

Jump-penalized least absolute values estimation of scalar or circle-valued signals

Martin Storath*, Andreas Weinmann†, Michael Unser‡

September 26, 2016

Abstract

We study jump-penalized estimators based on least absolute deviations which are often referred to as Potts estimators. They are estimators for a parsimonious piecewise constant representation of noisy data having a noise distribution which has heavier tails or which leads to many severe outliers. We consider real-valued data as well as circle-valued data which appear, for instance, as time series of angles or phase signals. We propose efficient algorithms that compute Potts estimators for real-valued scalar as well as for circle-valued data. The real-valued version improves upon the state-of-the-art solver w.r.t. to computational time. In particular for quantized data, the worst case complexity is improved. The circle-valued version is the first efficient algorithm of this kind. As an illustration, we apply our method to estimate the steps in the rotation of the bacterial flagella motor based on real biological data, and to the estimation of wind directions.

Keywords: Circle-valued data, jump penalization, least absolute deviation, piecewise constant signal, Potts estimator, step detection

1 Introduction

Signals or time-series with underlying piecewise constant groundtruth appear in various biological and medical applications; for example, the cross-hybridization of DNA [Snijders et al. (2001); Drobyshev et al. (2003); Hupé et al. (2004)], the reconstruction of brain stimuli [Winkler et al. (2005)], single-molecule fluorescence resonance energy transfer [Joo et al. (2008)], cellular ion channel functionalities [Hotz et al. (2013)], and photo-emission spectroscopy [Frick et al. (2014a)]. A classical engineering example based on piecewise constant functions is crack detection in fracture mechanics. For further applications, we refer to Little and Jones (2011a,b) and Frick et al. (2014a), and the references therein. Circle-valued time-series with underlying piecewise constant model appear for example in the rotations of the bacterial flagellar motor; see Sowa et al. (2005); Sowa and Berry (2008); Mora et al. (2009).

*Image Analysis and Learning Group, Universität Heidelberg, Germany.

†Department of Mathematics and Natural Sciences, Darmstadt University of Applied Sciences, and Institute of Computational Biology, Helmholtz Center Munich, Germany.

‡Biomedical Imaging Group, École Polytechnique Fédérale de Lausanne, Switzerland.

For such data, jump-penalized estimators are powerful tools. A classical instance is the least squares Potts estimator. It measures the complexity of the signal in terms of the number of jumps and penalizes the sums of the squared deviations from data. While the least squares approach is optimally matched to the Gaussian noise model, estimators based on least absolute deviations are more robust to noise with heavier tailed distributions such as Laplacian noise, and they perform better in the presence of outliers. The least absolute values Potts estimator is given by

$$\hat{x} = \arg \min_{x \in X^N} \gamma \|\nabla x\|_0 + \sum_{n=1}^N w_n d(x_n, y_n). \quad (\mathcal{P}_\gamma)$$

The first term $\|\nabla x\|_0 = |\{n : x_{n+1} \neq x_n\}|$ measures the complexity of x in terms of the number of jumps while the second term measures fidelity to data $y \in X^N$, where $d(x, y)$ is a metric on the space X and w is a vector of positive weights. Here, we consider the spaces $X = \mathbb{R}$ with the standard distance, and $X = \mathbb{S}^1 = \mathbb{T}$ with the arc length distance. The regularization parameter $\gamma > 0$ controls the balance between data deviation and jump penalty.

For scalar data ($X = \mathbb{R}$ with the standard distance), the least absolute values Potts estimator (\mathcal{P}_γ) takes the familiar form

$$\hat{x} = \arg \min_{x \in \mathbb{R}^N} \gamma \|\nabla x\|_0 + \|x - y\|_{1,w}, \quad (1)$$

which we refer to as L^1 -Potts estimator. Here the L^1 norm is formed with respect to the weight vector w . Due to their appearance in context of any cyclic data such as phase data or periodic data, we also consider circle-valued data, i.e., the data space $X = \mathbb{S}^1 = \mathbb{T}$ with d being the arc length distance. This results in

$$\hat{x} = \arg \min_{x \in \mathbb{T}^N} \gamma \|\nabla x\|_0 + \sum_{n=1}^N w_n \text{arclength}(x_n, y_n). \quad (2)$$

A Potts estimate is a candidate signal that minimizes the corresponding Potts functional; thus the estimates enjoy an optimal tradeoff – optimal with respect to the functional – between data fidelity and complexity measured in terms of the number of jumps. From a Bayesian perspective, Potts functionals are negative posterior loglikelihood functions with an improper prior. Minimal points correspond to the respective maximum a posteriori estimates. These connections are rather of formal than of analytic nature when considering non-finite data spaces; we refer to Winkler et al. (2005) for a detailed discussion.

Prior and related work. Potts estimators are named after Renfrey B. Potts who used the prior, respectively, the regularizing term of the functional, in his work on statistical mechanics; see Potts (1952). Related investigations can even be traced back to Ising (1925). Geman and Geman (1984) were one of the first to utilize the Potts model with least squares data fidelity (L^2 -Potts estimator) in the context of statistical image processing. In the work of Blake (1983) and Blake and Zisserman (1987), the L^2 -Potts model appears as limit case of the nowadays called Blake-Zisserman model. In image processing, the Potts model is also known as piecewise constant Mumford-Shah model after the seminal work of Mumford and Shah (1989). Donoho (1999) considers related bivariate estimators which are called wedgelets. For a thorough account on the history, we refer to Winkler et al. (2005) and the references therein.

More recently, Boysen et al. (2007, 2009) have given convergence rates and have shown consistency for the L^2 -Potts estimator. Wittich et al. (2008) have shown uniqueness of the L^2 -Potts estimate for almost all input data and idempotence of the corresponding solution operator. In the context of inverse problems, convergence rates have been investigated by Frick et al. (2014b), and iterative algorithms by Fornasier and Ward (2010) and Weinmann and Storath (2015). It is well known that an L^2 -Potts estimate can be computed exactly using dynamic programming; see Yao (1984); Mumford and Shah (1985, 1989); Auger and Lawrence (1989); Chambolle (1995); Winkler and Liebscher (2002); Jackson et al. (2005). The current state-of-the-art solver is the one of Friedrich et al. (2008) which has quadratic time and linear memory complexity. Under the additional assumption that the number of jumps of the underlying signals grows linearly with the signal length, Killick et al. (2012) could improve this to a linear time algorithm.

In contrast to the L^2 scenario, significantly less is known on Potts estimators with L^1 data terms. One of the reasons is that the L^1 norm is typically more intricate due to the non-differentiability and the missing relation to the inner product. This stands in contrast to the fact that the L^1 Potts estimator has several advantages over the least squares variant: as already mentioned, L^1 data terms are more robust to noise with heavy tailed distributions and in the presence of outliers. Friedrich et al. (2008) observed that they give in practice more robust estimates on the change points (i.e., the jumps) for the considered gene expression data. Remarkably, the L^1 -Potts estimator has a certain blind deconvolution property: it reconstructs piecewise constant signals x from given blurred data $y = \kappa * x$, whenever the unknown positive symmetric kernel κ is sufficiently narrow supported; for details see Weinmann et al. (2015). Such properties are not shared by the L^2 -Potts and L^1 -TV estimators. Concerning the computation of minimizers, Friedrich et al. (2008) were the first to give an efficient algorithm for computing an estimate of $O(N^2 \log N)$ time and $O(N^2)$ memory complexity. The authors of the present article improved this to $O(N^2)$ time and $O(N)$ space complexity; see Weinmann et al. (2015). Friedrich et al. (2008) also consider computing the Potts estimator for all parameters γ simultaneously; they propose an $O(N^3)$ time and $O(N^2)$ space algorithm for real-valued data. (We note that the present work deals with quantized and non-quantized data with values on the real line or on the circle.)

Closely related estimators are the TV estimators which penalize the total variation $\|\nabla x\|_1 = \sum_n |x_{n+1} - x_n|$ instead of the number of jumps. In particular, L^1 -TV estimators have been considered in, e.g., Fu et al. (2006); Dümbgen and Kovac (2009); Clason et al. (2009); Dong et al. (2009); Chambolle and Pock (2011); Kolmogorov et al. (2015). The Potts prior is a stronger prior in the sense that the underlying signal class of piecewise constant signals is smaller. However, whenever the signal under consideration belongs to this class, Potts estimators are more robust to higher noise levels and enjoy better reconstruction properties. On the flipside, for real-valued data, the computation of a Potts estimate is more challenging because the Potts functional is, in contrast to the L^1 -TV functional, not convex.

The Potts estimators discussed above are related to the simultaneous change point estimator (SMUCE) recently proposed by Frick et al. (2014a). In contrast to the discussed Potts estimators, which measure the data fidelity in terms of the L^1 or the L^2 norm, SMUCE uses a certain multiscale test statistics. The computational procedure is based on the aforementioned dynamic programming scheme proposed by Friedrich et al. (2008) for Potts estimators.

A variant that controls the false discovery rate has been proposed by Li et al. (2014). Frick et al. (2012) consider statistical multiresolution estimators with ℓ^∞ constraints. Xia and Qiu (2015) study an information criterion for jump-penalized estimation. We refer to Little and Jones (2011a,b) and the references therein for further related jump-penalized estimators and algorithms for the reconstruction of piecewise constant real-valued signals.

