Bayesian selection for the regularization parameter in TVℓ_0 denoising problems

Jordan Frecon,* Nelly Pustelnik*, Nicolas Dobigeon†,
Herwig Wendt†, and Patrice Abry*

Abstract

Piecewise constant denoising can be solved either by deterministic optimization approaches, based on total variation (TV), or by stochastic Bayesian procedures. The former lead to low computational time but requires the selection of a regularization parameter, whose value significantly impacts the achieved solution, and whose automated selection remains an involved and challenging problem. Conversely, fully Bayesian formalisms encapsulate the regularization parameter selection into hierarchical models, at the price of large computational costs. This contribution proposes an operational strategy that combines hierarchical Bayesian and TVℓ_0 formulations, with the double aim of automatically tuning the regularization parameter and of maintaining computational efficiency. The proposed procedure relies on formally connecting a Bayesian framework to a TVℓ_0 minimization formulation. Behaviors and performance for the proposed piecewise constant denoising and regularization parameter tuning techniques are studied qualitatively and assessed quantitatively, and shown to compare favorably against those of a fully Bayesian hierarchical procedure, both in accuracy and in computational load.

*J. Frecon (Corresponding author), N. Pustelnik and P. Abry are with Univ Lyon, Ens de Lyon, Univ Claude Bernard, CNRS, Laboratoire de Physique, F-69342 Lyon, France (email: firstname.lastname@ens-lyon.fr).

†N. Dobigeon and H. Wendt are with IRIT/INP-ENSEEIHT, University of Toulouse, France (email: firstname.lastname@irit.fr).

1 Introduction

Piecewise constant denoising. Piecewise constant denoising (tightly related to change-point detection) is of considerable potential interest in numerous signal processing applications including, e.g, econometrics and biomedical analysis (see [2, 20], for an overview). An archetypal and most encountered in the literature formulation considers noisy observations as resulting from the additive mixture of a piecewise constant signal $\overline{\boldsymbol{x}} \in \mathbb{R}^N$ with a Gaussian noise $\boldsymbol{\epsilon} \in \mathcal{N}(0, \sigma^2 \boldsymbol{I}_N)$

$$\boldsymbol{y} = \overline{\boldsymbol{x}} + \boldsymbol{\epsilon}. \tag{1}$$

Detecting change-points or denoising the piecewise constant information has been addressed by several strategies, such as Cusum procedures [2], hierarchical Bayesian inference frameworks [13, 19], or functional optimization formulations, involving either the ℓ_1-norm or the ℓ_0-pseudo-norm of the first differences of $\boldsymbol{x}$ [17, 21, 28] (i.e, TVℓ_1 or TVℓ_0, respectively).

This later class frames the present contribution. Formally, it amounts to recovering the solution of an TVℓ_0-penalized least-square problem, namely,

$$\widehat{\boldsymbol{x}}_\lambda = \arg \min_{\boldsymbol{x} \in \mathbb{R}^N} \frac{1}{2} \|\boldsymbol{y} - \boldsymbol{x}\|_2^2 + \lambda \|\boldsymbol{L}\boldsymbol{x}\|_0, \tag{2}$$

where $\boldsymbol{L} \in \mathbb{R}^{(N-1) \times N}$ models the first difference operator, i.e., $\boldsymbol{L}\boldsymbol{x} = (x_{i+1} - x_i)_{1 \leq i \leq N-1}$, the ℓ_0-pseudo-norm counts the non-zeros elements in $\boldsymbol{L}\boldsymbol{x}$, and $\lambda > 0$ denotes the regularization parameter which adjusts the respective contributions of the data-fitting and penalization terms in the overall objective function. Note that the TVℓ_0 problem (2) is often encountered as the univariate Potts model [30].

Such a formulation however suffers from a major limitation: its actual solution depends drastically on the regularization parameter λ. The challenging question of automatically estimating λ from data constitutes the core issue addressed in the present contribution.

Related works. Bayesian hierarchical inference frameworks have received considerable interests for addressing change-point or piecewise denoising problems [13, 19]. This mostly results from their ability to include the hyperparameters within the Bayesian modeling and to jointly estimate them with the parameters of interest. In return for this intrinsic flexibility, approximating the Bayesian estimators associated with this hierarchical model generally requires the use of Markov chain Monte Carlo (MCMC) algorithms, which is often known as excessively demanding in terms of computational burden and times.

Remaining in the class of deterministic functional minimization, the non-convexity of the objective function underlying (2) has sometimes been alleviated by a convex relaxation, i.e., the ℓ_0-pseudo-norm is replaced with ℓ_1-norm

$$\widehat{\boldsymbol{x}}_\tau = \arg \min_{\boldsymbol{x} \in \mathbb{R}^N} \frac{1}{2} \|\boldsymbol{y} - \boldsymbol{x}\|_2^2 + \tau \|\boldsymbol{L}\boldsymbol{x}\|_1. \tag{3}$$

In essence, such an approach preserves the same intuition of piecewise constant denoising, with the noticeable advantage of ensuring convexity of the resulting function to be minimized, which ensures the convergence of the minimization algorithms [1, 3, 6, 32] or straightforward computations [7, 9]. In particular, the formulation proposed in (3) has received considerable interest, because, besides the existence and performance of sound algorithmic resolution procedures, it can offer some convenient ways to handle the automated tuning of the regularization parameter $\tau > 0$. For instance, the *Stein unbiased risk estimate*

(SURE) [10, 29] aims at producing an unbiased estimator that minimizes the mean squared error between $\bar{\boldsymbol{x}}$ and $\widehat{\boldsymbol{x}}_\tau$. While practically effective, implementing SURE requires the prior knowledge of the variance σ^2 of the residual error $\boldsymbol{\epsilon}$, often unavailable a priori (see, a contrario, [4, 26] for hyperspectral denoising or image deconvolution involving frames where σ^2 has been estimated).

Alternatively, again, the problem in (3) can be tackled within a fully Bayesian framework, relying on the formulation of (3) as a statistical inference problem. Indeed, in the right-hand side of (3), the first term can be straightforwardly associated with a negative log-likelihood function by assuming an additive white Gaussian noise sequence $\boldsymbol{\epsilon}$, i.e., $\boldsymbol{y}|\boldsymbol{x}$ is distributed according to the Gaussian distribution $\mathcal{N}(\boldsymbol{x}, \sigma^2 \boldsymbol{I}_N)$. Further, the second term refers to a Laplace prior distribution for the first difference $\boldsymbol{Lx}$ of the unobserved signal. Under such Bayesian modeling, the corresponding maximum a posteriori (MAP) criterion reads

$$\underset{\boldsymbol{x} \in \mathbb{R}^N}{\text{maximize}} \left\{ \left(\frac{1}{2\pi\sigma^2} \right)^{N/2} e^{-\frac{1}{2\sigma^2} \|\boldsymbol{y} - \boldsymbol{x}\|_2^2} \frac{1}{Z(\tau/\sigma^2)} e^{-\frac{\tau}{\sigma^2} \|\boldsymbol{Lx}\|_1} \right\} \tag{4}$$

whose resolution leads to the solution (3) and where $Z(\tau/\sigma^2)$ is the normalizing constant associated with the prior distribution. Following a hierarchical strategy, the hyperparameters τ and σ^2 could be included into the Bayesian model to be jointly estimated with $\boldsymbol{x}$. However, in the specific case of (3), the prior distribution related to the penalization is not separable with respect to (w.r.t.) the individual components of $\boldsymbol{x}$: The partition function $Z(\tau/\sigma^2)$ can hence not be expressed analytically. Therefore, estimating τ within a hierarchical Bayesian framework would require either to choose a heuristic prior for τ as proposed in [5, 23, 24] or to conduct intensive approximate Bayesian computations as in [25].

Goals, contributions and outline. Elaborating on the Bayesian interpretation of the TVℓ_1 denoising problem stated in (3), the present contribution chooses to stick to the original TVℓ_0 formulation (2). Capitalizing on efficient dynamic programming algorithms [14, 30, 31, 33, 34] which allow $\widehat{\boldsymbol{x}}_\lambda$ to be recovered for a predefined value of λ, the main objective of this work resides in the joint estimation of the denoised signal and the *optimal* hyperparameter λ, without assuming any additional prior knowledge regarding the residual variance σ^2. Formally, this problem can be formulated as an extended counterpart of (2), i.e., a minimization procedure involving $\boldsymbol{x}$, λ, and σ^2 as stated in what follows.

