Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence

Outlier-Robust Convex Segmentation

Itamar Katz and **Koby Crammer**
Department of Electrical Engineering
The Technion – Israel Institute of Technology
Haifa, 32000 Israel
(itamark@tx,koby@ee).technion.ac.il

Abstract

We derive a convex optimization problem for the task of segmenting sequential data, which explicitly treats presence of outliers. We describe two algorithms for solving this problem, one exact and one a top-down novel approach, and we derive a consistency results for the case of two segments and no outliers. Robustness to outliers is evaluated on two real-world tasks related to speech segmentation. Our algorithms outperform baseline segmentation algorithms.

1 Introduction

Segmentation of sequential data, also known as change-point detection, is a fundamental problem in the field of unsupervised learning, and has applications in diverse fields such as speech processing (Brent 1999; Qiao, Shimomura, and Minematsu 2008; Shriberg et al. 2000), text processing (Beeferman, Berger, and Lafferty 1999), bioinformatics (Olshen et al. 2004) and network anomaly detection (Lévy-Leduc and Roueff 2009), to name a few. We are interested in formulating the segmentation task as a convex optimization problem that avoids issues such as local-minima or sensitivity to initializations. In addition, we want to explicitly incorporate robustness to outliers. Given a sequence of samples $\{\mathbf{x}_i\}_{i=1}^n$, for $\mathbf{x}_i \in \mathbb{R}^d$, our goal is to segment it into a few subsequences, where each subsequence is homogeneous under some criterion. Our starting point is a convex objective that minimizes the sum of squared distances of samples $\mathbf{x}_i$ from each sample's associated 'centroid', $\boldsymbol{\mu}_i$. Identical adjacent $\boldsymbol{\mu}_i$s identify their corresponding samples as belonging to the same segment. In addition, some of the samples are allowed to be identified as outliers, allowing reduced loss on these samples. Two regularization terms are added to the objective, in order to constrain the number of detected segments and outliers, respectively.

We propose two algorithms based on this formulation, both alternate between detecting outliers, which is solved analytically, and solving the problem with modified samples , which can be solved iteratively. The first algorithm, denoted by Outlier-Robust Convex Sequential (ORCS) segmentation, solves the optimization problem exactly, while the second is a top-down hierarchical version of the algorithm, called TD-ORCS. We also derive a weighted version of this algorithm, denoted by WTD-ORCS. We show that for the case of $K = 2$ segments and no outliers, a specific choice of the weights leads to a solution which recovers the exact solution of an un-relaxed optimization problem.

We evaluate the performance of the proposed algorithms on two speech segmentation tasks, for both clean sources and sources contaminated with added non-stationary noise. Our algorithms outperform other algorithms in both the clean and outlier-contaminated setting. Finally, based on the empirical results, we propose a heuristic approach for approximating the number of outliers.

Notation: The samples to be segmented are denoted by $\{\mathbf{x}_i \in \mathbb{R}^d\}_{i=1}^n$, and their associated quantities (both variables and solutions) are $\boldsymbol{\mu}_i, \boldsymbol{z}_i \in \mathbb{R}^d$. The same notation with no subscript, $\boldsymbol{\mu}$, denotes the collection of all $\boldsymbol{\mu}_i$'s. The same holds for $\mathbf{x}, \boldsymbol{z}$. We abuse notation and use $\boldsymbol{\mu}_i$ to refer to both the 'centroid' vector of a segment (these are not center of mass, due to the regularization term), and to the indexes of measurements assigned to that segment.

2 Outlier-Robust Convex Segmentation

Segmentation is the task of dividing a sequence of n data samples $\{\mathbf{x}_i\}_{i=1}^n$, into K groups of consecutive samples, or segments, such that each group is homogeneous with respect to some criterion. A common choice of such a criterion often involves minimizing the squared Euclidean distance of a sample to some representative sample $\boldsymbol{\mu}_i$. This criterion is highly sensitive to outliers and indeed, as we show empirically below, the performance of segmentation algorithms degrades drastically when the data is contaminated with outliers. It is therefore desirable to incorporate robustness to outliers into the model. We achieve this by allowing some of the input samples $\mathbf{x}_i$ to be identified as outliers, in which case we do not require $\boldsymbol{\mu}_i$ to be close to these samples. To this end we propose to minimize:

$$\min_{\boldsymbol{\mu},\boldsymbol{z}} \left\{ \frac{1}{2} \sum_{i=1}^{n} \|\mathbf{x}_i - \boldsymbol{z}_i - \boldsymbol{\mu}_i\|^2 \right\}$$
$$\text{s.t.} \sum_{i=1}^{n-1} \mathbb{1}\left\{ \|\boldsymbol{\mu}_{i+1} - \boldsymbol{\mu}_i\|_p > 0 \right\} \leq K - 1 \,,$$
$$\sum_{i=1}^{n} \mathbb{1}\left\{ \|\boldsymbol{z}_i\|_q > 0 \right\} \leq M \,,$$