Besides real-valued data, there is an emerging interest in estimation of circle-valued data. Regression of circular data has been considered by Fisher and Lewis (1983) and by Downs and Mardia (2002), and LASSO/TV type problems by Giaquinta et al. (1993); Cremers and Strelakovski (2013); Lellmann et al. (2013); Weinmann et al. (2014); Bergmann et al. (2015); Storath et al. (2016). This is motivated by their appearance as data spaces in various contexts including phase data, orientation data, as well as nonlinear color spaces. We refer to Fisher (1995) and to Jammalamadaka and SenGupta (2001) for an overview on statistics of circle-valued data. Statistical issues on Riemannian manifolds in general are the topic of Oller and Corcuera (1995); Bhattacharya and Patrangenaru (2003, 2005); Pennec (2006); Fletcher (2013). For manifold-valued data, and thus for circle-valued data in particular, an algorithm for the Potts problem has been proposed by Weinmann et al. (2016). However, for circle-valued data, it is not guaranteed that the algorithm converges to a global minimizer. Moreover, the resulting algorithm is of iterative nature and a faster algorithm would be desirable.

Contributions. In this paper, we study least absolute values Potts estimators for real-valued as well as for circle-valued data. To our knowledge, Potts estimators for circle-valued data have not been studied yet. Our contributions are as follows: *(i)* We provide new algorithms which yield global minimizers of the Potts problem for real-valued data (cf. Equation (1)) as well as for circle-valued data (cf. Equation (2)) where the data can be quantized or non-quantized; *(ii)* we propose a strategy to compute Potts estimates for all regularization parameters $\gamma > 0$ simultaneously with moderate extra effort, and *(iii)* we provide numerical experiments on synthetic and real life data.

Concerning *(i)*, the key ingredient is the reduction of the search space to a not a priori fixed data dependent set V^N . More precisely, for the case of circle-valued data, we show that we can reduce the search space to V^N where V contains the unique values of the data y and its antipodal points. The reduction of the problem size allows us to employ the Viterbi-algorithm; see Viterbi (1967); Bellman and Roth (1969); Blake and Zisserman (1987); Felzenszwalb and Zabih (2011). By exploiting the special structure of the Potts penalty as described in Felzenszwalb and Huttenlocher (2006), this leads to an algorithm of complexity $O(KN)$ where $K \leq 2N$ is the number of elements in the discrete set V . To our knowledge, this is the first exact algorithm for the Potts problem with circle-valued data. For scalar real-valued data, we proceed analogously to obtain a solver of complexity $O(KN)$. We show that the algorithm has complexity of linearithmic order in probability when the components of the data vector y are independent, discrete random variables which are not too heavy tailed. We stress that the proposed approach is fundamentally different from those of Friedrich et al. (2008) and Weinmann et al. (2015). The latter ones employ a different dynamic programming scheme, and they do not use a strategy to reduce the search space.

Concerning *(ii)*, we consider the closely related problem of jump-constrained least absolute values approximation, which we call J -jump sparsity problem. We propose an algorithm that

solves the J -jump sparsity problem for all possible number of jumps in $O(KNL)$, where $L \leq N - 1$ denotes the number of jumps of the data. Based on this complete scanning, we are able to determine all intervals for γ , where the solution does not change, and all corresponding estimates. Again, our approach relies on a different discrete search space than the $O(KN^2)$ algorithm of Friedrich et al. (2008) for solving for all model parameters γ simultaneously. In particular, our method improves upon the computational complexity of Friedrich et al. (2008) for quantized data.

Concerning (iii), we provide a series of numerical experiments on the proposed algorithm. For real-valued (non-quantized) data, we observe that our approach is significantly faster than the previously proposed solver of Weinmann et al. (2015). (We recall that the algorithm of Weinmann et al. (2015) is an $O(N^2)$ time and $O(N)$ space algorithm improving the $O(N^2 \log N)$ time and $O(N^2)$ space algorithm of Friedrich et al. (2008).) As a tradeoff, the memory consumption is $O(KN)$ as opposed to the linear memory complexity of Weinmann et al. (2015). Concerning circle-valued data we observe that the proposed algorithm, which is the first one proposed for this problem, yields reasonable runtimes. Further, we apply our method to real data: we estimate the steps in the rotation of the bacterial flagella motor [Sowa et al. (2005)] and time series of wind directions.

Outline of the paper. In Section 2, we propose efficient algorithms for computing Potts estimates with scalar or circle-valued data. Section 3 deals with the closely related J -jump sparsity problem. In Section 4, we compute solutions of the Potts problem for all parameter simultaneously. Numerical experiments on synthetic data and on real data are conducted in Section 5.

2 Efficient computation of the Potts estimator

We develop efficient solvers for the Potts problem (\mathcal{P}_γ) for circle-valued and for real-valued data. We first reduce the infinite search space to a finite set, and then utilize the Viterbi algorithm with an acceleration proposed by Felzenszwalb and Huttenlocher (2006).

2.1 Dimensionality reduction for circle-valued and scalar data spaces

Our first step towards computing a global minimizer of the Potts functional (\mathcal{P}_γ) with circle-valued and real-valued data is the reduction of the search space to a finite set. In the following, we use the notation $\text{Val}(y)$ to denote the set of values of a tuple $y \in X^N$, i.e.,

$$\text{Val}(y) = \{v \in X : \text{there is } 1 \leq n \leq N \text{ s.t. } y_n = v\}.$$

We recall that a (weighted) median μ of $y \in X^N$ is defined as minimizer of the functional $z \mapsto \sum_{n=1}^N w_n d(z, y_n)$; see e.g. Fletcher et al. (2009). For scalar data it is well known that the set of values of y contains at least one median. More precisely, there is $\mu \in \arg \min_{z \in \mathbb{R}} \sum_{n=1}^N w_n |z - y_n|$ such that $\mu \in \text{Val}(y)$. In the case of circle-valued data, a median is contained in the set of values $\text{Val}(y)$ or the set of antipodal points $\text{Val}(-y)$:

Lemma 1. *Let $y \in \mathbb{T}^N$. The set $V = \text{Val}(y) \cup \text{Val}(-y)$ contains a (weighted) median of y .*

Lemma 1 is most probably known. Since we did not find it in the literature, we provide a short proof in Appendix A.

The key to a reduction of the search space is the following statement on solutions of the Potts problem.

Theorem 2. *Let $\gamma > 0$. For real-valued data $y \in \mathbb{R}^N$, we let $V = \text{Val}(y)$; for circle-valued data $y \in \mathbb{T}^N$, we let $V = \text{Val}(y) \cup \text{Val}(-y)$. Then the Potts problem (\mathcal{P}_γ) has a solution \hat{x} fulfilling $\hat{x} \in V^N$.*

The proof is given in Appendix A.

We next discuss how to determine of the reduced search space V when data y is given. In many setups, we can a priori declare a set U that is only slightly larger than V or, in the best case, equal to V . Then it is convenient to use this set U . For example, we can use $U = \{0, \dots, 255\}$ for an 8-bit signal. If there is no such set we have to determine V algorithmically. A simple method to do so is based on sorting of y . Then, duplicate elements are adjacent and can be easily removed. However, this procedure may exceed the complexity of the subsequent step (see Section 2.2), because general purpose sorting algorithms are $O(N \log N)$. To avoid sorting, we utilize the following simple procedure: We create an empty list V . Then we iterate over the components of y , that is, for $n = 1, \dots, N$. If y_n is not in the list, we append it to V . In case of circle-valued data, we additionally append the antipodal points. The procedure needs at most $K = |V|$ comparisons at each step, thus the procedure is $O(KN)$.

2.2 A Viterbi-type algorithm for the reduced problem

Theorem 2 allows us to reduce the infinite search spaces \mathbb{R}^N and \mathbb{T}^N in (\mathcal{P}_γ) to the finite set V^N , which is specified in that theorem. Thus, it remains to solve

$$\arg \min_{x \in V^N} \gamma \|\nabla x\|_0 + \sum_{n=1}^N w_n d(x_n, y_n). \quad (3)$$

To this end, we use dynamic programming. Recall that the basic idea of dynamic programming is to solve a complex problem by reducing it to smaller subproblems; see Bellman (1957). In the following, we denote the cardinality of V by K . It is useful to equip the set V with indices, i.e., $V = \{v_1, \dots, v_K\}$. (Note that any ordering is valid.)