Problem 1.1 *Let $\boldsymbol{y} \in \mathbb{R}^N$ and $\phi\colon \mathbb{R}_+ \times \mathbb{R}_+ \to \mathbb{R}$. We aim to*

$$\underset{\boldsymbol{x} \in \mathbb{R}^N, \lambda > 0, \sigma^2 > 0}{\text{minimize}} \frac{1}{2\sigma^2} \|\boldsymbol{y} - \boldsymbol{x}\|_2^2 + \frac{\lambda}{\sigma^2} \|\boldsymbol{Lx}\|_0 + \phi(\lambda, \sigma^2). \tag{5}$$

The main challenge for handling Problem 1.1 lies in the design of an appropriate function ϕ that leads to a relevant penalization of the overall criterion w.r.t. the set of nuisance parameters (λ, σ^2). To that end, Section 2 provides a natural parametrization of $\boldsymbol{x}$ and a reformulation of Problem 1.1. In Section 3, a closed-form expression of ϕ and an interpretation of λ will be derived from an usual yet sound hierarchical Bayesian inference framework. For this particular choice of the function ϕ, Section 4 proposes an efficient algorithmic strategy to approximate a solution of the Problem 1.1. In Section 5, the relevance and performance of this procedure are qualitatively illustrated and quantitatively assessed, and shown to compare favorably against the Markov chain Monte Carlo algorithm resulting from the hierarchical Bayesian counterpart of (5), both in terms of accuracy and in computational load.

2 Problem parametrization

Following [11–13, 18, 19], piecewise constant signals $\boldsymbol{x} \in \mathbb{R}^N$ can be explicitly parametrized via change-point locations $\boldsymbol{r}$ and amplitudes of piecewise constant segments $\boldsymbol{\mu}$. These reparametrizations are derived in Sections 2.1 and 2.2. They are in turn used in Section 2.3 to bring Problem 1.1 in a form more amenable for explicit connection to a hierarchical Bayesian model.

2.1 Change-point location parametrization r

To locate the time instants of the change-points in the denoised signal $\boldsymbol{x}$, an indicator vector $\boldsymbol{r} = \left(r_i\right)_{1 \leq i \leq N} \in \{0,1\}^N$ is introduced as follows

$$r_i = \begin{cases} 1, & \text{if there is a change-point at time instant } i, \\ 0, & \text{otherwise.} \end{cases} \tag{6}$$

By convention, $r_i = 1$ indicates that x_i is the last sample belonging to the current segment, and thus that x_{i+1} belongs to the next segment. Moreover, stating $r_N = 1$ ensures that the number K of segments is equal to the number of change-points, i.e., $K = \sum_{i=1}^{N} r_i$.

For each segment index $k \in \{1, \ldots, K\}$, the set $\mathcal{R}_k \subset \{1, \ldots, N\}$ is used to denote the set of time indices associated with the k-th segment. In particular, it is worthy to note that $\mathcal{R}_k \cap \mathcal{R}_{k'} = \{\emptyset\}$ for $k \neq k'$ and $\cup_{k=1}^{K} \mathcal{R}_k = \{1, \cdots, N\}$. Hereafter, the notation $K_{\boldsymbol{r}}$ will be adopted to emphasize the dependence of the number K of segments on the indicator vector $\boldsymbol{r}$, i.e., $K = \|\boldsymbol{r}\|_0$.

2.2 Segment amplitude parametrization μ

The amplitudes of each segment of the piecewise constant signal can be encoded by introducing the vector $\boldsymbol{\mu} = (\mu_k)_{1 \leq k \leq K_{\boldsymbol{r}}}$ such that

$$(\forall k \in \{1, \ldots, K_{\boldsymbol{r}}\})(\forall i \in \mathcal{R}_k) \quad x_i = \mu_k. \tag{7}$$

2.3 Reformulation of Problem 1.1

In place of $\boldsymbol{x}$, the parameter vector $\boldsymbol{\theta} = \{\boldsymbol{r}, \boldsymbol{\mu}\}$ will now be used to fully specify the piecewise constant signal $\boldsymbol{x}$. An important issue intrinsic to the TVℓ_0 denoising problem and thus to this formulation stems from the fact that the unknown parameter $\boldsymbol{\theta}$ belongs to the space $\{0,1\}^N \times \mathbb{R}^{K_{\boldsymbol{r}}}$ whose dimension is a priori unknown, as it depends on the number $K_{\boldsymbol{r}}$ of change-points. Moreover, this parametrization leads to the following lemma.

Lemma 2.1 *Let $\boldsymbol{y} \in \mathbb{R}^N$ and $\phi \colon \mathbb{R}_+ \times \mathbb{R}_+ \to \mathbb{R}$. Problem 1.1 is equivalent to*

$$\underset{\substack{\boldsymbol{\theta} = \{\boldsymbol{r}, \boldsymbol{\mu}\} \in \{0,1\}^N \times \mathbb{R}^{K_{\boldsymbol{r}}} \\ \lambda > 0, \sigma^2 > 0}}{\text{minimize}} \left\{ \frac{1}{2\sigma^2} \sum_{k=1}^{K_{\boldsymbol{r}}} \sum_{i \in \mathcal{R}_k} (y_i - \mu_k)^2 \right.$$

$$\left. + \frac{\lambda}{\sigma^2}(K_{\boldsymbol{r}} - 1) + \phi(\lambda, \sigma^2) \right\} \tag{8}$$

where $(\mathcal{R}_k)_{1 \leq k \leq K_{\boldsymbol{r}}}$ is related to $\boldsymbol{r}$ as indicated in Section 2.1.

Indeed, the data fidelity term in the minimization Problem 1.1 can be equivalently written as

$$\|\boldsymbol{y} - \boldsymbol{x}\|^2 = \sum_{k=1}^{K_{\boldsymbol{r}}} \sum_{i \in \mathcal{R}_k} (y_i - x_i)^2 = \sum_{k=1}^{K_{\boldsymbol{r}}} \sum_{i \in \mathcal{R}_k} (y_i - \mu_k)^2. \tag{9}$$

Moreover, the $\text{TV}\ell_0$ penalization can be rewritten as

$$\|\boldsymbol{L}\boldsymbol{x}\|_0 = \|\boldsymbol{r}\|_0 - 1 = K_{\boldsymbol{r}} - 1. \tag{10}$$

Lemma 2.1 implies that estimating the piecewise constant signal $\boldsymbol{x}$ can be equivalently formulated as estimating the parameter vector $\boldsymbol{\theta}$.

3 Bayesian derivation of ϕ

Assisted by the reformulation of Problem 1.1 and a hierarchical Bayesian framework detailed in Section 3.1, a relevant penalization function ϕ will be derived in Section 3.2.

3.1 Hierarchical Bayesian model

In [13, 19], the problem of detecting change-points in a stationary sequence has been addressed following a Bayesian inference procedure which aims at deriving the posterior distribution of the parameter vector $\boldsymbol{\theta}$ from the likelihood function associated with the observation model and the prior distributions chosen for the unknown parameters. In what follows, a similar approach is proposed to produce a hierarchical Bayesian model that can be tightly related to the Problem 1.1 under a joint MAP paradigm.