where $p, q \geq 1$. Considering samples $\mathbf{x}_i$ for which $\boldsymbol{z}_i = 0$, the objective measures the loss of replacing a point $\mathbf{x}_i$ with some shared point $\boldsymbol{\mu}_i$, and can be thought of as minus the log-likelihood under Gaussian noise. Samples i with $\boldsymbol{z}_i \neq 0$ are intuitively identified as outliers. The first constraint bounds the number of segments by K, while the second constraint bounds the number of outliers by M. The optimal value for a nonzero $\boldsymbol{z}_i$ is to set $\boldsymbol{z}_i = \mathbf{x}_i - \boldsymbol{\mu}_i$, making the contribution to the objective zero, and thus in practice ignoring this sample, treating it as an outlier. We note that a similar approach to robustness was employed by (Forero, Kekatos, and Giannakis 2011) in the context of robust clustering, and by (Mateos and Giannakis 2012) in the context of robust PCA. Since the ℓ_0 constraints results in a non convex problem, we use a common practice and replace it with a convex surrogate ℓ_1 norm which induces sparsity. For the $\boldsymbol{\mu}_i$ variables it means that for most samples we will have $\boldsymbol{\mu}_{i+1} - \boldsymbol{\mu}_i = 0$, allowing the identification of the corresponding samples as belonging to the same segment. For the $\boldsymbol{z}_i$ variables, it means that most will satisfy $\boldsymbol{z}_i = 0$ and for some of them, the outliers, otherwise. We now incorporate the relaxed constraints into the objective, and in addition consider a slightly more general formulation in which we allow weighting of the summands in the first constraint. We get the following optimization problem:

$$\min_{\boldsymbol{\mu},\boldsymbol{z}} \left\{ \frac{1}{2} \sum_{i=1}^{n} \left\| \mathbf{x}_i - \boldsymbol{z}_i - \boldsymbol{\mu}_i \right\|^2 + \lambda \sum_{i=1}^{n-1} w_i \left\| \boldsymbol{\mu}_{i+1} - \boldsymbol{\mu}_i \right\|_p + \gamma \sum_{i=1}^{n} \left\| \boldsymbol{z}_i \right\|_q \right\}, \tag{1}$$

where w_i are weights to be determined. The parameter $\lambda > 0$ can be thought of as a tradeoff parameter between the first term which is minimized with n segments, and the second term which is minimized with a single segment. As λ is decreased, it crosses values at which there is a transition from K segments to $K+1$ segments, in a phase-transition like manner where $1/\lambda$ is the analog of temperature. The parameter $\gamma > 0$ controls the amount of outliers, where for $\gamma = \infty$ we enforce $\boldsymbol{z}_i = 0$ for all samples, and for $\gamma = 0$ the objective is optimal for $\boldsymbol{z}_i = \mathbf{x}_i - \boldsymbol{\mu}_i$, and thus all samples are in-fact outliers. Alternatively, one can think of λ, γ as the Lagrange multipliers of a constrained optimization problem. In what follows we consider $p = 2$ and $q = 1, 2$, focusing empirically on $q = 2$. Note that $q = 1$ encourages sparsity of coordinates of $\boldsymbol{z}_i$, and not of the vector as a whole. This amounts to outliers being modeled as noise in few features or samples, respectively.

2.1 Algorithms

The decoupling between $\boldsymbol{\mu}$ and $\boldsymbol{z}$ allows us to optimize Eq. (1) in an alternating manner, and we call this algorithm Outlier-Robust Convex Sequential (ORCS) segmentation. Holding $\boldsymbol{\mu}$ constant, optimizing over $\boldsymbol{z}$ is done analytically by noting that Eq. (1) becomes the definition of the proximal operator evaluated at $\mathbf{x}_i - \boldsymbol{\mu}_i$, for which a closed-form solution exists. For $q = 1$ the objective as a function of $\boldsymbol{z}$ is separable both over coordinates and over data samples, and the proximal operator is the shrinkage-and-threshold operator evaluated at each coordinate k:

$$\text{prox}_\gamma\left(v_k\right) = \text{sign}\left(v_k\right) \cdot \max\left\{0, |v_k| - \gamma\right\}.$$

However, we are interested in zeroing some of the $\boldsymbol{z}_i$'s as a whole, so we set $q = 2$. In this case, the objective is separable over data samples, and the proximal operator is calculated to be:

$$\text{prox}_\gamma\left(\boldsymbol{v}\right) = \boldsymbol{v} \cdot \max\left\{0, 1 - \frac{\gamma}{\left\|\boldsymbol{v}\right\|_2}\right\}. \tag{2}$$

Holding $\boldsymbol{z}$ constant, optimizing over $\boldsymbol{\mu}$ is done by defining $\hat{\mathbf{x}}_i \triangleq \mathbf{x}_i - \boldsymbol{z}_i$, which results in the following optimization problem:

$$\min_{\boldsymbol{\mu}} \sum_{i=1}^{n} \left\| \hat{\mathbf{x}}_i - \boldsymbol{\mu}_i \right\|^2 + \lambda \sum_{i=1}^{n-1} w_i \left\| \boldsymbol{\mu}_{i+1} - \boldsymbol{\mu}_i \right\|_2. \tag{3}$$

Note that Eq. (3) is equivalent to Eq. (1) with no outliers present. We also note that if we plug the analytical solution for the $\boldsymbol{z}_i$s into Eq. (3) (via the $\hat{\mathbf{x}}_i$s), the loss term turns out to be the multidimensional equivalent of the Huber loss of robust regression. We now discuss two approaches for solving Eq. (3), either exactly or approximately.