We utilize a dynamic programming scheme developed by Viterbi (1967); see also Forney Jr (1973). Related algorithms have been proposed by Bellman and Roth (1969) and Blake and Zisserman (1987). Here, we review a special instance of the Viterbi algorithm proposed by Felzenszwalb and Huttenlocher (2006); see also Felzenszwalb and Zabih (2011). Assume that we aim at minimizing an energy functional of the form

$$E(x_1, \dots, x_N) = \gamma \sum_{n=1}^{N-1} h(x_n, x_{n+1}) + \sum_{n=1}^N w_n d(x_n, y_n) \quad (4)$$

where x_1, \dots, x_N can take values in a finite set V and where h is an energy functional on V^2 . The Viterbi algorithm solves this problem in two steps: tabulation of energies and reconstruction by backtracking. For the tabulation, the starting point is the table $B^1 \in \mathbb{R}^K$ given by

$$B_k^1 = w_1 d(v_k, y_1) \quad \text{for } k = 1, \dots, K.$$

We successively compute for $n = 2, \dots, N$ the tables $B^n \in \mathbb{R}^K$ which are given by

$$B_k^n = w_n d(v_k, y_n) + \min_l \{B_l^{n-1} + \gamma h(v_k, v_l)\}, \quad \text{for } k = 1, \dots, K. \quad (5)$$

The entry B_k^n represents the energy of a minimizer on data (y_1, \dots, y_n) whose endpoint is equal to v_k . For the backtracking step, it is convenient to introduce an auxiliary tuple $l \in \mathbb{N}^N$ which stores minimizing indices. We initialize the last entry of l by $l_N = \arg \min_k B_k^N$. Then we successively compute the entries of l for $n = N - 1, N - 2, \dots, 1$ by

$$l_n = \arg \min_k B_k^n + \gamma h(v_k, v_{l_{n+1}}). \quad (6)$$

Eventually, we reconstruct a minimizer \hat{x} from the indices in l by $\hat{x}_n = v_{l_n}$, for $n = 1, \dots, N$. The result \hat{x} is a global minimizer of the energy (4); see Felzenszwalb and Zabih (2011).

As proposed by Felzenszwalb and Huttenlocher (2006), the special structure of the Potts prior can be exploited to streamline the computation. First note that the reduced Potts problem (3) corresponds to the functional defined by $h(v, w) = 1$ for $v \neq w$ and $h(v, v) = 0$. This allows to simplify the second term in (5) to

$$\min_l \{B_l^{n-1} + \gamma h(v_k, v_l)\} = \min\{\gamma + \min_l B_l^{n-1}; B_k^{n-1}\}.$$

It follows that the table B^n can be filled by the two-step procedure

$$\begin{cases} z = \min_l B_l^{n-1}, \\ B_k^n = w_n d(y_n, v_k) + \min\{\gamma + z; B_k^{n-1}\}, \quad \text{for } k = 1, \dots, K. \end{cases} \quad (7)$$

Now z can be precomputed and reused for all table entries B_k^n for $k = 1, \dots, K$. Since both lines of (7) are $O(K)$, the computation of the table B^n is $O(K)$. Hence, this algorithm is $O(KN)$.

2.3 Complete solver for the Potts problem

The computation of the global minimizer of the Potts problem (\mathcal{P}_γ) is performed in two steps: computation of the reduced search space V followed by the above variant of the Viterbi algorithm. A pseudocode of the complete algorithm is given in Algorithm 1. We record the following:

Theorem 3. *Algorithm 1 computes a global minimizer of the Potts problem (\mathcal{P}_γ) with scalar or circle-valued data in $O(KN)$.*

The proof is given in the Appendix.

For the Potts problem with scalar data, a solver of complexity $O(N^2 \log N)$ has been proposed by Friedrich et al. (2008) and improved to $O(N^2)$ by Weinmann et al. (2015). These methods use also the dynamic programming principle but with a fundamentally different tabulation scheme. It is based on computing the median deviations on all $O(N^2)$ intervals. In contrast, the tabulation in Algorithm 1 is based on the distances of single data points y_n to all the K median candidates. For the Potts problem with circle-valued data, no efficient solver has been proposed yet to the best of our knowledge.

Algorithm 1: Solver for the Potts problem (\mathcal{P}_γ) with real-valued or circle-valued data

Input: $y \in \mathbb{R}^N$ or $y \in \mathbb{T}^N$: data; $\gamma > 0$: model parameter; $w \in (\mathbb{R}_0^+)^N$: weights;
Output: Global minimizer \hat{x} of (\mathcal{P}_γ);
begin
 / Search space reduction */*
 Determine reduced search space V (see Section 2.1);
 / Tabulation */*
 for $k \leftarrow 1$ **to** K **do**
 $B_k^1 \leftarrow w_1 d(v_k, y_1)$;
 end
 for $n \leftarrow 2$ **to** N **do**
 $z \leftarrow \min_{k=1, \dots, K} B_k^{n-1}$;
 for $k \leftarrow 1$ **to** K **do**
 $B_k^n \leftarrow w_n d(v_k, y_n) + \min(z + \gamma, B_k^{n-1})$;
 end
 end
 / Backtracking */*
 $l \leftarrow \arg \min_{k=1, \dots, K} B_k^N$;
 $\hat{x}_n \leftarrow v_l$;
 for $n \leftarrow N - 1, N - 2, \dots, 1$ **do**
 $B_l^n \leftarrow B_l^n - \gamma$;
 $l \leftarrow \arg \min_{k=1, \dots, K} B_k^n$;
 $\hat{x}_n \leftarrow v_l$;
 end
 return \hat{x} ;
end

2.4 Algorithmic complexity for quantized data.

In general, the number of elements of V , denoted by K , is less or equal to the number of data points N for scalar data and less or equal to $2N$ for circle-valued data. Thus the worst case complexity of Algorithm 1 is $O(N^2)$. In many situations K is significantly smaller than N . For quantized data on the unit circle, the number of discrete values K is bounded. Thus we obtain linear complexity in this case.

In the case of quantized data on the real line, the number of values might not always be bounded. However, we can give a bound for the following typical situation. Let f be a time-continuous piecewise constant function on the interval $[0, 1]$ with finitely many jumps, denoted by $\|\nabla f\|_0$. We further assume that f is a càdlàg function which means that the limits $f(t+)$ and $f(t-)$ exist and $f(t+) = f(t)$ for all $t \in [0, 1]$. Let \bar{y}^N be the (point or integral) sampling of f at density $1/N$; that is, $\bar{y}_n^N = f\left(\frac{n}{N+1}\right)$ or $\bar{y}_n^N = \frac{1}{N} \int_{\frac{n-1}{N}}^{\frac{n}{N}} f(t) dt$. Data at sampling density $1/N$ is given by

$$y_n^N = \bar{y}_n^N + \eta_n^N \quad (8)$$

where η^N is an N -dimensional vector of independent random variables. Since

$$K = |\text{Val}(y_n^N)| \leq |\text{Val}(\bar{y}_n^N)| |\text{Val}(\eta_n^N)| \leq C_f |\text{Val}(\eta_n^N)|$$

	Circle-valued data	Real-valued data
Non-quantized data	$\mathcal{O}(N^2)$	$\mathcal{O}(N^2)$
Quantized data	$\mathcal{O}(N)$	$\mathcal{O}(N^2)$ general case $\mathcal{O}_p(N \log N)$ for noise with exp. decay (Prop. 4) $\mathcal{O}(N)$ for bounded noise

Table 1: Complexity of the proposed solver for the least absolute values Potts problem (Algorithm 1) for quantized and non-quantized, real-valued and circle-valued data. The last two rows for quantized real-valued data refer to data of the form (8).

with the constant $C_f = 2(\|\nabla f\|_0 + 1)$, the number of expected values is bounded by the number of expected values of the noise η . If the supports of the distribution of the components of η are uniformly bounded with respect to n, N , then K is bounded and we obtain linear complexity of Algorithm 1. For unbounded supports we get a slightly higher algorithmic complexity.

Proposition 4. *Let $C > 0$ and $0 < p < 1$. If η^N is an N -dimensional vector of independent random variables where each component follows an integer-valued probability distribution satisfying the exponential decay condition $P(\eta_n^N = k) \leq Cp^{|k|}$, then $|\text{Val}(\eta^N)| \in \mathcal{O}_p(\log N)$.*

The proof is given in the Appendix. Examples of distributions that satisfy the hypothesis of Proposition 4 are the discrete Laplace distribution, the discrete Gaussian distribution, and the geometric distribution. We summarize the complexity of Algorithm 1 for the discussed data in Table 1.

Remark 1. The statement of Proposition 4 remains true when each of the probability distributions of the independent random variables η_n^N are discrete with a global minimum separation distance, i.e., there is $\varepsilon > 0$ such that any pair of neighboring values a, b in the range of any random variable η_n^N fulfill $|a - b| > \varepsilon$. In this case the decay condition of Proposition 4 reads $P(\eta_n^N = a) \leq Cp^{|a|}$ where a is in the range of the random variable η_n^N .

3 The J -jump sparsity problem

Next we consider the J -jump sparsity problem given by

$$\operatorname{argmin}_{\|\nabla x\|_0 \leq J} \sum_{n=1}^N w_n d(x_n, y_n), \quad (\mathcal{P}_J)$$

where $J \in \mathbb{N}_0$. Its solution will be useful for solving the Potts problem (\mathcal{P}_γ) for all parameters simultaneously (see next section). Apart from this, (\mathcal{P}_J) is an interesting problem in its own right which has applications, for example, in minimum error quantization; see Bruce (1965).

3.1 Basic properties of the J -jump sparsity problem

In the following, we denote by L the number of jumps of the data y ; that is, $L = \|\nabla y\|_0$. For $J \geq L$, the J -jump sparsity problem (\mathcal{P}_J) has the trivial solution y . Thus, we can restrict ourselves to the case of J not being larger than L .