First, the noise samples $(\boldsymbol{\epsilon}_i)_{1 \leq i \leq N}$ are assumed to be independent and identically distributed (i.i.d.) zero mean Gaussian variables with common but unknown variance σ^2, i.e., $\boldsymbol{\epsilon}|\sigma^2 \sim \mathcal{N}\left(\boldsymbol{0}, \sigma^2 \boldsymbol{I}_N\right)$. The resulting joint likelihood function of the observations $\boldsymbol{y}$ given the piecewise constant model $\{\boldsymbol{r}, \boldsymbol{\mu}\}$ and the noise variance σ^2 reads

$$f\left(\boldsymbol{y}|\boldsymbol{r}, \boldsymbol{\mu}, \sigma^2\right) = \prod_{k=1}^{K_{\boldsymbol{r}}} \prod_{i \in \mathcal{R}_k} \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(\mu_k - y_i)^2}{2\sigma^2}\right). \tag{11}$$

Then to derive the posterior distribution, prior distributions are elected for the parameters $\boldsymbol{r}$ and $\boldsymbol{\mu}$, assumed to be a priori independent. Following well-admitted choices such as those in [12, 13, 19, 27], the components r_i of the indicator vector $\boldsymbol{r}$ are assumed to be a priori independent and identically distributed (i.i.d.) according to a Bernoulli distribution with hyperparameter p, quantifying the prior probability of occurrence of a change, independently of the location,

$$\begin{aligned} f(\boldsymbol{r}|p) &= \prod_{i=1}^{N-1} p^{r_i}(1-p)^{1-r_i} \\ &= p^{\sum_{i=1}^{N-1} r_i}(1-p)^{(N-1-\sum_{i=1}^{N-1} r_i)} \\ &= \left(\frac{p}{1-p}\right)^{(K_{\boldsymbol{r}}-1)}(1-p)^{(N-1)}. \end{aligned} \tag{12}$$

From a Bayesian perspective, a natural choice for $f(\boldsymbol{\mu}|\boldsymbol{r})$ consists in electing independent conjugate Gaussian prior distributions $\mathcal{N}\left(\mu_0, \sigma_0^2\right)$ for the segment amplitudes μ_k ($k = 1, \ldots, K_{\boldsymbol{r}}$), i.e.,

$$f(\boldsymbol{\mu}|\boldsymbol{r}) = \prod_{k=1}^{K_{\boldsymbol{r}}} \frac{1}{\sqrt{2\pi\sigma_0^2}} e^{\frac{-(\mu_k - \mu_0)^2}{2\sigma_0^2}}. \tag{13}$$

Indeed, this set of conjugate priors ensures that the conditional posterior distributions of the segment amplitudes are still Gaussian distributions.

Moreover, within a hierarchical Bayesian paradigm, nuisance parameters, such as the noise variance, and other hyperparameters defining the prior distributions can be included within the model to be estimated jointly with $\boldsymbol{\theta}$ [13, 19]. In particular, to account for the absence of prior knowledge on the noise variance σ^2, a scale-invariant non-informative Jeffreys prior is assigned to σ^2

$$f\left(\sigma^2\right) \propto \frac{1}{\sigma^2}. \tag{14}$$

Finally, a conjugate Beta distribution $\mathcal{B}(\alpha_0, \alpha_1)$ is assigned to the unknown hyperparameter p, which is a natural choice to model a $(0, 1)$-constrained parameter

$$f(p) = \frac{\Gamma\left(\alpha_0 + \alpha_1\right)}{\Gamma\left(\alpha_0\right)\Gamma\left(\alpha_1\right)} p^{\alpha_1 - 1}(1 - p)^{\alpha_0 - 1}. \tag{15}$$

3.2 Joint MAP criterion

From the likelihood function (11) and prior and hyper-prior distributions (12)–(15) introduced above, the joint posterior distribution reads

$$f(\boldsymbol{\Theta}|\boldsymbol{y}) \propto f\left(\boldsymbol{y}|\boldsymbol{r}, \boldsymbol{\mu}, \sigma^2\right) f\left(\boldsymbol{\mu}|\boldsymbol{r}\right) f(\boldsymbol{r}|p) f(p) f\left(\sigma^2\right) \tag{16}$$

with $\boldsymbol{\Theta} = \left\{\boldsymbol{r}, \boldsymbol{\mu}, \sigma^2, p\right\}$. Deriving the Bayesian estimators, such as the minimum mean square error (MMSE) and MAP estimators associated with this posterior distribution is not straightforward, mainly due to the intrinsic combinatorial problem resulting from the dimension-varying parameter space $\{0, 1\}^N \times \mathbb{R}^{K_{\boldsymbol{r}}}$. In particular, a MAP approach would consist in maximizing the joint posterior distribution (16), which can be reformulated as the following minimization problem by taking the negative logarithm of (16).

Problem 3.1 *Let $\boldsymbol{y} = (y_i)_{1 \le i \le N} \in \mathbb{R}^N$ and let $\boldsymbol{\Phi} = \left\{\alpha_0, \alpha_1, \sigma_0^2\right\}$ a set of hyperparameters. We aim to*

$$\begin{aligned}
\underset{\boldsymbol{\Theta} = \{\boldsymbol{r}, \boldsymbol{\mu}, \sigma^2, p\}}{\text{minimize}} \quad & \frac{1}{2\sigma^2} \sum_{k=1}^{K_{\boldsymbol{r}}} \sum_{i \in \mathcal{R}_k} (y_i - \mu_k)^2 \\
& + (K_{\boldsymbol{r}} - 1)\left(\log\left(\frac{1 - p}{p}\right) + \frac{1}{2}\log(2\pi\sigma_0^2)\right) \\
& + \frac{N}{2}\log(2\pi\sigma^2) - (N - 1)\log(1 - p) + \log\sigma^2 \\
& - (\alpha_1 - 1)\log p - (\alpha_0 - 1)\log(1 - p) \\
& + \frac{1}{2\sigma_0^2} \sum_{k=1}^{K_{\boldsymbol{r}}} (\mu_k - \mu_0)^2 + \frac{1}{2}\log(2\pi\sigma_0^2).
\end{aligned} \tag{17}$$

Despite apparent differences in parametrization between Problem 1.1 and Problem 3.1, we prove hereafter that both are equivalent for specific choices of λ and ϕ.

Proposition 3.1 *For σ_0^2 large enough, Problem 1.1 with*

$$\lambda = \sigma^2 \left(\log\left(\frac{1-p}{p}\right) + \frac{1}{2}\log(2\pi\sigma_0^2) \right) \tag{18}$$

and

$$
\begin{aligned}
\phi(\lambda, \sigma^2) = {} & \frac{N}{2}\log(2\pi\sigma^2) + \log(\sigma^2) \\
& - \frac{\lambda}{\sigma^2}(N + \alpha_0 - 2) + \frac{N + \alpha_0 - 1}{2}\log(2\pi\sigma_0^2) \\
& + (N + \alpha_0 + \alpha_1 - 3)\log\left(1 + \exp\left(\frac{\lambda}{\sigma^2} - \frac{1}{2}\log(2\pi\sigma_0^2)\right)\right)
\end{aligned}
\tag{19}
$$

matches Problem 3.1.

The sketch of the proof consists in identifying three terms in the criterion (17): the data fidelity term (9), a term proportional to the TVℓ_0 regularization (10), and a third one independent of the indicator vector $\boldsymbol{r}$. We observe that this is almost the case under the condition that we could neglect the term $\frac{1}{2\sigma_0^2}\sum_{k=1}^{K_r}(\mu_k - \mu_0)^2$ which explicitly depends on $\boldsymbol{r}$ through $K_{\boldsymbol{r}}$. Thus, if we choose σ_0^2 sufficiently large such that

$$\frac{1}{2\sigma_0^2}\sum_{k=1}^{K_r}(\mu_k - \mu_0)^2 \ll \frac{1}{2}\log(2\pi\sigma_0^2), \tag{20}$$

then the following equalities can be stated

$$\lambda = \sigma^2\left(\log\left(\frac{1-p}{p}\right) + \frac{1}{2}\log(2\pi\sigma_0^2)\right), \tag{21}$$

and

$$
\begin{aligned}
\phi(p, \sigma^2) = {} & \frac{N}{2}\log(2\pi\sigma^2) - (N-1)\log(1-p) + \log\sigma^2 \\
& - (\alpha_1 - 1)\log p - (\alpha_0 - 1)\log(1-p) \\
& + \frac{1}{2}\log(2\pi\sigma_0^2).
\end{aligned}
\tag{22}
$$

Therefore, it follows that p can be parametrized as a function of $\{\lambda, \sigma^2\}$:

$$p = \left(1 + \exp\left(\frac{\lambda}{\sigma^2} - \frac{1}{2}\log(2\pi\sigma_0^2)\right)\right)^{-1}, \tag{23}$$

which permits to reformulate ϕ as a function of λ and σ^2 such as in (19).

