear functions allows to handle more complicated trends. See e.g. (Fea05).

Piecewise linear (or other) *continuous* regression is more complicated. Assume that μ_p in (11) does not denote the level of the whole segment p , but its level at the right boundary, which together with μ_{p-1} determines the linear function in segment p . Only after fixing μ_p , left and right side decouple. So the recursion analogous to (14) now involves a quantity Q which in addition to (i,j) also depends on (μ_l, μ_m) . This functional recursion may approximately be solved by discretizing $\{(\mu_l, \mu_m) \in \mathbb{R}^2\}$, or by approximating Q by a 2-dimensional Gaussian in (μ_l, μ_m) and storing only the 2 means and the 2×2 covariance matrix for each (i,j) . The following two simpler heuristic approaches may work sufficiently well in practice: One could ignore the continuity constraint when determining the boundaries, and only take them into account in the subsequent (much simpler) regression problem with known boundaries. Another possibility is to consider instead of the continuous piecewise linear function f its piecewise constant derivative f' , i.e. use BPCR on Δ_t and finally integrate the result.

It is also not necessary to use a parametric model for the noise. If different segments can have different noise distributions, we could compute the in-segment evidence, mean, and variance A_{ij}^r based on some (fast) non-parametric model. If all segments have the same distribution, we could non-parametrically estimate a single density for the differences Δ and then deconvolve the density (e.g. by $\text{FT}^{-1}(\sqrt{\text{FT}(\text{density})})$), and henceforth use this as prior for σ in EstGeneral(), where FT is the fourier transform. As non-parametric density estimator we could use the fast (linear-time) exact Bayesian tree model (Hut05b).

Finally, for (very) large n , say > 1000 , the $O(k_{max}n^2)$ algorithm is too slow. Fortunately, there is nearly no interaction between distant segments; boundary t_k is often practically independent of where $t_{k\pm 2}, t_{k\pm 3}$, etc. are placed. This suggests to break the whole data set into smaller overlapping pieces, where each piece should be long enough to contain at least four segments. Then boundaries $t_2^{piece}, \dots, t_{k-2}^{piece}$ of each piece are used, and appropriately merged. For the Bayesian regression curve one should use some blending on the overlap. If single segments are very long, one could coarsen (locally

lump together) the data and later refine around the boundaries. Alternatively one can, of course, resort to simulation based (Monte-Carlo) methods (BH93; Fea06).

13 Comparison to other work

As discussed in the introduction, there are many non-Bayesian approaches to change point detection in general and PC regression in particular, e.g. (SS75; OVLW04; Jon03; Pic05). There are also simulation (Monte-Carlo) approaches like (BH93). In this work we developed an exact Bayesian solution. In the following we compare our model to a different exact Bayesian model developed in (Yao84; BH92; Fea06). The authors consider a renewal process: Given that there is a break at i , the prior probability p_{ij} that the next break is at $j > i$ is independent of where and how many change points occurred before i , that is $P(t_{1l}) = p_{0t_1} p_{t_1 t_2} \dots p_{t_{l-1} t_l}$. They develop recursions in n for the quantities of interest, while our recursion is over k . Since the number of segments in their model is variable without a recursion over k , their algorithm runs in time $O(n^2)$.

Yao (Yao84) considers identically geometrically distributed breaks with $p_{ij} = q(1-q)^{j-i-1}$, i.e. for all i , the probability of a break at i is q , independent of the breaks at other $j \neq i$. In the following we compare Yao's model (P_q) to our model with uniform k -prior (P_u) and our model with fixed segment number (P_k). Table 3 summarizes the most important quantities, derived from the segment prior $P(\mathbf{t}|k)P(k)$ by exploiting various binomial identities.