Although the J -jump sparsity problem (\mathcal{P}_J) and the Potts problem (\mathcal{P}_γ) are closely related, they are not fully equivalent in the sense that a minimizer of (\mathcal{P}_J) may not appear as minimizer of (\mathcal{P}_γ), no matter of the value of γ . To see this, consider the scalar data $y = (0, 1, 0)$ and as data term the Euclidean distance with uniform weights. The possible solutions for $J = 1$ are the one-jump solutions of the form $(0, a, a)$, $(a, a, 0)$, $a \in [0, 1]$, and the zero-jump solution $(0, 0, 0)$. However, one-jump solutions cannot be minimizers of the Potts functional: their Potts functional values are equal to $\gamma + 1$, but this is for all $\gamma > 0$ strictly higher than the functional value of the zero jump solution which is equal to 1. An analogous example can be given for circle-valued data.

The above example also shows that the J -jump sparsity problem can have solutions with $(J - 1)$ jumps. Next, we argue that a solution of the J -jump sparsity problem cannot have less than $(J - 1)$ jumps. Assume that a solution x^J has $(J - 2)$ or fewer jumps. Then we find an index n such that $x_n^J \neq y_n$. Setting x_n^J equal to y_n adds at most two jumps. Thus, we have still at most J jumps. But the approximation error of that candidate is lower than that of x^J , which contradicts the fact that x^J minimizes the J -jump sparsity problem.

3.2 An efficient solver based on dynamic programming

As with the Potts problem, the key to a fast solver is a reduction of the search space to the finite set V^N :

Theorem 5. *Let $J \in \mathbb{N}_0$ and let V be as in Theorem 2. The J -jump sparsity problem (\mathcal{P}_J) has a solution \hat{x} that fulfills $\hat{x} \in V^N$.*

The proof is an obvious modification of the proof of Theorem 2.

The computational solution is also similar to the solver of the Potts problem. The main difference is that we now require JN tables of length K instead of N tables of length K . A table entry $B_k^{n,j}$ ($1 \leq n \leq N$, $1 \leq j \leq J$, $1 \leq k \leq K$) stores the approximation energy of the optimal j -jump solution x for data (y_1, \dots, y_n) with $x_n = v_k$. The first J tables $B^{1,j}$ are given by $B_k^{1,j} = w_1 d(v_k, y_1)$, for $j = 0, \dots, J$. Using the dynamic programming principle, we fill the tables $B^{n,j}$ successively for $n = 1, \dots, N$ and for $j = 0, \dots, J$ by

$$B_k^{n,j} = \min \left\{ z; w_n d(v_k, y_n) + B_k^{n-1,j} \right\}$$

where $z = \min_{l=1, \dots, K} B_l^{n-1, j-1}$. As z does not depend on k , it can be computed separately. Filling the table $B^{n,j}$ requires K steps; hence, the total complexity for the tabulation is $\mathcal{O}(NKJ)$. For the backtracking step, we first determine the minimizer l_N of the table $B^{N,J}$, i.e., $l_N = \arg \min_k B_k^{N,J}$. (As for the Potts problem, the tuple $l \in \mathbb{N}^N$ stores minimizing indices.) We further initialize the number of remaining jumps by $i_N = J$. For $n = N - 1, \dots, 1$, we decide – based on the tables – whether it is advantageous to take a jump or to continue with the same value via

$$(l_{n-1}, i_{n-1}) = \underset{(k,j) \in \{(1, i_n - 1), \dots, (K, i_n - 1), (l_n, i_n)\}}{\operatorname{argmin}} B_k^{n,j}.$$

The procedure terminates if $i_n = 0$, and then we set $l_1 = \dots = l_n$. Having computed the minimizing indices l_n , we obtain a minimizer \hat{x} of the J -jump sparsity problem by $\hat{x}_n = v_{l_n}$, for $n = 1, \dots, N$. The complexity of the backtracking step is $\mathcal{O}(KN)$.

A minor modification of the backtracking step allows us to compute the minimizers for the jump constrained problem for all parameters $J' \leq J$ without having to build up new tables. We just have to initialize the backtracking step with $i_N = J'$ instead of $i_N = J$. In particular, if we want to compute minimizers for all possible number of jumps $J = 0, \dots, L$, we only need compute the tabulation for $J = L$ and perform $L + 1$ backtracking steps with the initializations $i_N = 0, \dots, L$. This amounts to the complexity $O(KNL)$.

We eventually summarize the results of the above derivation:

Theorem 6. *The above algorithm computes a global minimizer of (\mathcal{P}_J) in $O(KNJ)$. Computing minimizers for all $J = 0, \dots, L$ simultaneously is $O(KNL)$.*

4 Solving the Potts problem for all parameters simultaneously

We can solve the Potts problem (\mathcal{P}_γ) for all parameters $\gamma > 0$ simultaneously with a moderate effort. To that end, we exploit the following relation between the Potts and the J -jump sparsity problems.

Lemma 7. *Let $\gamma > 0$ and let \hat{x} be a solution of the Potts problem (\mathcal{P}_γ) . Then \hat{x} is also a solution of the jump constrained problem (\mathcal{P}_J) with parameter $J = \|\nabla \hat{x}\|_0$.*

The proof is given in Appendix A. The example in the first paragraph of Section 3 shows that the opposite direction of Lemma 7 is not true in general. Now assume that we have computed for each $J = 0, \dots, L$, a minimizer for the J -jump sparsity problem, which we denote by x^J . (Recall that this can be done in $O(KNL)$ by Theorem 6). Lemma 7 tells us that the set $\{x^0, \dots, x^L\}$ already contains a solution of (\mathcal{P}_γ) for each $\gamma > 0$. This implies that we obtain a solution \hat{x} of (\mathcal{P}_γ) by scanning through the J -jump solutions; that is,

$$\hat{x} \in \operatorname{argmin}_{x \in \{x^0, \dots, x^L\}} \gamma \|\nabla x\|_0 + \sum_{n=1}^N w_n d(x_n, y_n).$$

It remains to identify the connection between γ and J . We denote the minimal Potts functional value for some $\gamma > 0$ by $\phi(\gamma)$, and we rewrite it as

$$\phi(\gamma) = \min_x P_\gamma(x) = \min_{x \in \{x^0, \dots, x^L\}} \gamma \|\nabla x\|_0 + \sum_{n=1}^N w_n d(x_n, y_n) = \min_{J=0, \dots, L} \gamma J + \epsilon_J, \quad (9)$$

where $\epsilon_J = \sum_{n=1}^N w_n d(x_n^J, y_n)$ is the approximation error of a J -jump solution x^J . Of even greater interest is the function that maps γ to the minimizing argument J :

$$\hat{J}(\gamma) = \operatorname{argmin}_{J=0, \dots, L} \gamma J + \epsilon_J. \quad (10)$$

With the help of \hat{J} , we obtain a minimizer of (\mathcal{P}_γ) by $\hat{x} = x^{\hat{J}(\gamma)}$.

Let us now explain how to compute the mappings ϕ and \hat{J} given by (9) and (10) efficiently. Being the pointwise infimum of $L + 1$ affine linear functions, ϕ is concave and piecewise linear (with at most $L + 1$ pieces). (It is the largest concave minorant of these affine linear functions.) The graph of ϕ is a polygon which is determined by at most L points; see Figure 1. Let us




Figure 1: The function $\phi(\gamma)$ (blue solid line) is the largest minorant of the affine functions $\gamma J + \epsilon_J$ for $J = 0, \dots, L$ (red dashed lines).

denote the affine function $\gamma J + \epsilon_J$ by g^J . We first compute the intersection of g^0 and g^1 and store that point on a stack. We proceed inductively as follows. Assume that we have computed the concave minorant, ϕ^J , of the first J affine functions. Also assume that we have stored the determining nodes on a stack S where at the top of the stack is the leftmost node. In the step $J+1$ we compute the intersection of the line g^{J+1} with ϕ^J . Since the slope of g^{J+1} is larger than all slopes of ϕ^J there is at most one point of intersection. We first delete the nodes of ϕ^J that lie above the line g^{J+1} ; that is, we pop the stack until we encounter a point which is below the new line. Then we compute the intersection of the line through the point we have popped last and the point at top of the stack. We push that point to the stack. We proceed until no more lines have to be added. The points in the stack then determine the graph of ϕ . (The slope to ∞ is equal to 0 and the slope to 0 is given by the number of jumps of the data, $\|\nabla y\|_0$.) The mapping \hat{J} is determined by the slopes of ϕ .

We eventually summarize the results of this section:

Theorem 8. *There is an $O(KNL)$ algorithm that computes minimizers of (\mathcal{P}_γ) for all parameters $\gamma > 0$ simultaneously.*

5 Numerical experiments

We present numerical results on synthetic data with known ground truth, and we show results on real data. The experiments were conducted on a desktop computer (Intel Xeon E5, 3.5 GHz, 32 GB RAM). The proposed algorithms were implemented in Matlab.