Remark 1 *A sufficient condition to guarantee (20) reads*

$$\sigma_0^2 \gg N(\max \boldsymbol{y} - \min \boldsymbol{y})^2 + \frac{\pi}{2}. \tag{24}$$

Indeed, on one side, since the segments amplitudes follow (13), we can upper bound $(\mu_k - \mu_0)^2$ by $(\max \boldsymbol{y} - \min \boldsymbol{y})^2$ and thus

$$\frac{1}{2\sigma_0^2} \sum_{k-1}^{K_r} (\mu_k - \mu_0)^2 \leq \frac{K_r}{2\sigma_0^2} (\max \boldsymbol{y} - \min \boldsymbol{y})^2 \tag{25}$$

$$\leq \frac{N}{2\sigma_0^2} (\max \boldsymbol{y} - \min \boldsymbol{y})^2. \tag{26}$$

On the other side, we can make use of the lower bound $\frac{1}{2} \log 2\pi\sigma_0^2 \geq \frac{1}{2}(1 - \frac{1}{2\pi\sigma_0^2})$. By combining the two, it yields that (24) implies (20).

Remark 2 *Interesting behaviors can be noticed from the parametrization* (18) *of the regularization parameter* λ *in terms of noise variance* σ^2 *and of prior probability* p *of having a change-point. In particular, as expected,* λ *varies with* σ^2, *which classically means that more regularization should be considered for noisier signals. Conversely,* λ *is a decreasing function of* p *which, as we will illustrate, indicates that* λ *should be selected as more and more finely as* p *grows.*

4 Algorithmic solution

Thanks to Proposition 3.1, a function ϕ has been derived which allows the choice of the regularization parameter λ to be penalized in Problem 1.1. In this section, an algorithmic solution is proposed to approximate $\{\widehat{\boldsymbol{x}}, \widehat{\lambda}, \widehat{\sigma}^2\}$, a solution of Problem 1.1.

An alternate minimization over $\boldsymbol{x}$, λ and σ^2 would not be efficient due to the non-convexity of the underlying criterion. To partly alleviate this problem, we propose to estimate λ on a grid Λ. Therefore, a candidate solution can be obtained by solving ($\forall \lambda \in \Lambda$)

$$(\widehat{\boldsymbol{x}}_\lambda, \widehat{\sigma}_\lambda^2) \in \underset{\boldsymbol{x} \in \mathbb{R}^N, \sigma^2 > 0}{\mathrm{Argmin}} \ \underbrace{\frac{1}{2\sigma^2} \|\boldsymbol{y} - \boldsymbol{x}\|_2^2 + \frac{\lambda}{\sigma^2} \|L\boldsymbol{x}\|_0 + \phi(\lambda, \sigma^2)}_{F(\boldsymbol{x}, \lambda, \sigma^2)}, \tag{27}$$

which we propose to approximate by

$$(\forall \lambda \in \Lambda) \quad \begin{cases} \widehat{\boldsymbol{x}}_\lambda &= \arg \min_{\boldsymbol{x} \in \mathbb{R}^N} \ \frac{1}{2} \|\boldsymbol{y} - \boldsymbol{x}\|_2^2 + \lambda \|L\boldsymbol{x}\|_0, \\ \widehat{\sigma}_\lambda^2 &= \frac{\|\boldsymbol{y} - \widehat{\boldsymbol{x}}_\lambda\|^2}{N-1}, \end{cases} \tag{28}$$

and then selecting the triplet $\{\widehat{\boldsymbol{x}}_{\widehat{\lambda}}, \widehat{\lambda}, \widehat{\sigma}_{\widehat{\lambda}}^2\}$ such that

$$\widehat{\lambda} = \arg \min_{\lambda \in \Lambda} F(\widehat{\boldsymbol{x}}_\lambda, \lambda, \widehat{\sigma}_\lambda^2). \tag{29}$$

Note that the approximation used in (28) amounts in using the solution of (2) for different $\lambda \in \Lambda$ to probe the space $\{\boldsymbol{x}, \sigma^2\} \in \mathbb{R}^N \times \mathbb{R}_+$. Therefore, the iterations of the proposed full algorithmic scheme (reported in Algo. 1) are very succinct and the overall algorithm complexity mainly depends on the ability to solve (2) efficiently for any $\lambda \in \Lambda$. In this work, we propose to resort to *Pottslab*, a dynamic programming algorithm developed in [30].

Algorithm 1 Bayesian driven resolution of the TVℓ_0 problem

Input: Observed signal $\boldsymbol{y} \in \mathbb{R}^N$.

   The predefined set of regularization parameters Λ.

   Hyperparameters $\boldsymbol{\Phi} = \left\{ \alpha_0, \alpha_1, \sigma_0^2 \right\}$.

Iterations:

 1: **for** $\lambda \in \Lambda$ **do**

 2:    Compute $\widehat{\boldsymbol{x}}_\lambda = \arg \min_{\boldsymbol{x} \in \mathbb{R}^N} \frac{1}{2} \|\boldsymbol{y} - \boldsymbol{x}\|_2^2 + \lambda \|\boldsymbol{L}\boldsymbol{x}\|_0$.

 3:    Compute $\widehat{\sigma}_\lambda^2 = \|\boldsymbol{y} - \widehat{\boldsymbol{x}}_\lambda\|^2 / (N-1)$.

 4: **end for**

Output: Solution $\{\widehat{\boldsymbol{x}}_{\widehat{\lambda}}, \widehat{\lambda}, \widehat{\sigma}_{\widehat{\lambda}}^2\}$ with $\widehat{\lambda} = \arg \min_{\lambda \in \Lambda} F(\widehat{x}_\lambda, \lambda, \widehat{\sigma}_\lambda^2)$

---

5 Automated selection of λ: Illustration and performance

5.1 Performance evaluation and hyperparameter settings

5.1.1 Synthetic data

The performance of the proposed automated selection of λ are illustrated and assessed using Monte Carlo numerical simulation based on synthetic data $\boldsymbol{y} = \overline{\boldsymbol{x}} + \boldsymbol{\epsilon}$, where the noise $\boldsymbol{\epsilon}$ consists of i.i.d. samples drawn from $\boldsymbol{\epsilon} \sim \mathcal{N}(0, \sigma^2 \boldsymbol{I}_N)$, and signal $\overline{\boldsymbol{x}}$ is piece-wise constant, with i.i.d. change-points, occurring with probability p, and i.i.d. amplitudes drawn from a uniform distribution (between $\overline{x}_{\min}$ and $\overline{x}_{\max}$).

5.1.2 Performance quantification

Performance are quantified by relative mean square error (MSE) and Jaccard error. While the former evaluates performance in the overall (shape and amplitude) estimation $\widehat{\boldsymbol{x}}$ of $\overline{\boldsymbol{x}}$, the latter focuses on the accuracy of change-point location estimation $\boldsymbol{r}$. The Jaccard error between the true change-point vector $\overline{\boldsymbol{r}}$ and its estimate $\widehat{\boldsymbol{r}}$ (both in $\{0,1\}^N$), is defined as [15, 16]

$$J(\overline{\boldsymbol{r}}, \widehat{\boldsymbol{r}}) = 1 - \frac{\sum_{i=1}^N \min(\overline{r}_i, \widehat{r}_i)}{\sum_{\substack{1 \leq i \leq N \\ \overline{r}_i > 0, \widehat{r}_i > 0}} \frac{\overline{r}_i + \widehat{r}_i}{2} + \sum_{\substack{1 \leq i \leq N \\ \widehat{r}_i = 0}} \overline{r}_i + \sum_{\substack{1 \leq i \leq N \\ \overline{r}_i = 0}} \widehat{r}_i}. \tag{30}$$

$J(\overline{\boldsymbol{r}}, \widehat{\boldsymbol{r}})$ ranges from 0 when $\overline{\boldsymbol{r}} = \widehat{\boldsymbol{r}}$, to 1, when $\overline{\boldsymbol{r}} \cap \widehat{\boldsymbol{r}} = \emptyset$. The Jaccard error is a demanding measure: when one half of non-zero values of two given binary sequences coincides while the other half does not, then $J(\overline{\boldsymbol{r}}, \widehat{\boldsymbol{r}}) = 2/3$.