The first observation is that, given k , the breaks in Yao's model are uniformly distributed as in ours (row 2). Hence his model can only differ from ours in the prior of k . Indeed, his prior is strongly concentrated around $k = qn$ for typical (large) n ($P_q(k)$ in row 1). It more resembles a model with fixed k . Note that in all three models, probabilities involving t_{lm} only depend on t_l and t_m and are independent of t_{l+1}, \dots, t_{m-1} . Yao's model suppresses long segments exponentially ($P_q(\text{row 4}) = q(1-q)^{\text{length}-1}$) with fixed expected length $\mathbf{E}_q[\text{row 5}] = \frac{1}{q} - 1$, while our model has a wide polynomial tail ($P_u(\text{row 4}) \sim l \cdot t_{l-1}^l / t_l^{l+1}$) with expected length $\mathbf{E}_u[\text{row 5}] \sim (t_{l-1} + l) / (l - 1)$ adapting to the past average length and no prior bias ($\mathbf{E}_u[t_1]$ is undefined). Similarly, the probability that there is no break in the range from i to j (row 8) and the probability of a break at j given the previous one is at i (row D) decay exponentially in Yao's model but only harmonically respectively inverse cubic in our model.

Barry and Hartigan (BH92) consider renewal processes with general transition probabilities p_{ij} . We can also use our algorithm for renewal processes by setting $T_{ij} = p_{ij}$ in (34). Interestingly, Barry and Hartigan's favorite model (stated without derivation or motivation) essentially coincides with the long-tail p_{ij} (row A-D) of our model. Nevertheless, their model is completely different from ours, since our model with uniform prior is *not* a renewal process. In our model $p_{ij}^l := P_u(t_l = j | t_{l-1} = i)$ (row 4) depends on l , i.e. $P_u(t_{1l}) = p_{0t_1}^1 p_{t_1 t_2}^2 \dots p_{t_{l-1} t_l}^l \neq p_{0t_1} p_{t_1 t_2} \dots p_{t_{l-1} t_l}$. This l -dependence also prevents us from a recursion algorithm in n . This also explains why the high probability of a break at i ($P(\text{row 9}) = 1/2$) and of segments of length one ($p_{i-1,i} = \frac{4}{1.2.3}$) in both models, imply a high bias towards short segments in Barry's model, but not in ours. This explains

Table 3: **(Quantities of Interest for various models)** Rows (1)–(9) contain various quantities of interest (first column) for our model (P_u , fourth column), our model with fixed k (P_k , third column), and Yao’s model with geometrically distributed segment lengths (P_q , third column). The second column contains the definition or a relation or way to compute the quantity. Section 5 explains notation. Note that conditional on k all three models are the same, hence $P_k(\cdot|k) = P_u(\cdot|k) = P_q(\cdot|k) = P_k(\cdot)$ can be read off from the third column.