5.1 Synthetic real-valued data

An extensive numerical studies on the denoising performance of the L^1 -Potts estimator has been given in earlier works by Friedrich et al. (2008) and Weinmann et al. (2015). Therefore,

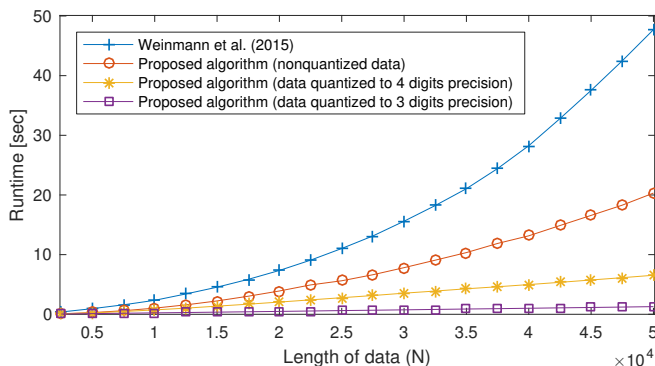


Figure 2: Runtime of our algorithm for the Potts problem with real-valued data. The proposed method is faster than the state-of-the-art solver of Weinmann et al. (2015). With a prior quantization, we are able to process even very large data within a few seconds.

we here focus on comparison to the runtimes of the solver by Weinmann et al. (2015).¹ (Recall that this solver improved upon the first solver of Friedrich et al. (2008) both in time and memory complexity.) As a basis for our experiment we use the signal previously employed in Figure 1 of Weinmann et al. (2015) as ground truth; it has eight jumps and a dynamic range of $[0, 1]$. We corrupted it by Laplacian noise of standard deviation $\sigma = 0.1$. As a result, the final signal has $K = N$ different values. In Figure 2, we show the average runtime with respect to ten different realizations of the noise. We observe that Algorithm 1 is about twice as fast as the state-of-the-art solver for general non-quantized data. For quantized data, the proposed algorithm becomes significantly faster as the length of the data increases. Thus, it might be reasonable to quantize the data to a desired precision before applying the proposed solver when dealing with very large data. We note that the runtimes of the algorithm of Weinmann et al. (2015) and the proposed algorithm for nonquantized data are by construction virtually independent of the number of jumps of the underlying signal.

In Figure 3, we compare the runtimes for different noise distributions where data is quantized to three decimals. Uniform noise is bounded, so the runtime grows linearly in N . Laplacian noise leads to linearithmic scaling as it satisfies the hypothesis of Proposition 4. Even for Cauchy noise, the complexity grows significantly slower than quadratic, although it has infinite variance and its probability density has very slow decay. In all cases, the computational time is significantly lower than the one of the benchmark method of Weinmann et al. (2015) which exhibits almost no dependence of the runtimes on the noise.

5.2 Synthetic circle-valued data

We illustrate the performance of the Potts estimator for denoising synthetic circle-valued signals (with known ground truth). To this end we create random jump-sparse signals as follows: We create a compound Poisson distributed random vector $s \in \mathbb{R}^N$; that is, $s_n = 0$ with probability $e^{-\lambda}$ and s_n is uniformly distributed in $[-a, a]$ with probability $1 - e^{-\lambda}$; see Unser and Tafti (2014). Here we use $\lambda = 0.05$ and $a = 1/2$. The (true) signal \bar{y} is given as the summation

¹The reference implementation of the solver by Weinmann et al. (2015) is available at <http://potts1ab.de>.




Figure 3: Runtime of our algorithm for the Potts problem for data with different noise distributions. All signals are quantized to three decimals. The runtimes are averages over 100 realizations. We observe that the runtime grows linearly in N for uniform noise. Laplacian noise leads to linearithmic scaling as suggested by Proposition 4. Even for Cauchy noise, the computational effort grows significantly slower than quadratic.

process of the innovation s interpreted as angle; that is, $\bar{y}_n = \exp(i \sum_{j=1}^n s_j)$. We add noise to the phase by $y = \bar{y} \exp(i\eta)$, where η is a Laplacian distributed random vector with standard deviation $\sigma = 0.4$. In Figure 4 we report a realization of the random signal and its Potts estimate for different model parameters γ . In Figure 5 we show the deviation of the estimates to the ground truth with respect to the sum over the pointwise arc length distances and with respect to the number of jumps in dependence of the model parameter γ . (Recall that the Potts estimate for all model parameter γ can be computed with moderate effort by the algorithm described in Section 4.) Figure 6 shows the results for a different noise distribution, namely a wrapped Gaussian distribution with $\sigma = 0.3$.

Eventually we note that the runtimes for circle-valued data are of the same order of magnitude as those of the real-valued case for boundedly quantized and unquantized data, respectively.

5.3 Results on real-world data

Next we apply our method to real-world data. First, we look at the estimation of steps in the rotation of the bacterial flagellar motor. The bacterial flagellar motor is a rotary molecular machine that is embedded in the bacterial cell envelope [Sowa et al. (2005); Sowa and Berry (2008)]. It propels many species of swimming bacteria. Sowa et al. (2005) observed steps in the rotation of the flagellar motor over time; in average, they found approximately 26 steps per revolution. The present data set is a time series of the angular position of the flagellar motor. The data was acquired using back-focal plane interferometry (BFP). The measurement uncertainty of the data acquisition results in a slightly blurred signal. In view of the built-in deconvolution property of the L^1 -Potts functional (cf. Theorem 4.4 in Weinmann et al. (2015)), we expect that a slight blur does not negatively affect the performance of the Potts estimator. In Figure 7, we report the estimate for the model parameter $\gamma = 1$.

Our second real-life data set consists of wind directions at the station WPOW1 (West Point,

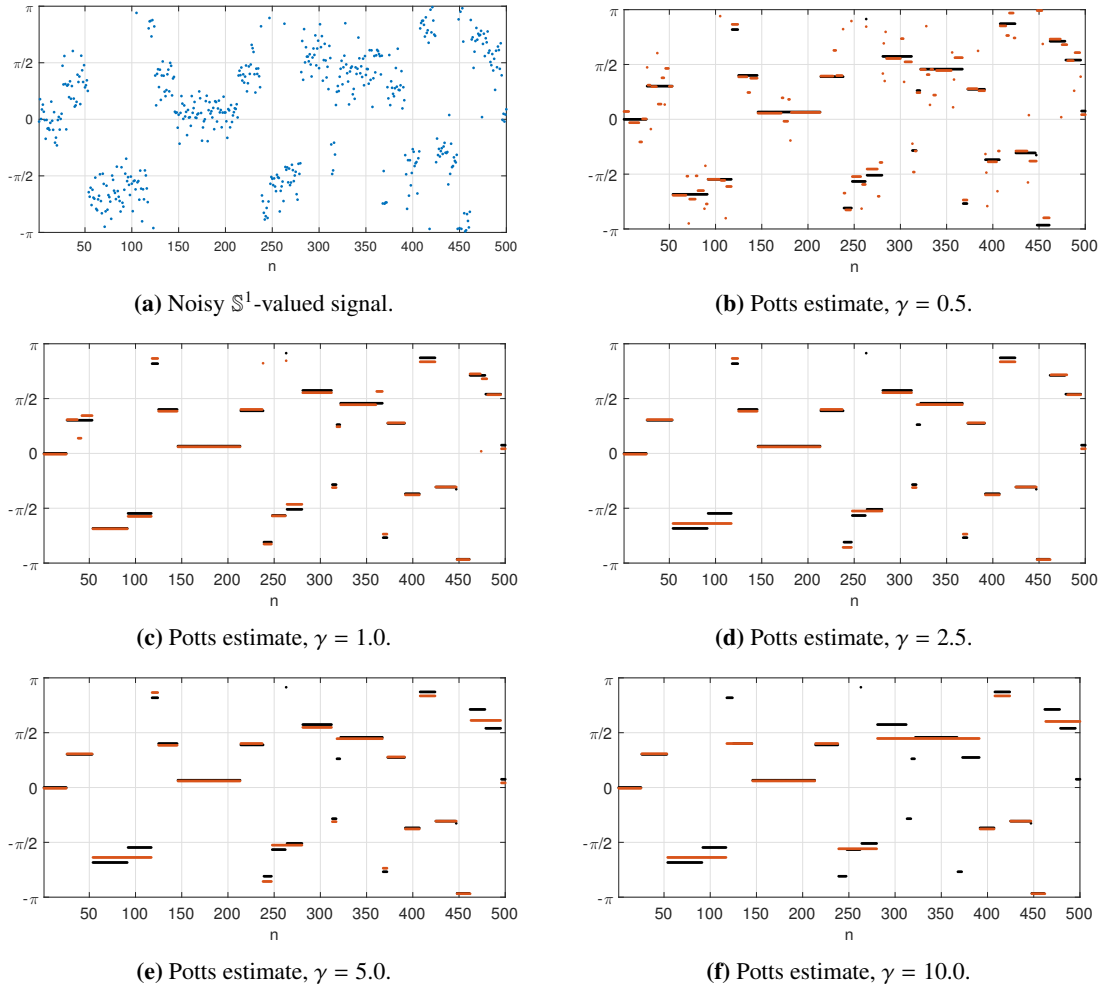


Figure 4: (a) A circle-valued signal of length $N = 500$ in angular representation corrupted with wrapped Laplacian noise with $\sigma = 0.4$. (b) – (f) Potts estimate visualized for various values of the model parameter γ , plotted as red line. The ground truth is drawn in black. (See also Figure 5 for the errors in dependance of the γ parameter.)




Figure 5: The plots illustrate the total absolute deviation of the Potts estimate \hat{x} and to ground truth \bar{y} (left) and the difference w.r.t. to the number of jumps in dependance of model parameter γ (right). The graphs correspond to the signal in Figure 4.