Exact solution of Eq. (3): The common proximal-gradient approach (Bach et al. 2011; Beck and Teboulle 2009b) for solving non-smooth convex problems has in this case the disadvantage of convergence time which grows linearly with the number of samples n. The reason is that the Lipschitz constant of the gradient of the first term in Eq. (3) grows linearly with n, which results in a decreasing step size. An alternative approach is to derive the dual optimization problem to Eq. (3), analogously to the derivation of (Beck and Teboulle 2009a) in the context of image denoising. The resulting objective is smooth and has a bounded Lipschitz constant independent of n. Yet another approach was proposed by (Bleakley and Vert 2011) for the task of change-point detection, who showed that under a suitable change of variables Eq. (3) can be formulated as a group-LASSO regression (Tibshirani 1996; Yuan and Lin 2006).

Approximate solution of Eq. (3): Two reasons suggest that deriving an alternative algorithm for solving Eq. (3) might have an advantage. First, the parameter λ does not allow direct control of the resulting number of segments, and in many use-cases such a control is a desired property. Second, as mentioned above, (Bleakley and Vert 2011) showed that Eq. (3) is equivalent to group-LASSO regression, under a suitable change of variables. It is known from the theory of LASSO regression that certain conditions on the design matrix must hold in order for perfect detection of segment boundaries to be possible. Unfortunately, these conditions are violated for the objective in Eq. (3) ; see (Levy-leduc and others 2007) and references therein. Therefore a non-exact solution has a potential of performing better, at least in some

situations. We indeed encountered this phenomenon empirically, as is demonstrated in Sec. 3. Therefore we also derive an alternative top-down, greedy algorithm, which finds a segmentation into K segments, where K is a user-controlled parameter. The algorithm works in rounds. On each round it picks a segment of a current segmentation, and finds the optimal segmentation of it into two subsequences. We start with the following lemma, which gives an analytical rule which solves Eq. (3) for the case of $K = 2$ segments.

Lemma 1 *Consider the optimal solution of Eq. (3) for the largest parameter λ for which there are $K = 2$ segments, and denote this value of the parameter by λ^*. Denote by i^* the associated splitting point into 2 segments, i.e. samples $\hat{\mathbf{x}}_i$ with $i \leq i^*$ belong to the first segment, and otherwise belong to the second segment. Then $i^*(\mathbf{x}) = \underset{1 \leq i \leq n-1}{\operatorname{argmax}} g(i, \mathbf{x})$, where:*

$$g(i, \mathbf{x}) = \left\{ \frac{i(n-i)}{w_i n} \left\| \bar{\mathbf{x}}_2(i) - \bar{\mathbf{x}}_1(i) \right\|_2 \right\}, \tag{4}$$

and $\bar{\mathbf{x}}_{1,2}(i)$ are the means of the first and second segments, respectively, given that the split occurs after the ith sample. In addition, $\lambda^*(\mathbf{x}) = g(i^*, \mathbf{x})$.

The proof is given in the supplementary material of (Katz and Crammer 2014). This result motivates a top-down hierarchical segmentation algorithm, which chooses at each iteration to split the segment which results in the maximal decrease of the sum-of-squared-errors criterion. Note that we cannot use the criterion of minimal increment to the objective in Eq. (3), since by continuity of the solution path, there is no change in the objective at the splitting from $K = 1$ to $K = 2$ segments. The top-down algorithm can be implemented in $\mathcal{O}(nK)$. Note that no search in the solution path is needed in case K is known, and this search can be made efficiently in case where K is not known. The top-down approach is used in the algorithm presented in Sec. 2.1.

From the functional form of $g(i, \mathbf{x})$ in Eq. (4) it is clear that in the unweighted case ($w_i = 1$ for all i), the solution is biased towards segments of approximately the same length, because of the $i(n-i)$ factor. We now show that a specific choice of w_i exactly recovers the solution to the unrelaxed optimization problem, where the regularization term in Eq. (3) is replaced with the ℓ_0 constraint, that is $\sum_{i=1}^{n-1} \mathbb{1}\{\|\boldsymbol{\mu}_{i+1} - \boldsymbol{\mu}_i\|_2 > 0\} = 1$. This is formulated by the following lemma:

Lemma 2 *Consider the case of two segments, $K = 2$, and denote by n_w^* the minimizer of Eq. (4) with $w_i = \sqrt{i(n-i)}$. Then the split into two segments found by solving the following:*

$$\begin{aligned} &\underset{\boldsymbol{\mu}}{\operatorname{argmin}} \left\{ \sum_{i=1}^{n} \|\mathbf{x}_i - \boldsymbol{\mu}_i\|^2 \right\}, \\ &\text{s.t.} \sum_{i=1}^{n-1} \mathbb{1}\{\|\boldsymbol{\mu}_{i+1} - \boldsymbol{\mu}_i\|_2 > 0\} = 1 \,, \end{aligned} \tag{5}$$

is also given by n_w^*.