Model	Function	definition or relation or way to compute	Fixed	Uniform k	Markov or geometric
			#segments k		
	$P(\dots)$		$P_k(\dots)$	$P_u(\dots)$	$P_q(\dots)$
1)	$P(k)$	$\sum_{0 < t_1 < \dots < t_{k-1} < n, t_k \geq n} P(t_{1k})$	“ δ_{kk} ”	$1/n$	$\binom{n-1}{k-1} q^{k-1} (1-q)^{n-k}$
2)	$P(t_{1,k-1} k)$	$P(t_{1,k-1} \wedge t_k \geq n) / P(k)$	$\binom{n-1}{k-1}^{-1}$	$\binom{n-1}{k-1}^{-1}$	$\binom{n-1}{k-1}^{-1}$
3)	$P(t_{lm})$	$\sum_{k=m+1}^{n-t_m+1} P(t_{lm} k) P(k)$	$\frac{\binom{t_l-1}{l-1} \binom{n-t_m-1}{k-m-1}}{\binom{n-1}{k-1}}$	$\frac{\binom{t_l-1}{l-1}}{(t_m+1) \binom{t_m}{m}}$	$\binom{t_l-1}{l-1} q^m (1-q)^{t_m-m}$
4)	$P(t_l t_{1,l-1})$	$\frac{P(t_{1l})}{P(t_{1,l-1})}$	$\frac{\binom{n-t_l-1}{k-l-1}}{\binom{n-t_{l-1}-1}{k-l}}$	$\frac{t_{l-1}+1}{t_l+1} \frac{\binom{t_{l-1}-1}{l-1}}{\binom{t_l}{l}}$	$q(1-q)^{t_l-t_{l-1}-1}$
5)	$\mathbf{E}[t_l - t_{l-1} t_{1,l-1}]$	$\sum_{t_l=l}^{n-1} (t_l - t_{l-1}) P(t_l t_{1,l-1})$	$\frac{n-t_{l-1}}{k-l+1}$	$\frac{t_{l-1}+l}{l-1} \quad (n = \infty)$	$\frac{1}{q} - 1$
6)	$P(t_l)$	$P(t_{1l})$	$\uparrow P(t_{1l})$	$\frac{l}{(t_l+1)t_l}$	$\uparrow P(t_{1l})$
7)	$P(t_{l-1} < i \wedge t_l \geq j)$	$\sum_{t_{l-1}=l-1}^{i-1} \sum_{t_l=j+1}^{n-1} P(t_{l-1}, t_l)$	$\frac{\binom{i-1}{l-1} \binom{n-j}{k-i}}{\binom{n-1}{k-1}}$	$\frac{1}{l} \frac{\binom{i-1}{l-1}}{\binom{j}{j}}$ $(n = \infty)$	$\binom{i-1}{l-1} q^{l-1} (1-q)^{j-l}$
8)	$P(\forall l: i \not\leq t_l \not\leq j)$	$\sum_{l=1}^i P(t_{l-1} < i \wedge t_l \geq j)$	$\frac{\binom{n-j+i-1}{k-1}}{\binom{n-1}{k-1}}$	$\frac{1}{j-i+1} \quad (n = \infty)$	$(1-q)^{j-i}$
9)	$P(\exists l: t_l = i)$	$\sum_{l=1}^i P(t_l = i)$	$\frac{k-1}{n-1}$	$1/2$	q
A)	p_{0n}	$P(k=1)$	δ_{k1}	$1/n$	$(1-q)^{n-1}$
B)	p_{0j}	$P(t_1 = j)$	$\binom{n-j-1}{k-2} / \binom{n-1}{k-1}$	$\frac{1}{j(j+1)}$	$q(1-q)^{j-1}$
C)	p_{in}	$\frac{\sum_{k=2}^{i+1} P(t_{k-1}=i k) P(k)}{\sum_{l=2}^{i+1} P(t_{l-1}=i)}$	$\binom{i-1}{k-2} / \binom{n-2}{k-2}$	$\frac{2}{(n-i)(n-i+1)}$	$(1-q)^{n-i-1}$
D)	p_{ij}	$\frac{\sum_{l=2}^{i+1} P(t_{l-1}=i \wedge t_l=j)}{\sum_{l=2}^{i+1} P(t_{l-1}=i)}$	$\binom{n+i-j-2}{k-3} / \binom{n-2}{k-2}$	$\frac{4}{(j-i)(j-i+1)(j-i+2)}$	$q(1-q)^{j-i-1}$

why their method overestimates the number of segments.

In general, any renewal process p_{ij} has some expected or median segment length, which leads to a strongly peaked k -prior, in contrast to our uniform prior. If we would perform a Bayesian average over $q \in [0,1]$ with uniform prior, Yao's model would reduce to ours. Both algorithms are $O(n^2)$ with a factor k_{max} in our case and the grid size for integrating over q in Yao's case. Actually Yao (Yao84) and Fearnhead (Fea06) use a maximum likelihood estimate for q .

So the major difference of our work from (Yao84; BH92; Fea06) is that we start with a non-informative uniform prior over k , while they start with a renewal process necessarily corresponding to a highly informed prior for k . In contrast to their recursion over n , our recursion over k allows for arbitrary (in particular uniform) prior over k . This also naturally allows us to find the MAP \hat{k} .

14 Summary

We considered Bayesian regression of piecewise constant functions with unknown segment number, location and level. We derived an efficient algorithm that works for any noise and segment level prior, e.g. Cauchy which can handle outliers. We derived simple but good estimates for the in-segment variance. We also proposed a Bayesian regression curve as a better way of smoothing data without blurring boundaries. The Bayesian approach also allowed us to straightforwardly determine the global evidence, break probabilities and error estimates, useful for model selection and significance and robustness studies. We discussed the performance on synthetic and real-world examples. Many possible extensions have been discussed.

Acknowledgements. Thanks to IOSI for providing the gene copy # data and to Ivo Kwee for discussions.