In the present study, to account for the fact that a solution with a change point position mismatch by a few time indices remains useful and of practical interest, the Jaccard error is computed between smoothed versions $\overline{\boldsymbol{r}} * \mathcal{G}$ and $\widehat{\boldsymbol{r}} * \mathcal{G}$ of the true and estimated sequences $\overline{\boldsymbol{r}}$ and $\widehat{\boldsymbol{r}}$. The convolution kernel $\mathcal{G}$ is chosen here as a Gaussian filter (with a standard deviation of 0.5) truncated to a 5-sample support.

Performance are averaged over 50 realizations, except for comparisons with the MCMC-approximated Bayesian estimators (see Section 5.4) where only 20 realizations are used because of MCMC procedure's high computational cost.

5.1.3 Hyperparameter setting

The prior probability p for change-point is chosen as a uniform distribution over $(0, 1)$, obtained with hyperparameters set to $\alpha_0 = \alpha_1 = 1$. Indeed, for such a choice, the Beta distribution in (15) reduces to a uniform distribution, hence leading to a non-informative prior for the change-point probability. Amplitudes $\boldsymbol{\mu}$ for $\overline{\boldsymbol{x}}$ are parametrized with σ_0^2, which according to Proposition 3.1 should be chosen large enough. For the time being, we set $2\pi\sigma_0^2 = 10^4$, and further discuss the impact of this choice in Section 5.5.

5.1.4 Discretization of Λ

For practical purposes, we make use of a discretized subset Λ for λ (500 values equally spaced, in a $\log_{10}$-scale, between 10^{-5} and 10^5).

5.2 Illustration of the principle of the automated tuning of λ.

Fig. 1 illustrates the principle of the automated selection of λ, under various scenarios, with different values for p and different *amplitude-to-noise-ratios* [1] (ANR) where ANR= $\frac{\overline{x}_{\max} - \overline{x}_{\min}}{3\sigma}$.

For all scenarios, Fig. 1 shows that the automatically selected $\widehat{\lambda}$, obtained as the minimum of the devised criterion F (cf. (27), vertical red line in bottom row), satisfactorily falls within the ranges of λ achieving oracle estimations by minimizing either MSE (denoted Λ_{MSE} and marked with vertical lines in second row) or Jaccard error (denoted Λ_{MAP} and marked with vertical lines in third row): $\widehat{\lambda} \in \Lambda_{\mathrm{MSE}} \cap \Lambda_{\mathrm{MAP}}$. In addition, the corresponding solution $\widehat{\boldsymbol{x}}_{\widehat{\lambda}}$ (red) visually appears as a satisfactory estimator of $\overline{\boldsymbol{x}}$ (black), like the oracle propositions $\widehat{\boldsymbol{x}}_{\lambda_{\mathrm{MSE}}}$ (blue) and $\widehat{\boldsymbol{x}}_{\lambda_{\mathrm{Jac}}}$ (green). Solution $\widehat{\boldsymbol{x}}_{\widehat{\lambda}}$ indeed systematically benefits from lower relative MSE and Jaccard error than $\widehat{\boldsymbol{x}}_{\lambda}$ for any other λ. While, by construction, $\widehat{\boldsymbol{x}}_{\lambda_{\mathrm{MSE}}}$ and $\widehat{\boldsymbol{x}}_{\lambda_{\mathrm{Jac}}}$ are identical for all λ within Λ_{MSE} or Λ_{MAP}, the automated selection procedure for λ yields interestingly a single global minimizer.

When ANR decreases, a closer inspection of Fig. 1 (left column) further shows that the supports of oracle λ, Λ_{MSE} and Λ_{MAP} are drastically shrinking, yet the automated selection of λ remains satisfactory even in these more difficult contexts. The same holds when p increases (see Fig. 1, right column).

5.3 Estimation performance quantitative

To assess and quantify estimation performance of $\widehat{\lambda}$ as functions of data parameters σ^2, $\overline{x}_{\max} - \overline{x}_{\min}$, and p, we have performed Monte Carlo simulations under various settings.

First, Fig. 2 reports estimation performance for $\widehat{\lambda}$ as a function of the ANR. It shows that the estimated $\widehat{\lambda}$ (red), averaged over Monte Carlo simulations, satisfactorily remains within the range of oracles λ (dashed white lines) and well-follows the oracle mean-values (solid white line), as ANR varies. This holds for different $\overline{x}_{\max} - \overline{x}_{\min}$. As p grows (cf. Fig. 2 from top to bottom), the oracle regions in dashed white shrink, thus indicating that the selection of λ becomes more intricate when a larger amount of segments are to be detected. The proposed automated selection for λ still performs well in these more difficult situations.

---

[1] This measure allows amplitudes between successive segments to be compared w.r.t. to noise power. Since segments amplitude are drawn uniformly between $\overline{x}_{\min}$ and $\overline{x}_{\max}$ then the mean amplitude between successive segments is about $\frac{\overline{x}_{\max} - \overline{x}_{\min}}{3}$.
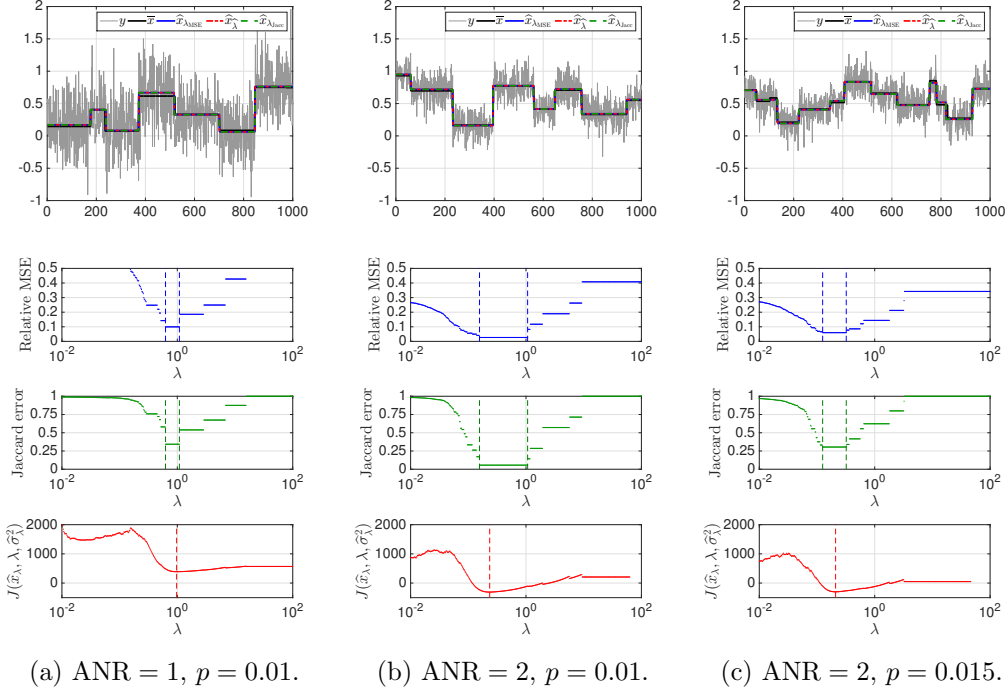
Figure 1: **Illustration of the automated tuning of** λ**:** Top: Data y to which are superimposed true signal $\overline{x}$ and oracle signals $\widehat{x}_{\lambda_{\mathrm{MSE}}}$ (blue) and $\widehat{x}_{\lambda_{\mathrm{Jac}}}$ (green) obtained for λ_{MSE} and λ_{Jac} minimizing the MSE and the Jaccard error, together with estimated $\widehat{x}_{\widehat{\lambda}}$ (red) obtained from automated selection of λ. Second and third lines: relative MSE and Jaccard error as functions of λ. Vertical lines locate λ_{MSE} and λ_{Jac}. Bottom line: Criterion F (cf. (27)) as a function of λ. Automatically selected $\widehat{\lambda}$ is indicated by vertical red lines and is satisfactorily located in between the vertical lines indicating λ_{MSE} and λ_{Jac}. (left) $p = 0.01$ and ANR $= 1$, (middle) $p = 0.01$ and ANR $= 2$, (right) $p = 0.15$ and ANR $= 2$. For all configurations, $\overline{x}_{\max} - \overline{x}_{\min} = 1$.