The proof appears in the supplementary material of (Katz and Crammer 2014). We note that the same choice for w_i was derived by (Bleakley and Vert 2011) from different considerations based on a specific noise model. In this sense our derivation is more general, as it does not make any assumptions about the data.

Robust top-down algorithm We now propose a robust top-down algorithm for approximately optimizing Eq. (1). For a fixed value of $\boldsymbol{\mu}$, using Eq. (2) we can calculate analytically which of the $\boldsymbol{z}_i$s represent a detected outlier. These are $\boldsymbol{z}_i$s which satisfy $\|\boldsymbol{z}_i\|_2 > \gamma$. This allows us to calculate the value γ^* for which the first outlier is detected as having a non-zero norm. Furthermore, for $\lambda = \lambda^*,\ \gamma = \gamma^*$ we know that $\boldsymbol{z}_i = 0$ for all $i = 1, \ldots, n$, and therefore $\boldsymbol{\mu}_i = \operatorname{mean}(\mathbf{x} - \boldsymbol{z}) = \bar{\mathbf{x}}$ for all $i = 1, \ldots, n$, and we can find γ^* analytically:

$$\gamma^* = \max_i \{\|\mathbf{x}_i - \bar{\mathbf{x}}\|_2\} \,, \tag{6}$$

where the index i^* at which the maximum is attained is the index to the first detected outlier. The value of λ^* is found as given in Lemma 1, with the replacement of each $\mathbf{x}_i$ with $\hat{\mathbf{x}}_i$ as defined above. We note that the values $\lambda^*,\ \gamma^*$ are helpful for finding a solution path, since they allow to exclude parameters which result in trivial solutions.

In the case where $\lambda = \lambda^*$ we can extend Eq. (6) for any number $M > 1$ of outliers, by simply looking for the first M vectors $\mathbf{x}_i - \bar{\mathbf{x}}$ with the largest norm. In this case it no longer holds true that $\boldsymbol{z}_i = 0$ for all $i = 1, \ldots, n$, so we have to use the alternating optimization in order to find a solution. However, each iteration is now solved analytically and convergence is fast compared to the case $\lambda < \lambda^*$ where we do not have an analytical solution for the optimization over $\boldsymbol{\mu}$. This result motivates the top-down version of the ORCS algorithm. We denote the algorithm by TD-ORCS for the unweighted case ($w_i = 1$, $i = 1, \ldots, n$), and by WTD-ORCS when using the weights given in Lemma 2. The number of required segments K and number of required outliers M is set by the user. In each iteration the algorithm chooses the segment-split which results in the maximal decrease in the squared loss. Whenever a segment is split, the number of outliers belonging to each sub-segment is kept and used in the next iteration, so the overall number of outliers equals M at all iterations. The algorithm is summarized in Alg. 1.

2.2 Analysis of Lemma 1 for K=2

We now bound the probability that the solution as given by Lemma 1 fails to detect the correct boundary. We use the derived bound to show that the weights given in Lemma 2 are optimal in a sense explained below. For simplicity we analyze the one dimensional case $\mathbf{x}_i \in \mathbb{R}$, and we show later how the results generalize to multidimensional data.

We assume now that the data sequence is composed of two subsequences of lengths n_1 and n_2, each composed of samples taken iid from some probability distributions with means μ_1 and μ_2 respectively, and define $\Delta\mu \triangleq \mu_2 - \mu_1$. We further assume that the samples are bounded, i.e. $|\mathbf{x}_i| \leq M/2,\ i = 1, \ldots, n$, for some positive constant M. We set $w_i = 1$ for all $i = 1, \ldots, n-1$, and quote results for the

Algorithm 1 Top-down outlier-robust hierarchical segmentation

Input: Data samples $\{\mathbf{x}_i\}_{i=1}^n$.
Parameters: Number of required segments K, number of required outliers M, weights $\{w_i\}_{i=1}^{n-1}$.
Initialize: Segments boundaries $\mathrm{B} = \{1, n\}$, current number of segments $k = |B| - 1$, number of outliers for each segment $M_j = M$ (for $j = 1$).
while $k < K$ **do**
  for $j = 1, ..., k$ **do**
    Set $S_j = \left\{\mathbf{x}_{B_j}, .., \mathbf{x}_{B_{j+1}-1}\right\}$.
    while not converged **do**
      Split segment S_j into sub-sequences $S_{1,2}$, using the solution to Eq. (4) at $\mathbf{x} - \boldsymbol{z}$, with weights $\{w_i\}$.
      Find γ for M_j outliers, using the extension of Eq. (6) to M_j outliers.
      Set $\boldsymbol{z}_i = \mathrm{prox}_\gamma\left(\mathbf{x}_i - \bar{\mathbf{x}}\right)$, for $i = 1, \ldots, n$.
    end while
    Calculate the mean $\bar{\mathbf{x}}_j$ of segment S_j, and the means of $S_{1,2}$, denoted by $\bar{\mathbf{x}}_1$ and $\bar{\mathbf{x}}_2$.
    Set $L(j) = \sum_{i \in S_1} \|\mathbf{x}_i - \bar{\mathbf{x}}_1\|^2 + \sum_{i \in S_2} \|\mathbf{x}_i - \bar{\mathbf{x}}_2\|^2 - \sum_{i \in S_j} \|\mathbf{x}_i - \bar{\mathbf{x}}_j\|^2$.
  end for
  Choose segment $j^* = \arg\max_j L(j)$ and calculate the splitting point n_{j^*}
  Add the new boundary $\mathbf{x}_{B_{j^*}} + n_{j^*} + 1$ to the (sorted) boundary list B.
  Set M_j and M_{j+1} to the number of $\boldsymbol{z}_i$s with non-zero norm in segment S_j and S_{j+1}, respectively.
end while