References

- [BH92] D. Barry and J. A. Hartigan. Product partition models for change point problems. *Annals of Statistics*, 20:260–279, 1992. 636, 660, 662
- [BH93] D. Barry and J. A. Hartigan. A Bayesian analysis for change point problems. *Journal of the American Statistical Association*, 88:309–319, 1993. 660
- [Bol04] W. M. Bolstad. *Introduction to Bayesian Statistics*. Wiley Interscience, New Jersey, 2004. 644
- [EF05] D. Endres and P. Földiák. Bayesian bin distribution inference and mutual information. *IEEE Transactions on Information Theory*, 51(11):3766–3779, 2005. 636
- [Fea05] P. Fearnhead. Exact Bayesian curve fitting and signal segmentation. *IEEE Transactions on Signal Processing*, 53:2160–2166, 2005. 636, 659
- [Fea06] P. Fearnhead. Exact and efficient Bayesian inference for multiple changepoint problems. *Statistics and Computing*, 16:203–213, 2006. 636, 642, 660, 662

- [Goo83] I. J. Good. Explicativity, corroboration, and the relative odds of hypotheses. In *Good thinking: The Foundations of Probability and its applications*. University of Minnesota Press, Minneapolis, MN, 1983. 656
- [Hut05a] M. Hutter. Additional material to article. <http://www.idsia.ch/~marcus/ai/pcreg.htm>, 2005. 648
- [Hut05b] M. Hutter. Fast non-parametric Bayesian inference on infinite trees. In *Proc. 10th International Conf. on Artificial Intelligence and Statistics (AISTATS-2005)*, pages 144–151. Society for Artificial Intelligence and Statistics, 2005. 659
- [Hut07] M. Hutter. Bayesian regression of piecewise constant functions. In J.M. Bernardo et al., editors, *Proc. Bayesian Statistics*, volume 8, Benidorm, 2007. Oxford University Press. 635
- [Jay03] E. T. Jaynes. *Probability Theory: The Logic of Science*. Cambridge University Press, Cambridge, MA, 2003. 656
- [Jon03] K. Jong et al. Chromosomal breakpoint detection in human cancer. In *Applications of Evolutionary Computing: EvoWorkshops'03*, volume 2611 of *LNCS*, pages 54–65. Springer, 2003. 636, 660
- [KW95] R. E. Kass and L. Wasserman. A reference Bayesian test for nested hypotheses with large samples. *Journal of the ACM*, 90:773–795, 1995. 636, 657
- [Mac03] D. J. C. MacKay. *Information theory, inference and learning algorithms*. Cambridge University Press, Cambridge, MA, 2003. 656
- [OVLW04] A. B. Olshen, E. S. Venkatraman, R. Lucito, and M. Wigler. Circular binary segmentation for the analysis of array-based DNA copy number data. *Biostatistics*, 5:557–572, 2004. 636, 660
- [Pic05] F. Picard et al. A statistical approach for array CGH data analysis. *BMC Bioinformatics*, 6(27):1–14, 2005. 636, 660
- [Pin98] D. Pinkel et al. High resolution analysis of DNA copy number variation using comparative genomic hybridization to microarrays. *Nature Genetics*, 20:207–211, 1998. 656
- [PRLD05] F. Picard, S. Robin, E. Lebarbier, and J. J. Daudin. A segmentation-clustering problem for the analysis of array CGH data. In *Proc. 11th International Symposium on Applied Stochastic Models and Data Analysis (ASMDA'05)*, pages 145–152, Brest, France, 2005. 658
- [Rin06] A. Rinaldi et al. Genomic and expression profiling identifies the B-cell associated tyrosine kinase Syk as a possible therapeutic target in mantle cell lymphoma. *British Journal of Haematology*, 132(3):303–316, 2006. 656
- [Sch78] G. Schwarz. Estimating the dimension of a model. *Annals of Statistics*, 6(2):461–464, 1978. 636, 657
- [SS75] A. Sen and M. S. Srivastava. On tests for detecting a change in mean. *Annals of Statistics*, 3:98–108, 1975. 636, 660

- [Wea99] D. L. Weakliem. A critique of the Bayesian information criterion for model selection. *Sociological Methods and Research*, 27:359–397, 1999. [636](#), [657](#)
- [Yao84] Y.-C. Yao. Estimation of a noisy discrete-time step function: Bayes and empirical Bayes approaches. *Annals of Statistics*, 12:1434–1447, 1984. [636](#), [660](#), [662](#)