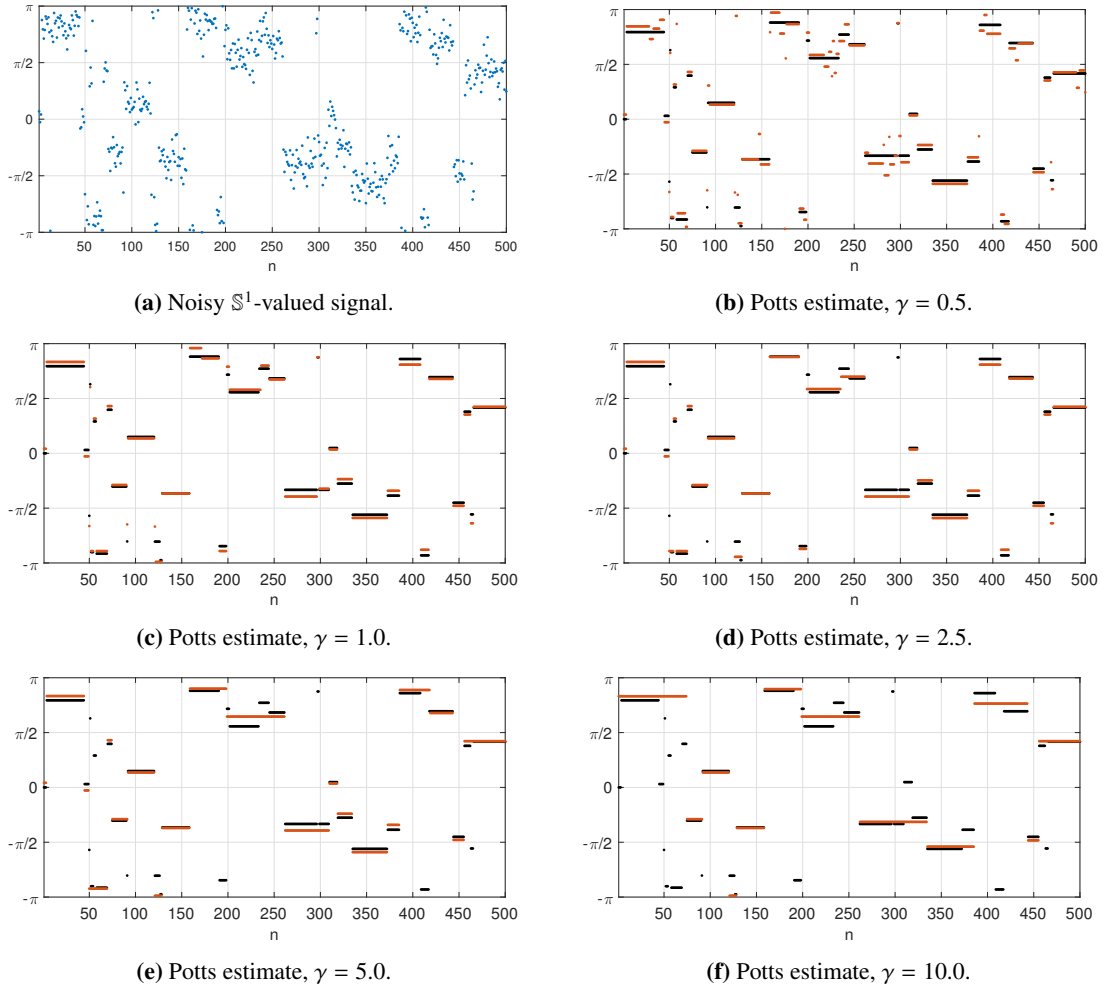


Figure 6: (a) A circle-valued signal of length $N = 500$ in angular representation corrupted with wrapped Gaussian noise with $\sigma = 0.3$. (b) – (f) Potts estimate visualized for various values of the model parameter γ , plotted as red line. The ground truth is drawn in black.

WA) recorded every hour in the year 2013.² The Potts estimate facilitates to identify the time intervals of approximately constant wind direction. The data is given quantized to integer angles in degrees, i.e., $K = 360$. In Figure 8, we report the result for the model parameter $\gamma = 10$.

6 Conclusion

We have studied the Potts estimator based on least absolute deviations for real-valued data and for circle-valued data. We have proposed exact solvers for the Potts problem with real-valued and circle-valued data. Here, the key observation was that solutions take values in a finite

²Data available at http://www.ndbc.noaa.gov/historical_data.shtml.

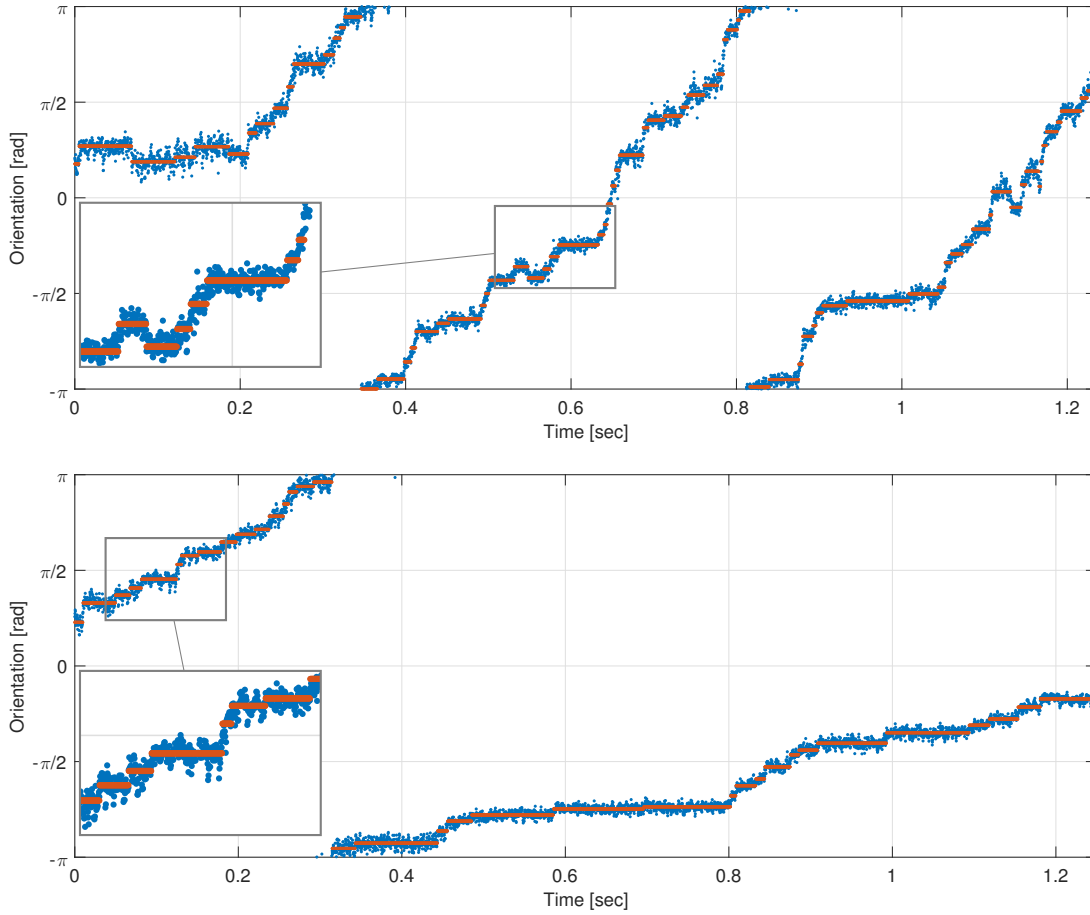


Figure 7: The blue points represent the angle of the bacterial flagellar motor over time. The red line is the least absolute values Potts estimate. (Original data by courtesy of Sowa et al. (2005)).

set (also for non-quantized data) which allowed us to employ a Viterbi-type dynamic program as proposed in Felzenszwalb and Huttenlocher (2006). Furthermore, we have obtained an algorithm for solving the Potts problem for all parameters γ simultaneously which was based on solving the related J -jump sparsity problem. The proposed algorithm for real-valued data improves upon the state-of-the-art solver with respect to computational time. The circle-valued variant is the first exact solver for the circle-valued Potts problem to our knowledge. We have illustrated the practical utility for the estimation of two real-life data sets: estimation of steps in the rotation of the bacterial flagellar motor and estimation of time series of wind directions. An interesting topic of future research are strategies for choosing the model parameter.

Acknowledgement

Martin Storath and Andreas Weinmann acknowledge the support by the German Research Foundation (STO1126/2-1, WE5886/4-1). Andreas Weinmann was supported by the Helmholtz




Figure 8: Wind directions at Station WPOW1 (West Point, WA) recorded every hour in the year 2013. (The data is given quantized to $K = 360$ angles.) The red line is the least absolute values Potts estimate for circle-valued data.

Association within the young investigator group VH-NG-526. Michael Unser was supported by the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007-2013)/ERC grant agreement no. 267439. We would like to thank Yoshiyuki Sowa, Alexander Rowe, Mark Leake, Toshiharu Yakushi, Michio Homma, Akihiko Ishijima, and Richard Berry for providing the data set of the bacterial flagellar motor. We further would like to thank the anonymous reviewers for their valuable comments and suggestions.