In addition, it can also be observed that $\widehat{\lambda}$ depends, as expected, on σ (or equivalently on $\overline{x}_{\max} - \overline{x}_{\min}$) cf. Fig. 2, from left to right.

Second, Fig. 3 focuses on the behavior of the estimated $\widehat{\lambda}$ as a function of σ, for different values of ANR. Again, it shows satisfactory performance of $\widehat{\lambda}$ compared to the oracle λ. Incidently, it also very satisfactorily reproduces the linear dependence of λ with σ^2, which can be predicted from a mere dimensional analysis of the TVℓ_0 criterion yielding:

$$\lambda \sim \frac{\sigma^2}{2p}. \tag{31}$$

5.4   Comparison with Bayesian estimators

The proposed method has been compared to classical Bayesian estimators associated with the hierarchical Bayesian model derived in Section 3.1 for which an MCMC procedure has been derived (cf. Appendix A). The number of Monte Carlo iterations is fixed to $T_{\mathrm{MC}} = 10^3$
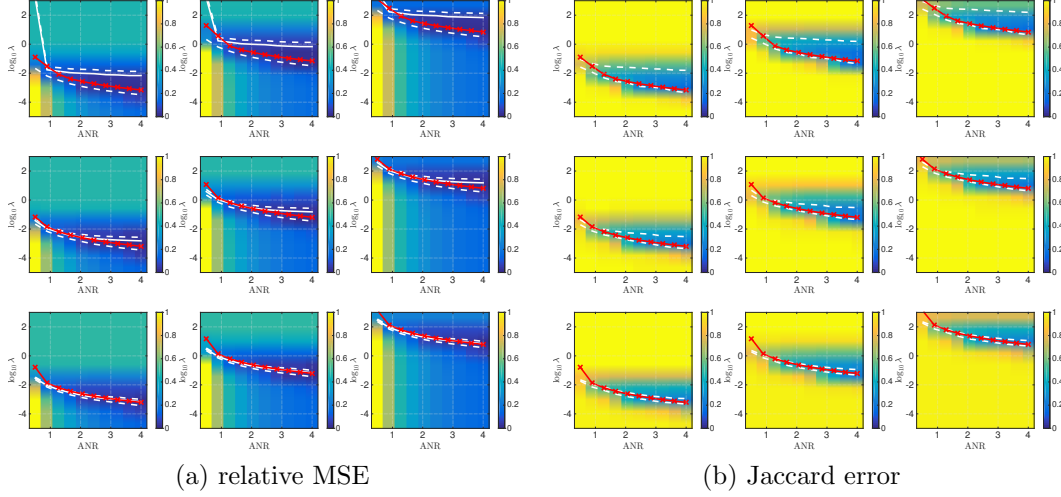
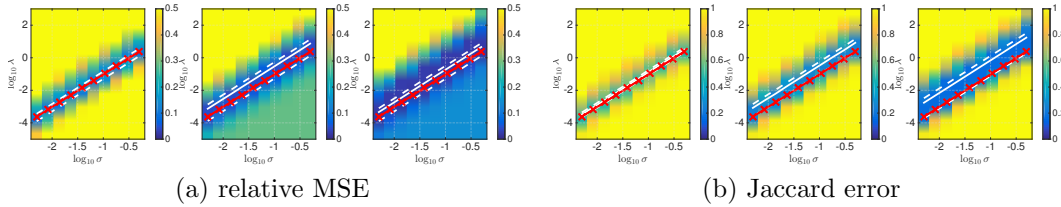(a) relative MSE   (b) Jaccard error

Figure 2: **Estimation performance: RMSE and Jaccard Error as functions of $\widehat{\lambda}$ and** ANR**.**
Estimate $\widehat{\lambda}$ (average over 50 realizations) is displayed in red, as a function of the ANR,
is shown to satisfactorily remain within the range of oracles λ, delimited by dashed white
lines and to closely follow oracle Monte Carlo average indicated by solid white lines (left:
relative MSE, right: Jaccard error). From top to bottom: $p = 0.005$, 0.010 and 0.015.
From left to right: $\overline{x}_{\max} - \overline{x}_{\min} = 0.1$, 1, and 10.

(a) relative MSE   (b) Jaccard error

Figure 3: **Estimation performance: RMSE and Jaccard Error as functions of $\widehat{\lambda}$ and σ.**
Estimate $\widehat{\lambda}$ (average over 50 realizations) is displayed in red, as a function of $\log_{10} \sigma$, is
shown to satisfactorily remain within the range of oracles λ, delimited by dashed white
lines and to closely follow oracle Monte Carlo average indicated by solid white lines (left:
relative MSE, right: Jaccard error). For each configuration $p = 0.01$ and from left to right:
ANR $= 1$, 2 and 4. This illustrates that $\widehat{\lambda}$ leads to solutions with same performance as
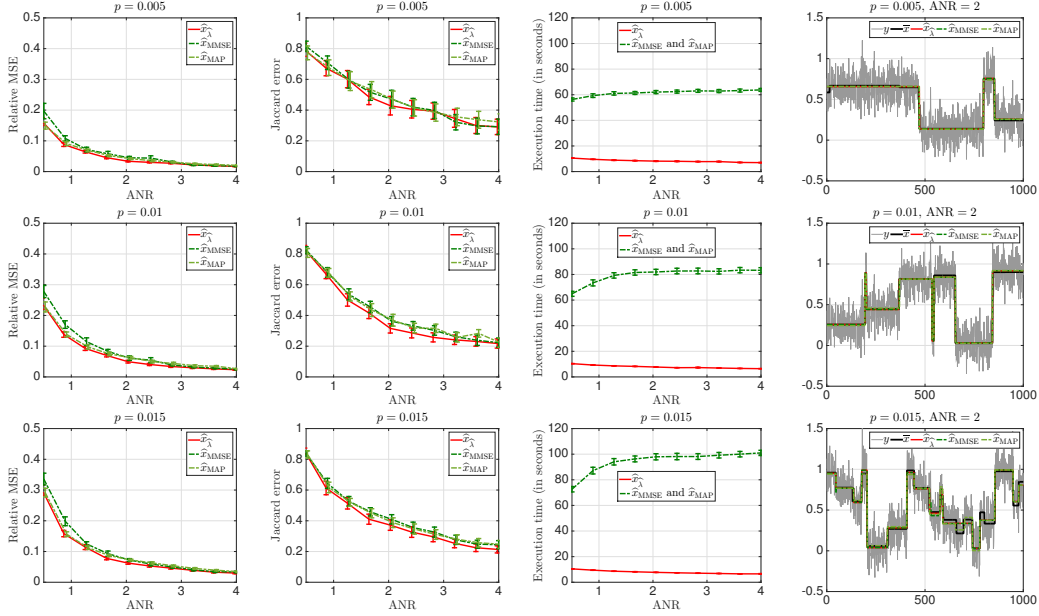oracle λ and highlight thats $\widehat{\lambda}$ varies linearly with σ^2, as expected.

Figure 4: **Comparison with Bayesian estimators.** For each configuration ANR = 2, $\overline{x}_{\max} - \overline{x}_{\min} = 1$ and from top to bottom: $p = 0.005$, 0.010 and 0.015. From left to right: relative MSE, Jaccard error, execution time and example of solutions. Overall, the proposed estimator (red) yields estimation performance comparable to Bayesian estimators (green) while benefiting from significantly lower computational costs.

and the amplitude hyperprior parameters are chosen as the mean of $\boldsymbol{y}$ for μ_0 and $\sigma_0^2 = \widehat{\mathrm{var}}(\boldsymbol{y})$, where $\widehat{\mathrm{var}}(\cdot)$ stands for the empirical variance.