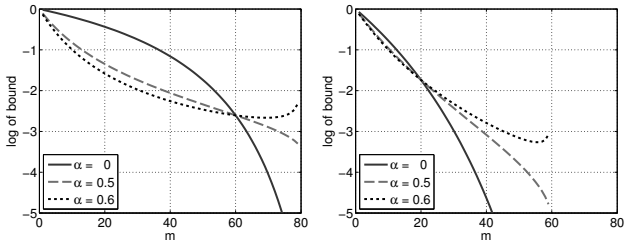

(a) $n_1 = 20$ (b) $n_1 = 40$

Figure 1: Bounds on $\mathbb{P}\left(g(0) < g(m)\right)$ as a function of the distance m from the true boundary n_1, for $n = 100$, two values of n_1, and various values of the weighting parameter α. The case of $\alpha = 0$ amounts to uniform weights.

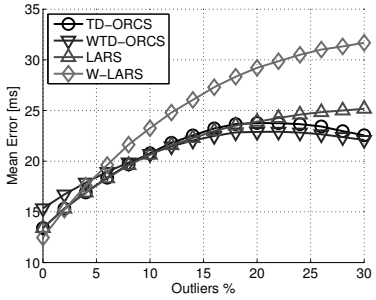

Figure 2: Mean error (absolute time difference to ground-truth boundary) of four algorithms on TIMIT phonetic data contaminated with outliers. See text for a list of algorithms compared.

weighted case where relevant. We note that n_1 and n_2 represent the ground-truth, and not a variable we have to optimize. We parameterize the sample-index argument of $g\left(\cdot\right)$ in Eq. (4) as $i = n_1 + m$ (and similarly $i^* = n_1 + m^*$), that is we measure it relatively to the true splitting point n_1. For ease of notation, in what follows we substitute $g\left(m\right)$ for $g\left(n_1 + m\right)$. Without loss of generality, we treat the case where $m \geq 0$. Note that $m^* \neq 0$ if $g(0) < g(m)$ for some $m > 0$. The probability of this event is bounded:

$$\mathbb{P}\left(g(0) < g(m)\right) \leq 2\exp\left(-Cm\right), \tag{7}$$

for $C = \left(2\Delta\mu^2 n_1^2\right) / \left(M^2 n^2\right)$. The proof is given in the supplementary material of (Katz and Crammer 2014).

Note that in order for the bound to be useful, the true segments should not be too long or too short, in agreement with the motivation for using weights given before Lemma 2. We now use Eq. (7) to prove the following theorem:

Theorem 3 *Consider a sequence of n variables as described above. Given $\delta \in (0, 1)$ set $m_0 = \frac{\log(2n_2/\delta)}{C}$. Then, the probability that the solution $i^* = n_1 + m^*$ as given in Lemma 1 is no less than m_0 samples away from the true boundary is bounded, $\mathbb{P}\left(m^* \geq m_0\right) \leq \delta$.*

The proof appears in the supplementary material of (Katz and Crammer 2014).

Considering the weighted case with arbitrary w_i, we repeated the calculation for the bound on $\mathbb{P}\left(g(0) < g(m)\right)$. To illustrate the influence of the weights on the bound, we heuristically parameterize $w_j = \left(j\left(n - j\right)\right)^\alpha$ for some $\alpha \in [0, 1]$[1]. The bound as a function of m is illustrated in Fig. 1 for several values of α. It is evident that indeed $\alpha > 0$ achieves a faster decaying bound for small n_1, and that $\alpha = 0.5$ is optimal in the sense that for $\alpha > 0.5$ the bound is no longer a monotonous function of m. This agrees with the weights given by Lemma 2. We note that this specific choice of the weights was derived as well by (Bleakley and Vert 2011) by assuming a Gaussian noise model, while our derivation is more general. Finally, we note that generalizing the results for multidimensional data is done by using the fact that for any two vectors $a, b \in \mathbb{R}^d$, it holds that $\mathbb{P}\left(\|a\|_2 \leq \|b\|_2\right) \leq \sum_i \mathbb{P}\left(|a_i| \leq |b_i|\right)$. Thus generalizing Eq. (7) for $d > 1$ is straightforward. While the bound derived in this way will have a multiplicative factor of d, it is still exponential in the number of samples n.