A Proofs

Proof of Lemma 1. Let $z \in \mathbb{T} \setminus V$ be arbitrary, and let us denote the median-defining functional by $f(v) = \sum_{n=1}^N w_n d(v, y_n)$. The diameter passing through z cuts the unit circle into two hemispheres. We denote the counterclockwise hemisphere relative to z by S^+ and the other one by S^- . Let us denote the total weight of the elements in S^+ by W^+ , i.e. $W^+ = \sum_{n: y_n \in S^+} w_n$, and analogously let $W^- = \sum_{n: y_n \in S^-} w_n$. Without loss of generality, we can assume that S^- is not “heavier” than S^+ , i.e., $W^- \leq W^+$. Let $z' \in V$ be the nearest neighbor of z in V on S^+ . The distance between each element in S^+ and z' is by $d(z', z)$ smaller than its distance to z . Hence, replacing z by z' , the functional value of f is decreased by $\sum_{n: y_n \in S^+} w_n d(z', z) = d(z', z)W^+$. With an analogous argument for the other hemisphere, f is increased at the same time by $d(z', z)W^-$. Thus, we have that

$$f(z') = f(z) + d(z', z)W^- - d(z', z)W^+ = f(z) + d(z', z)(W^- - W^+) \leq f(z),$$

where the inequality follows from $W^- \leq W^+$. So, for every $z \in \mathbb{T} \setminus V$, there is z' in V that has a smaller or equal functional value as z . This implies that there is a median in V . \square

Proof of Theorem 2. Let $Q = \{q_1, \dots, q_t\} \subset \{1, \dots, N-1\}$ be a set of t jump indices, and let X_Q be the set of $x \in X^N$ that jump at the index indicated by Q ; that is, $x_n \neq x_{n+1}$ for $n \in Q$ and

$x_n = x_{n+1}$ otherwise. The minimizer of the Potts restricted to X_Q may be expressed as

$$\min_{x \in X_Q} P_\gamma(x) = \gamma|Q| + \min_{x \in X_Q} \sum_{n=1}^N w_n d(x_n, y_n).$$

Since the elements of X_Q are constant between two jumps we get that

$$\min_{x \in X_Q} \sum_{n=1}^N w_n d(x_n, y_n) = \min_{\mu \in X} \sum_{n=1}^{q_1} w_n d(\mu, y_n) + \dots + \min_{\mu \in X} \sum_{n=q_r+1}^N w_n d(\mu, y_n).$$

Hence, $x \in X_Q$ minimizes $\min_{x \in X_Q} P_\gamma(x)$ if x equals a median between any two jumps. By Lemma 1, a median is contained in the set V . Considering all potential sets of jump indices (which are finitely-many) completes the proof. \square

Proof of Theorem 3. It follows from Theorem 2 and its subsequent paragraph that the search space reduction is $O(KN)$. The utilized variant of the Viterbi algorithm solves the reduced problem in $O(KN)$, see Felzenszwalb and Huttenlocher (2006). Thus, the overall method is $O(KN)$. \square

Proof of Proposition 4. We have that

$$P(|\eta_n^N| \geq m) \leq C \sum_{|j| \geq m} p^{|j|} = C \left[\left(2 \frac{1}{1-p} - p \right) - \left(2 \frac{1-p^m}{1-p} - p \right) \right] \leq C' p^m$$

for a constant $C' > 0$. Now let m_0 be such that $C' p^{m_0} < 1$. The cardinality of the unique values, i.e., $|\text{Val}(\eta^N)|$, can be estimated by $|\text{Val}(\eta^N)| \leq 2 \max_n |\eta_n^N| + 1$, where $|\eta_n^N|$ is the ordinary absolute value of η_n^N . It follows for $m \geq m_0$ that

$$\begin{aligned} P(|\text{Val}(\eta^N)| < m) &\geq P(\max_n |\eta_n^N| < (m-1)/2) = \prod_{n=1}^N P(|\eta_n^N| < (m-1)/2) \\ &\geq \prod_{n=1}^N (1 - C' p^{(m-1)/2}) = (1 - C' p^{(m-1)/2})^N \geq 1 - C' N p^{(m-1)/2}. \end{aligned}$$

Then, letting $m = -\frac{4}{\log p} \log N$ yields

$$\begin{aligned} P(|\text{Val}(\eta^N)| < -\frac{4}{\log p} \log N) &\geq 1 - C' p^{-1/2} N p^{\frac{2 \log N}{-\log p}} \\ &= 1 - C' p^{-1/2} N \exp\left(\frac{2 \log N}{-\log p} \log p\right) = 1 - C' p^{-1/2} N^{-1}. \end{aligned}$$

Hence $P(|\text{Val}(\eta^N)| \geq C'' \log N) \rightarrow 0$ as $N \rightarrow \infty$, with $C'' = -\frac{4}{\log p}$, which implies that $|\text{Val}(\eta^N)| \in \mathcal{O}_p(\log N)$. \square

Proof of Lemma 7. Since \hat{x} is a minimizer of the Potts functional, it is a minimizer of the problem

$$\min_x \sum_{n=1}^N w_n d(x_n, y_n), \quad \text{s.t. } \|\nabla x\|_0 = \|\nabla \hat{x}\|_0. \quad (11)$$

By the minimality of \hat{x} with respect to the Potts functional, we further get that $\sum_{n=1}^N w_n d(x_n, y_n) > \sum_{n=1}^N w_n d(\hat{x}_n, y_n)$ for all x with $\|\nabla x\|_0 < \|\nabla \hat{x}\|_0$. This means that all candidates with less than or equally many jumps as \hat{x} do not lead to a smaller value of the target functional in (11). Therefore, \hat{x} is a solution of (\mathcal{P}_J) for the parameter $J = \|\nabla \hat{x}\|_0$. \square

Proof of Theorem 8. It follows from Lemma 7 that the set of solutions of the Potts problem (\mathcal{P}_γ) is contained in the set of solutions of the jump constrained problem (\mathcal{P}_J). The latter can be computed for all number of jumps in $O(KNL)$ by (6). For mapping the parameter value γ to the smallest number of jumps of a corresponding Potts solution (\mathcal{P}_γ), we need to compute intersections of certain lines and graphs of concave functions as explained in the paragraph before Theorem 8. Namely, in step $J + 1$ ($0 \leq J \leq L$) we compute the intersection of the line g^{J+1} with the concave piecewise affine function ϕ^J . (Here, we use the notation as in the paragraph before Theorem 8.) The costs for this can be bounded from above by $O(L^2)$, and, as $L \leq N - 1$, also by $O(NL)$. Summing all up, we get the complexity $O(KNL)$. □

References

- I. Auger and C. Lawrence. Algorithms for the optimal identification of segment neighborhoods. *Bulletin of Mathematical Biology*, 51(1):39–54, 1989.
- R. Bellman. *Dynamic Programming*. Princeton University Press, 1957.
- R. Bellman and R. Roth. Curve fitting by segmented straight lines. *Journal of the American Statistical Association*, 64(327):1079–1084, 1969.
- R. Bergmann, F. Laus, G. Steidl, and A. Weinmann. Second order differences of cyclic data and applications in variational denoising. *SIAM J. Imaging Sci.*, 2015. to appear.
- R. Bhattacharya and V. Patrangenaru. Large sample theory of intrinsic and extrinsic sample means on manifolds I. *Annals of Statistics*, 31(1):1–29, 2003.
- R. Bhattacharya and V. Patrangenaru. Large sample theory of intrinsic and extrinsic sample means on manifolds II. *Annals of Statistics*, 33(3):1225–1259, 2005.
- A. Blake. The least-disturbance principle and weak constraints. *Pattern Recognition Letters*, 1:393–399, 1983.
- A. Blake and A. Zisserman. *Visual reconstruction*. MIT Press Cambridge, 1987.
- L. Boysen, V. Liebscher, A. Munk, and O. Wittich. Scale space consistency of piecewise constant least squares estimators: Another look at the regressogram. *IMS Lecture Notes Monograph Series*, 55:65–84, 2007.
- L. Boysen, A. Kempe, V. Liebscher, A. Munk, and O. Wittich. Consistencies and rates of convergence of jump-penalized least squares estimators. *Annals of Statistics*, 37(1):157–183, 2009.
- J. Bruce. Optimum quantization. Technical report, Massachusetts Institute of Technology, 1965.
- A. Chambolle. Image segmentation by variational methods: Mumford and Shah functional and the discrete approximations. *SIAM Journal on Applied Mathematics*, 55(3):827–863, 1995.
- A. Chambolle and T. Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of Mathematical Imaging and Vision*, 40(1):120–145, 2011.
- C. Clason, B. Jin, and K. Kunisch. A duality-based splitting method for ℓ^1 -TV image restoration with automatic regularization parameter choice. *SIAM Journal on Scientific Computing*, 32(3):1484–1505, 2009.
- D. Cremers and E. Strelakovski. Total cyclic variation and generalizations. *Journal of Mathematical Imaging and Vision*, 47(3):258–277, 2013.
- Y. Dong, M. Hintermüller, and M. Neri. An efficient primal-dual method for L^1 TV image restoration. *SIAM Journal on Imaging Sciences*, 2(4):1168–1189, 2009.
- D. Donoho. Wedgelets: Nearly minimax estimation of edges. *Annals of Statistics*, 27(3):859–897, 1999.
- T. Downs and K. Mardia. Circular regression. *Biometrika*, 89(3):683–698, 2002.
- A. Drobyshev, C. Machka, M. Horsch, M. Seltmann, V. Liebscher, M. de Angelis, and J. Beckers. Specificity assessment from fractionation experiments (safe): a novel method to evaluate microarray probe specificity based on hybridisation stringencies. *Nucleic acids research*, 31(2):1–10, 2003.
- L. Dümbgen and A. Kovac. Extensions of smoothing via taut strings. *Electronic Journal of Statistics*, 3:41–75, 2009.
- P. Felzenszwalb and R. Zabih. Dynamic programming and graph algorithms in computer vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(4):721–740, 2011.