In Fig. 4, estimation performance for the proposed procedure (solid red) are compared against MAP and MMSE hierarchical Bayesian estimators, as functions of ANR. Overall, $\widehat{\boldsymbol{x}}_{\widehat{\lambda}}$ is equivalent to $\widehat{\boldsymbol{x}}_{\mathrm{MAP}}$ and $\widehat{\boldsymbol{x}}_{\mathrm{MMSE}}$ in terms of MSE (first column) and Jaccard error (second column), while benefiting of significantly lower computational costs. Interestingly, when p increases (large number of change-points), the larger the gain in using the proposed procedure. This is also the case when the sample size N increases: For $N = 10^4$, the MCMC approach takes more than an hour while the method we propose here provides a relevant solution in a few minutes.

5.5 Impact of hyperparameter σ_0^2

Let us finally come back to the study of the impact of σ_0^2 on the achieved solution.

Fig. 5 displays $\widehat{\lambda}$ (red circle) as a function of σ_0^2, for different $\overline{x}_{\max} - \overline{x}_{\min}$. It shows that using $\log 2\pi\sigma_0^2 \in [0, 5]$ systematically leads to satisfactory estimates minimizing the relative MSE (left) or Jaccard Error (right). This clearly indicates that σ_0^2 does not depend on data dynamics ($\overline{x}_{\max} - \overline{x}_{\min}$), which is what is expected from a hyperparameter.

Consequently, we further set $\overline{x}_{\max} - \overline{x}_{\min} = 1$ and focus on the range $\log 2\pi\sigma_0^2 \in [0, 5]$ to explore dependencies on p or ANR. Fig. 6 shows that selecting any value of σ_0 such that $\log 2\pi\sigma_0^2 \in [0, 5]$ leads to satisfactory estimates minimizing the relative MSE (left) or Jaccard error (right), irrespective of the actual values of p or of the ANR.
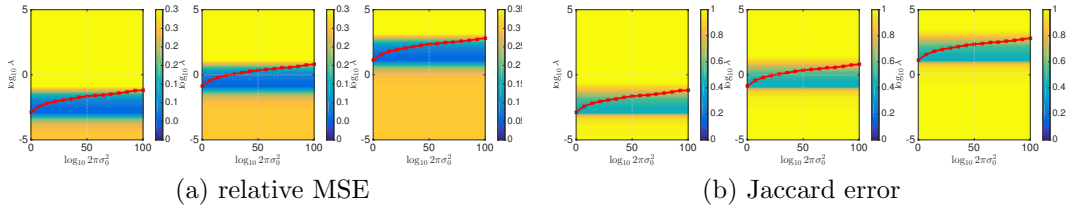
13

(a) relative MSE   (b) Jaccard error

Figure 5: **Estimation performance: RMSE and Jaccard Error as functions of $\widehat{\lambda}$ and σ_0^2.**
The estimate $\widehat{\lambda}$, which a priori explicitly depends on the choice of the hyperprior σ_0^2, is averaged over 50 realizations and displayed in red as a function of $2\pi\sigma_0^2$. From left to right: $\overline{x}_{\max} - \overline{x}_{\min} = 0.1$, 1 and 10. For each configuration $p = 0.01$, ANR = 1. Choosing $\log 2\pi\sigma_0^2 \in [0,5]$ systematically leads to satisfactory estimates minimizing the relative MSE (left) or Jaccard Error (right).

6 Conclusions and perspectives

This contribution studied a change-point detection strategy based on the TVℓ_0 minimization, whose performance depend crucially on the selection of a regularization parameter. Using an equivalence between a variational formulation and a hierarchical Bayesian formulation of the change-point detection problem, the present contribution proposed and assessed an efficient automated selection of this regularization parameter. It shows that estimation performance of the proposed procedure (evaluated in terms of global relative-MSE and Jaccard error) match satisfactorily those achieved with oracle solutions. Moreover, when compared to fully Bayesian strategies, the proposed procedure achieved equivalent performance at significantly lower computational costs. Future work could aim to extend the present framework to non Gaussian noises, such as Laplace or Poisson.

A Bayesian estimators

The maximum a posteriori (MAP) or minimum mean squared error (MMSE) estimators associated with the joint posterior $f(\mathbf{\Theta}|\mathbf{y})$ in (16) can be approximated by using MCMC procedures that essentially rely on a partially collapsed Gibbs sampler [8] similar to the algorithm derived in [13]. It consists in iteratively drawing samples (denoted $\cdot^{[t]}$) according to conditional posterior distributions that are associated with the joint posterior (16). The resulting procedure, detailed in Algo. 2, provides a set of samples $\boldsymbol{\vartheta} = \left\{ \boldsymbol{r}^{[t]}, \boldsymbol{\mu}^{[t]}, \sigma^{2[t]}, p^{[t]} \right\}_{t=1}^{T_{\mathrm{MC}}}$ that are asymptotically distributed according to (16). These samples can be used to approximate the MAP and MMSE estimators of the parameters of interest [22]. The corresponding solutions are referred to as $\widehat{\boldsymbol{x}}_{\mathrm{MAP}}$ and $\widehat{\boldsymbol{x}}_{\mathrm{MMSE}}$ in Section 5.4.

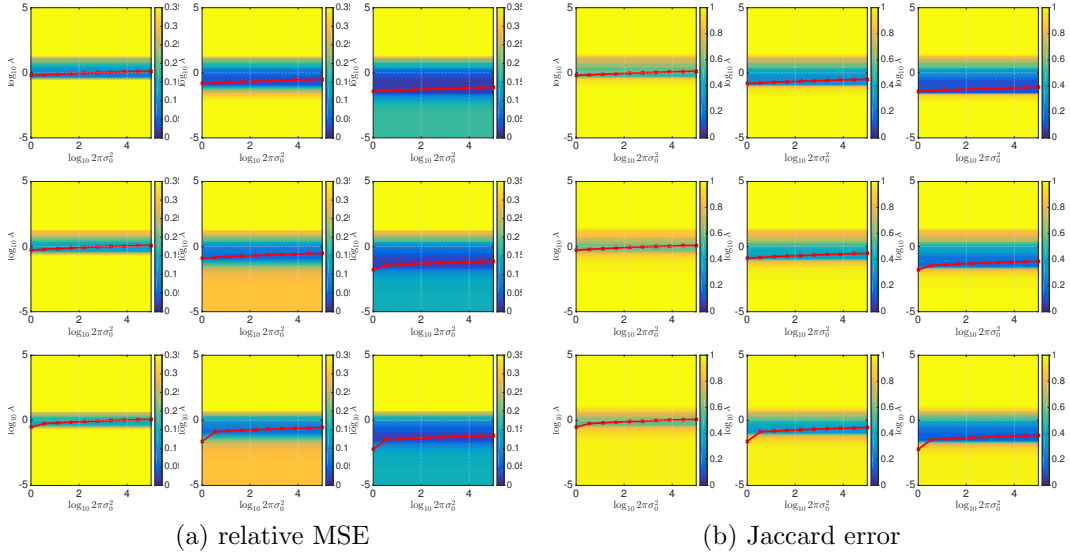(a) relative MSE   (b) Jaccard error

Figure 6: **Estimation performance: RMSE and Jaccard Error as functions of $\widehat{\lambda}$ and σ_0^2.**
The estimate $\widehat{\lambda}$, which a priori explicitly depends on the choice of the hyperprior σ_0^2, is averaged over 50 realizations and displayed in red as a function of $2\pi\sigma_0^2$. Choosing $\log 2\pi\sigma_0^2 \in [0,5]$ leads to satisfactory estimation performance independently of p and the ANR. For each configuration $\overline{x}_{\max} - \overline{x}_{\min} = 1$. From top to bottom: $p = 0.005$, 0.010 and 0.015. From left to right: ANR $= 1$, 2 and 4.