3 Empirical Study

We compared the unweighted (TD-ORCS) and weighted (WTD-ORCS) versions of our top-down algorithm to LASSO and group fused LARS [2] of (Bleakley and Vert 2011), which are based on reformulating Eq. (3) as group LASSO regression, and solving the optimization problem either exactly or approximately. Both the TD-ORCS and

[1] This parametrization is motivated by Eq. (4)
[2] http://cbio.ensmp.fr/ jvert/svn/GFLseg/html/

LARS algorithms have complexity of $\mathcal{O}(nK)$. We also report results for a Bayesian change-point detection algorithm (BCP), as formulated by (Erdman and Emerson 2008). We note that we experimented with a left-to-right Hidden Markov Model (HMM) for segmentation. We do not report results for this model, as its performance was inferior to the other baselines.

3.1 Biphones subsequences segmentation

In this experiment we used the TIMIT corpus (Garofolo and others 1988). Data include $4,620$ utterances with annotated phoneme boundaries, amounting to more than $170,000$ boundaries. Audio was divided into frames of 16ms duration and 1 ms hop-length, each represented with 13 MFCC coefficients. The task is to find the boundary between two consecutive phonemes (biphones), and performance is evaluated as the mean absolute distance between the detected and ground-truth boundaries. Since the number of segments is $K = 2$ the ORCS and the TD-ORCS algorithms are essentially identical, and the same holds for LASSO and LARS. Outliers were incorporated by adding short (0.25 frame-length) synthetic transients to the audio source. The percentage of outliers reflects the percentage of contaminated frames. Results are shown in Fig. 2 as the mean error as a function of outlier percentage. For low fraction of outliers, all algorithms perform the same, except WTD-ORCS, which is slightly worse. For about 15% outliers, the performance of W-LARS degrades to ~27ms mean error vs ~22ms for the rest. For 30% outliers both TD-ORCS algorithms outperform all other algorithms. The counter-intuitive drop of error at high outliers rate for the TD-ORCS algorithms might be the result of over-estimating the number of outliers. We plan to further investigate this phenomenon in future work.

We also compared our algorithm to RD, which (Qiao, Shimomura, and Minematsu 2008) found to be the best among five different objective functions, and was not designed for treating outliers. In this setting (no outliers) the RD algorithm achieved 15.1ms mean error, while TD-ORCS achieved 13.4ms, with 95% confidence interval (not reported for the RD algorithm) of 0.1.

3.2 Radio show segmentation

In this experiment we used a 35 minutes, hand-annotated audio recording of a radio talk show, composed of different sections such as opening title, monologues, dialogs, and songs. A detected segment boundary is considered a true positive if it falls within a tolerance window of two frames around a ground-truth boundary. Segmentation quality is commonly measured using the F measure, which is defined as $2pr/(p+r)$, where p is the precision and r is the recall. Instead, we used the R measure introduced by (Räsänen, Laine, and Altosaar 2009), which is more robust to over-segmentation. It is defined as $R \triangleq 1 - 0.5\left(|s_1| + |s_2|\right)$, where $s_1 \triangleq \sqrt{(1-r)^2 + (r/p-1)^2}$ and $s_2 \triangleq (r - r/p)/\sqrt{2}$. The R measure satisfies $R \leq 1$, and $R = 1$ only if $p = r = 1$.

Signal representation A common representation in speech analysis is the MFCC coefficients mentioned in Sec. 3.1. However, this representation is computed over time windows of tens of milliseconds, and therefore it is not designed to capture the characteristics of a segment with length in the order of seconds or minutes. We therefore apply post-processing on the MFCC representation. First, the raw audio is divided into N non-overlapping blocks of 5 seconds duration, and the MFCC coefficients are computed for all blocks $\{S_j\}_{j=1}^{N}$. We used 13 MFCC coefficient with 25ms window length and 10ms hop length. Then a Gaussian Mixture Model (GMM) T_j with 10 components and a diagonal covariance matrix is fitted to the jth block S_j. These parameters of the GMM were selected using the Bayesian Information Criterion (BIC). The log-likelihood matrix A is then defined by $A_{ij} = \log \mathbb{P}(S_j|T_i)$. The clean feature matrix (no outliers) is shown in Fig. 3(a), where different segments can be discerned. Since using the columns of A as features yields a dimension growing with N, we randomly choose a subset of $d = 100$ rows of A, and the columns of the resulting matrix $X \in \mathbb{R}^{d \times N}$ are the input to the segmentation algorithm. We repeat the experiment for different number of outliers, ranging between 0% and 16% with intervals of 2%. Outliers were added to the raw audio. A given percentage of outliers refers to the relative number of blocks randomly selected as outliers, to which we add a 5 seconds recording of repeated hammer strokes, normalized to a level of 0dB SNR.