- P. Felzenszwalb and D. Huttenlocher. Efficient belief propagation for early vision. *International Journal of Computer Vision*, 70(1):41–54, 2006.
- N. Fisher. *Statistical Analysis of Circular Data*. Cambridge University Press, 1995.
- N. Fisher and T. Lewis. Estimating the common mean direction of several circular or spherical distributions with differing dispersions. *Biometrika*, 70(2):333–341, 1983.
- P. Fletcher. Geodesic regression and the theory of least squares on Riemannian manifolds. *International Journal of Computer Vision*, 105(2):171–185, 2013.
- P. T. Fletcher, S. Venkatasubramanian, and S. Joshi. The geometric median on Riemannian manifolds with application to robust atlas estimation. *NeuroImage*, 45(1):S143–S152, 2009.
- M. Fornasier and R. Ward. Iterative thresholding meets free-discontinuity problems. *Foundations of Computational Mathematics*, 10(5):527–567, 2010.
- G. Forney Jr. The Viterbi algorithm. *Proceedings of the IEEE*, 61(3):268–278, 1973.
- K. Frick, P. Marnitz, and A. Munk. Statistical multiresolution Dantzig estimation in imaging: fundamental concepts and algorithmic framework. *Electronic Journal of Statistics*, 6:231–268, 2012.
- K. Frick, A. Munk, and H. Sieling. Multiscale change point inference. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 76(3):495–580, 2014a.
- S. Frick, T. Hohage, and A. Munk. Asymptotic laws for change point estimation in inverse regression. *Statistica Sinica*, 24:555–575, 2014b.
- F. Friedrich, A. Kempe, V. Liebscher, and G. Winkler. Complexity penalized M-estimation: Fast computation. *Journal of Computational and Graphical Statistics*, 17(1):201–224, 2008.
- H. Fu, M. Ng, M. Nikolova, and J. Barlow. Efficient minimization methods of mixed ℓ^1 - ℓ^1 and ℓ^2 - ℓ^1 norms for image restoration. *SIAM Journal on Scientific Computing*, 27(6):89–97, 2006.
- S. Geman and D. Geman. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6(6):721–741, 1984.
- M. Giaquinta, G. Modica, and J. Souček. Variational problems for maps of bounded variation with values in S^1 . *Calculus of Variations and Partial Differential Equations*, 1(1):87–121, 1993.
- T. Hotz, O. M. Schutte, H. Sieling, T. Polupanow, U. Diederichsen, C. Steinem, and A. Munk. Idealizing ion channel recordings by a jump segmentation multiresolution filter. *IEEE Transactions on NanoBioscience*, 12(4):376–386, 2013.
- P. Hupé, N. Stransky, J. Thiery, F. Radvanyi, and E. Barillot. Analysis of array CGH data: from signal ratio to gain and loss of DNA regions. *Bioinformatics*, 20(18):3413–3422, 2004.
- E. Ising. Beitrag zur Theorie des Ferromagnetismus. *Zeitschrift für Physik A: Hadrons and Nuclei*, 31(1):253–258, 1925.
- B. Jackson, J. D. Scargle, D. Barnes, S. Arabhi, A. Alt, P. Gioumoussis, E. Gwin, P. Sangtrakulcharoen, L. Tan, and T. T. Tsai. An algorithm for optimal partitioning of data on an interval. *IEEE Signal Processing Letters*, 12(2):105–108, 2005.
- S. Jammalamadaka and A. SenGupta. *Topics in Circular Statistics*. World Scientific Publishing Company, 2001.
- C. Joo, H. Balci, Y. Ishitsuka, C. Buranachai, and T. Ha. Advances in single-molecule fluorescence methods for molecular biology. *Annual Review of Biochemistry*, 77:51–76, 2008.
- R. Killick, P. Fearnhead, and I. Eckley. Optimal detection of changepoints with a linear computational cost. *Journal of the American Statistical Association*, 107(500):1590–1598, 2012.
- V. Kolmogorov, T. Pock, and M. Rolinek. Total variation on a tree. *Preprint arXiv:1502.07770*, 2015.
- J. Lellmann, E. Strelakovski, S. Koetter, and D. Cremers. Total variation regularization for functions with values in a manifold. In *IEEE International Conference on Computer Vision (ICCV)*, pages 2944–2951, 2013.
- H. Li, A. Munk, and H. Sieling. FDR-control in multiscale change-point segmentation. *Preprint arXiv:1412.5844*, 2014.
- M. Little and N. Jones. Generalized methods and solvers for noise removal from piecewise constant signals. I. Background theory. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Science*, 467(2135):3088–3114, 2011a.
- M. Little and N. Jones. Generalized methods and solvers for noise removal from piecewise constant signals. II. New methods. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Science*, 467(2135):3115–3140, 2011b.
- T. Mora, H. Yu, Y. Sowa, and N. Wingreen. Steps in the bacterial flagellar motor. *PLoS Computational Biology*, 5(10):e1000540, 2009.

- D. Mumford and J. Shah. Boundary detection by minimizing functionals. In *IEEE Conference on Computer Vision and Pattern Recognition*, volume 17, pages 137–154, 1985.
- D. Mumford and J. Shah. Optimal approximations by piecewise smooth functions and associated variational problems. *Communications on Pure and Applied Mathematics*, 42(5):577–685, 1989.
- J. Oller and J. Corcuera. Intrinsic analysis of statistical estimation. *Annals of Statistics*, 23(5):1562–1581, 1995.
- X. Pennec. Intrinsic statistics on Riemannian manifolds: Basic tools for geometric measurements. *Journal of Mathematical Imaging and Vision*, 25(1):127–154, 2006.
- R. Potts. Some generalized order-disorder transformations. *Mathematical Proceedings of the Cambridge Philosophical Society*, 48(1):106–109, 1952.
- A. Snijders, N. Nowak, R. Segraves, et al. Assembly of microarrays for genome-wide measurement of DNA copy number by CGH. *Nature Genetics*, 29:263–264, 2001.
- Y. Sowa and R. Berry. Bacterial flagellar motor. *Quarterly Reviews of Biophysics*, 41(02):103–132, 2008.
- Y. Sowa, A. Rowe, M. Leake, T. Yakushi, M. Homma, A. Ishijima, and R. Berry. Direct observation of steps in rotation of the bacterial flagellar motor. *Nature*, 437(7060):916–919, 2005.
- M. Storath, A. Weinmann, and M. Unser. Exact algorithms for L^1 -TV regularization of real-valued or circle-valued signals. *SIAM Journal on Scientific Computing*, 38(1):A614–A630, 2016.
- M. Unser and P. Tafti. *An introduction to sparse stochastic processes*. Cambridge University Press, 2014.
- A. Viterbi. Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *IEEE Transactions on Information Theory*, 13(2):260–269, 1967.
- A. Weinmann and M. Storath. Iterative Potts and Blake-Zisserman minimization for the recovery of functions with discontinuities from indirect measurements. *Proceedings of the Royal Society A*, 471(2176), 2015.
- A. Weinmann, M. Storath, and L. Demaret. The L^1 -Potts functional for robust jump-sparse reconstruction. *SIAM Journal on Numerical Analysis*, 53(1):644–673, 2015.
- A. Weinmann, L. Demaret, and M. Storath. Mumford-Shah and Potts regularization for manifold-valued data. *Journal of Mathematical Imaging and Vision*, 55(3):428–445, 2016.
- A. Weinmann, L. Demaret, and M. Storath. Total variation regularization for manifold-valued data. *SIAM Journal on Imaging Sciences*, 7(4):2226–2257, 2014.
- G. Winkler and V. Liebscher. Smoothers for discontinuous signals. *Journal of Nonparametric Statistics*, 14(1-2): 203–222, 2002.
- G. Winkler, O. Wittich, V. Liebscher, and A. Kempe. Don’t shed tears over breaks. *Jahresbericht der Deutschen Mathematiker-Vereinigung*, 107:57–87, 2005.
- O. Wittich, A. Kempe, G. Winkler, and V. Liebscher. Complexity penalized least squares estimators: Analytical results. *Mathematische Nachrichten*, 281(4):582–595, 2008.
- Z. Xia and P. Qiu. Jump information criterion for statistical inference in estimating discontinuous curves. *Biometrika*, 102(2):397–408, 2015.
- Y.-C. Yao. Estimation of a noisy discrete-time step function: Bayes and empirical Bayes approaches. *Annals of Statistics*, 12(4):1434–1447, 1984.