Algorithm 2 Piecewise constant Bayesian denoising

Input: Observed signal $\boldsymbol{y} \in \mathbb{R}^N$.
   Hyperparameters $\boldsymbol{\Phi} = \left\{ \alpha_0, \alpha_1, \mu_0, \sigma_0^2 \right\}$.
Iterations:
 1: **for** $t = 1, \ldots, T_{\mathrm{MC}}$ **do**
 2:    **for** $i = 1, \ldots, N-1$ **do**
 3:       Draw $r_i^{[t]} \sim f\left(r_i | \boldsymbol{y}, \boldsymbol{r}_{\backslash i}, p, \sigma^2, \mu_0, \sigma_0^2 \right)$
 4:    **end for**
 5:    **for** $k = 1, \ldots, \sum_{i=1}^{N} r_i^{[t]}$ **do**
 6:       Draw $\mu_k^{[t]} \sim f\left(\mu_k | \boldsymbol{y}, \boldsymbol{r}, \sigma^2, \mu_0, \sigma_0^2 \right)$
 7:    **end for**
 8:    Draw $\sigma^{2[t]} \sim f\left(\sigma^2 | \boldsymbol{y}, \boldsymbol{r}, \boldsymbol{\mu} \right)$
 9:    Draw $p^{[t]} \sim f\left(p | \boldsymbol{r}, \alpha_0, \alpha_1 \right)$
10: **end for**
Output: $\boldsymbol{\vartheta} = \left\{ \boldsymbol{r}^{[t]}, \boldsymbol{\mu}^{[t]}, \sigma^{2[t]}, p^{[t]} \right\}_{t=1}^{T_{\mathrm{MC}}}$, $\widehat{\boldsymbol{x}}_{\mathrm{MAP}}$ and $\widehat{\boldsymbol{x}}_{\mathrm{MMSE}}$.

References

[1] A. Barbero and S. Sra. Fast Newton-type methods for total variation regularization. In L. Getoor and T. Scheffer, editors, *International Conference on Machine Learning*, ICML '11, pages 313–320, New York, NY, USA, Jun. 2011. ACM.

[2] M. Basseville and I. Nikiforov. *Detection of Abrupt Changes: Theory and Application.* Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1993.

[3] H. H. Bauschke and P. L. Combettes. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces.* Springer, New York, 2011.

[4] A. Benazza-Benyahia and J.-C. Pesquet. Building robust wavelet estimators for multi-component images using Stein's principles. *IEEE Trans. Image Process.*, 14(11):1814–1830, Nov. 2005.

[5] L. Chaari, J.-C. Pesquet, J.-Y. Tourneret, and P. Ciuciu. Parameter estimation for hybrid wavelet-total variation regularization. In *Proc. IEEE Workshop Stat. Sign. Proc.*, Nice, France, June, 28-30 2011.

[6] A. Chambolle. An algorithm for total variation minimization and applications. *J. Math. Imag. Vis.*, 20(1-2):89–97, Jan. 2004.

[7] L. Condat. A direct algorithm for 1D total variation denoising. *IEEE Signal Process. Lett.*, 20(11):1054–1057, 2013.

[8] D.A. van Dyk and T. Park. Partially collapsed Gibbs samplers: Theory and methods. *J. American Statist. Ass.*, 103(482):790–796, June 2008.

[9] P. Davies and A. Kovac. Local extremes, runs, strings and multiresolution. *Ann. Stat.*, 29(1):1–65, 2001.

[10] C. Deledalle, S. Vaiter, G. Peyré, J. Fadili, and C. Dossal. Unbiased risk estimation for sparse analysis regularization. In *Proc. Int. Conf. Image Process. (ICIP)*, pages 3053 – 3056, Orlando, FL, USA, Sept. 30-Oct. 3 2012.

[11] N. Dobigeon and J.-Y. Tourneret. Joint segmentation of wind speed and direction using a hierarchical model. *Comput. Stat. Data Anal.*, 51(12):5603–5621, Aug. 2007.

[12] N. Dobigeon, J.-Y. Tourneret, and M. Davy. Joint segmentation of piecewise constant autoregressive processes by using a hierarchical model and a Bayesian sampling approach. *IEEE Trans. Signal Process.*, 55(4):1251–1263, Apr. 2007.

[13] N. Dobigeon, J.-Y. Tourneret, and J. D. Scargle. Joint segmentation of multivariate astronomical time series: Bayesian sampling with a hierarchical model. *IEEE Trans. Signal Process.*, 55(2):414–423, Feb. 2007.

[14] F. Friedrich, A. Kempe, V. Liebscher, and G. Winkler. Complexity penalized M-estimation: fast computation. *J. Comput. Graph. Statist.*, 17(1):201–224, 2008.

[15] R. Hamon, P. Borgnat, P. Flandrin, and C. Robardet. Relabelling vertices according to the network structure by minimizing the cyclic bandwidth sum. Accepted to *Journal of Complex Networks, arXiv preprint arXiv:1410.6108*, 2016.

[16] P. Jaccard. Distribution de la flore alpine dans le bassin des Dranses et dans quelques régions voisines. *Bulletin de la Société Vaudoise des Sciences Naturelles*, (37):241–272, 1901.

[17] S.-J. Kim, K. Koh, S. Boyd, and D. Gorinevsky. ℓ_1-trend filtering. *SIAM Review*, 51(2):339–360, May 2009.

[18] M. Lavielle. Optimal segmentation of random processes. *IEEE Trans. Signal Process.*, 46(5):1365–1373, 1998.

[19] M. Lavielle and E. Lebarbier. An application of MCMC methods for the multiple change-points problem. *Signal Process.*, 81(1):39–53, Jan. 2001.

[20] M. Little and N. Jones. Generalized methods and solvers for noise removal from piecewise constant signals. I. Background theory. *Proc. R. Soc. A*, 467:3088–3114, 2011.

[21] M. A. Little and N. S. Jones. Sparse bayesian step-filtering for high-throughput analysis of molecular machine dynamics. In *Proc. IEEE Int. Conf. Acoust., Speech and Signal Process. (ICASSP)*, pages 4162–4165, Dallas, TX, March 2010.

[22] J. Marin and C. Robert. *Bayesian Core: A Practical Approach to Computational Bayesian Statistics.* Springer, New York, NY, USA, 2007.

[23] J. P. Oliveira, J. M. Bioucas-Dias, and M. A. T. Figueiredo. Adaptive total variation image deblurring: a majorization–minimization approach. *Signal Process.*, 89(9):1683–1693, 2009.

[24] M. Pereyra, J. M. Bioucas-Dias, and M. A. T. Figueiredo. Maximum-a-posteriori estimation with unknown regularisation parameters. In *Proc. Eur. Sig. Proc. Conference*, pages 230–234, Nice, France, Aug 2015.

[25] M. Pereyra, N. Dobigeon, H. Batatia, and J.-Y. Tourneret. Estimating the granularity coefficient of a Potts-Markov random field within an MCMC algorithm. *IEEE Trans. Image Process.*, 22(6):2385–2397, June 2013.

[26] J.-C. Pesquet, A. Benazza-Benyahia, and C. Chaux. A SURE approach for digital signal/image deconvolution problems. *IEEE Trans. Signal Process.*, 57:4616–4632, 2009.

[27] E. Punskaya, C. Andrieu, A. Doucet, and W. Fitzgerald. Bayesian curve fitting using MCMC with applications to signal segmentation. *IEEE Trans. Signal Process.*, 50(3):747–758, Mar. 2002.

[28] I. W. Selesnick and I. Bayram. Total variation filtering. Technical report, 2010.

[29] C. Stein. Estimation of the mean of a multivariate normal distribution. *Ann. Stat.*, 9(6):1135–1151, 1981.

[30] M. Storath, A. Weinmann, and L. Demaret. Jump-sparse and sparse recovery using Potts functionals. *IEEE Trans. Signal Process.*, 62(14):3654–3666, July 2014.

[31] M. Storath, A. Weinmann, and M. Unser. Exact algorithms for l^1-TV regularization of real-valued or circle-valued signals. *J. Sci. Comput.*, 38(1):614–630, 2016.

[32] C. Vogel and M. Oman. Iterative methods for total variation denoising. *J. Sci. Comput.*, 17:227–238, 1996.

[33] G. Winkler, A. Kempe, V. Liebscher, and O. Wittich. *Innovations in Classification, Data Science, and Information Systems: Proceedings of the 27th Annual Conference of the Gesellschaft für Klassifikation e.V., Brandenburg University of Technology, Cottbus, March 12–14, 2003*, chapter Parsimonious Segmentation of Time Series by Potts Models, pages 295–302. Springer Berlin Heidelberg, Berlin, Heidelberg, 2005.

[34] G. Winkler and V. Liebscher. Smoothers for discontinuous signals. *Journal of Nonparametric Statistics*, 14(1-2):203–222, 2002.