Algorithms We consider the Outlier-Robust Convex Sequential (ORCS) segmentation, and its top-down versions (weighted and unweighted) which we denote by WTD-ORCS and TD-ORCS, respectively. We compare the performance to three other algorithms. The first is a greedy bottom-up (BU) segmentation algorithm, which minimizes the sum of squared errors on each iteration. The bottom-up approach has been successfully used in tasks of speech segmentation (Qiao, Luo, and Minematsu 2012; Gracia and Binefa 2011). The second algorithm is the W-LARS algorithm of (Bleakley and Vert 2011). The third algorithm is a Bayesian change-point detection algorithm (BCP), as formulated by (Erdman and Emerson 2008). A solution path was found as follows. For the ORCS algorithm, a 35×35 parameter grid was used, where $0 < \gamma < \gamma^*$ was sampled uniformly, and for each γ value, $0 < \lambda < \lambda^*(\gamma)$ was sampled logarithmically, where $\lambda^*(\gamma)$ is the critical value for λ for a given choice of γ (see Sec. 2.1 for details). For the TD-ORCS, W-LARS, and BU algorithms, $K = 2, \ldots, 150$ number of segments were used as an input to the algorithms. For the TD-ORCS algorithm, where the number of required outliers is an additional input parameter, the correct number of outliers was used. For the BCP algorithm, a range of thresholds on the posterior probability of change-points was used to detect a range of number of segments. As is evident from the empirical results below, the ORCS algorithm can achieve high detection rate of the outliers even without knowing their exact number a-priori. Furthermore, we suggest below a way of estimating the number of outliers. For each algorithm, the maximal R measure over all parameters

range was used to compare all algorithms.

Results Results are shown in Fig. 3(b) as the maximal R measure achieved versus the percentage of outliers, for each of the algorithms considered. It is evident that the performance of the BU and BCP algorithms decreases significantly as more outliers are added, while the outlier-robust ORCS algorithm keeps an approximately steady performance. Our unweighted and weighted TD-ORCS algorithms achieve the best performance for all levels of outliers. Results for LARS algorithm are omitted as it did not perform better than other algorithms. We verified the ability of our algorithms to correctly detect outliers by calculating the R measure of the outliers detection of the ORCS algorithm, with zero length tolerance window, i.e a detection is considered a true-positive only if it exactly pinpoints an outlier. The R measure of the detection was evaluated on the γ, λ parameter grid, as well as the corresponding numbers of detected outliers. Results for the representative case of $p = 10\%$ outliers are shown in Fig. 4(a). It is evident that a high R measure (> 0.9) is attained on a range of parameters that yield around the true number of outliers. We conclude that one does not need to know the exact number of outliers in order to use the ORCS algorithm, and a rough estimate is enough. Some preliminary results suggest that such an estimate can be approximated from the histogram of number of detected outliers (i.e. Fig. 4(b)).

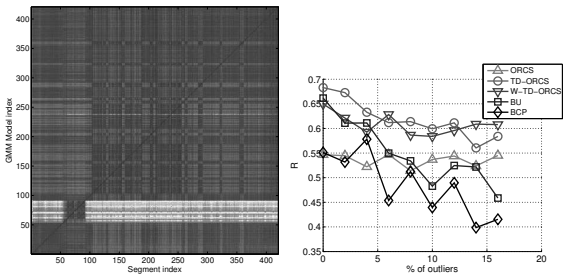

(a) Likelihood matrix (b) Maximal R vs. % of outliers

Figure 3: (a) A_{ij} is the log-probability of segment j given the GMM fitted to segment i. (b) Maximal R measure vs. percentage of outliers. See text for algorithms details.

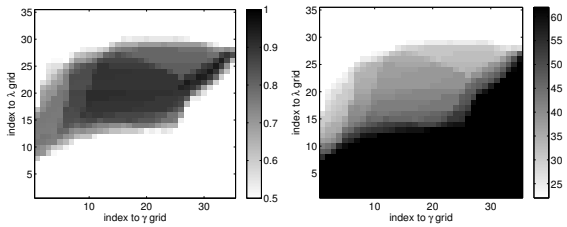

(a) R measure (b) Number of detected outliers

Figure 4: (a) Outliers detection quality by the ORCS algorithm, quantified using the R measure of outliers detection. R is evaluated on a 35×35 parameter grid, for 10% outliers. (b) The corresponding number of detected outliers.

4 Related Work and Conclusion

There is a large amount of literature on change-point detection, see for example (Basseville and Nikiforov 1993; Brodsky and Darkhovsky 1993). Optimal segmentation can be found using dynamic programming (Lavielle and Teyssière 2006); however, the complexity is quadratic in the number of samples n, which might be infeasible for large data sets. Some approaches with linear complexity in n (Levy-leduc and others 2007; Killick, Fearnhead, and Eckley 2012) treat only one dimensional data. Some related work is concerned with the objective in Eq. (3). In (Levy-leduc and others 2007) it was suggested to reformulate Eq. (3) for the one dimensional case as a LASSO regression problem (Tibshirani 1996; Yuan and Lin 2006), while (Bleakley and Vert 2011) extended this approach to multidimensional data, although not treating outliers directly. Another common approach is deriving an objective from a maximum likelihood criterion of a generative model (Qiao, Luo, and Minematsu 2012; Qiao, Shimomura, and Minematsu 2008; Gracia and Binefa 2011). We note that the two-dimensional version of Eq. (3) is used in image denoising applications, where it is known as the Total-Variation of the image (Rudin, Osher, and Fatemi 1992; Chambolle 2004; Beck and Teboulle 2009a). Finally, we note that all these approaches do not directly incorporate outliers into the model.

We formulated the task of segmenting sequential data and detecting outliers using convex optimization, which can be solved in an alternating manner. We showed that a specific choice of weighting can empirically enhance performance. We described how to calculate λ^* and γ^*, the critical values for the split into two segments and the detection of the first outlier, respectively. We also derived a top-down, outlier-robust algorithm which minimizes the objective in a greedy manner. This algorithm allows for directly controlling both the number of segments K and number of outliers M. Experiments with real-world audio data with outliers added manually demonstrated the superiority of our algorithms.

We consider a few possible extensions to the current work. One is deriving algorithms that will work on-the-fly. Another direction is to investigate more involved noise models, such as noise which corrupts a single feature along all samples, or a consecutive set of samples. Yet another interesting question is how to identify that different segments come from the same source, e.g. that the same speaker is present at different segments. We plan to investigate these directions in future work.

References

Bach, F.; Jenatton, R.; Mairal, J.; and Obozinski, G. 2011. Convex optimization with sparsity-inducing norms. *Optimization for Machine Learning* 19–53.

Basseville, M., and Nikiforov, I. V. 1993. Detection of abrupt changes: theory and application.

Beck, A., and Teboulle, M. 2009a. Fast gradient-based algorithms for constrained total variation image denoising and

deblurring problems. *Image Processing, IEEE Transactions on* 18(11):2419 –2434.

Beck, A., and Teboulle, M. 2009b. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J. Img. Sci.* 2(1).

Beeferman, D.; Berger, A.; and Lafferty, J. 1999. Statistical models for text segmentation. *Machine learning* 34(1-3):177–210.

Bleakley, K., and Vert, J.-P. 2011. The group fused lasso for multiple change-point detection. *arXiv preprint arXiv:1106.4199*.

Brent, M. R. 1999. Speech segmentation and word discovery: A computational perspective. *Trends in Cognitive Sciences* 3(8):294–301.

Brodsky, B. E., and Darkhovsky, B. S. 1993. *Nonparametric methods in change point problems*, volume 243. Springer.

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

Chi, E. C., and Lange, K. 2013. Splitting Methods for Convex Clustering. *ArXiv e-prints*.

Erdman, C., and Emerson, J. W. 2008. A fast bayesian change point analysis for the segmentation of microarray data. *Bioinformatics* 24(19):2143–2148.

Forero, P. A.; Kekatos, V.; and Giannakis, G. B. 2011. Outlier-aware robust clustering. In *Acoustics, Speech and Signal Processing (ICASSP), 2011 IEEE International Conference on*, 2244–2247. IEEE.

Garofolo, J. S., et al. 1988. Getting started with the darpa timit cd-rom: An acoustic phonetic continuous speech database. *National Institute of Standards and Technology (NIST), Gaithersburgh, MD* 107.

Gracia, C., and Binefa, X. 2011. On hierarchical clustering for speech phonetic segmentation. In *Eusipco*.

Katz, I., and Crammer, K. 2014. Outlier-robust convex segmentation. *arXiv preprint*. arXiv:1411.4503v1.

Killick, R.; Fearnhead, P.; and Eckley, I. 2012. Optimal detection of changepoints with a linear computational cost. *Journal of the American Statistical Association* 107(500):1590–1598.

Lavielle, M., and Teyssière, G. 2006. Detection of multiple change-points in multivariate time series. *Lithuanian Mathematical Journal* 46(3):287–306.

Levy-leduc, C., et al. 2007. Catching change-points with lasso. In *Advances in Neural Information Processing Systems*, 617–624.

Lévy-Leduc, C., and Roueff, F. 2009. Detection and localization of change-points in high-dimensional network traffic data. *The Annals of Applied Statistics* 637–662.

Mateos, G., and Giannakis, G. B. 2012. Robust PCA as bilinear decomposition with outlier-sparsity regularization. *Signal Processing, IEEE Transactions on* 60(10):5176–5190.

Olshen, A. B.; Venkatraman, E.; Lucito, R.; and Wigler, M. 2004. Circular binary segmentation for the analysis of array-based dna copy number data. *Biostatistics* 5(4):557–572.

Qiao, Y.; Luo, D.; and Minematsu, N. 2012. A study on unsupervised phoneme segmentation and its application to automatic evaluation of shadowed utterances. Technical report.

Qiao, Y.; Shimomura, N.; and Minematsu, N. 2008. Unsupervised optimal phoneme segmentation: Objectives, algorithm and comparisons. In *ICASSP*.

Räsänen, O. J.; Laine, U. K.; and Altosaar, T. 2009. An improved speech segmentation quality measure: the r-value. In *INTERSPEECH*.

Rudin, L. I.; Osher, S.; and Fatemi, E. 1992. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena* 60(1):259 – 268.

Shriberg, E.; Stolcke, A.; Hakkani-Tü, D.; and Tür, G. 2000. Prosody-based automatic segmentation of speech into sentences and topics. *Speech Communication* 32(1–2).

Tibshirani, R. 1996. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)* 267–288.

Yuan, M., and Lin, Y. 2006. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 68(1):49–67.