

Can automated smoothing significantly improve benchmark time series classification algorithms?

James Large, Paul Southam, and Anthony Bagnall

University of East Anglia, Norwich Research Park, UK
{James.Large, Paul.Southam, ajb}@uea.ac.uk
<http://www.timeseriesclassification.com>

Abstract. tl;dr: no, it cannot, at least not on average on the standard archive problems. We assess whether using six smoothing algorithms (moving average, exponential smoothing, Gaussian filter, Savitzky-Golay filter, Fourier approximation and a recursive median sieve) could be automatically applied to time series classification problems as a preprocessing step to improve the performance of three benchmark classifiers (1-Nearest Neighbour with Euclidean and Dynamic Time Warping distances, and Rotation Forest). We found no significant improvement over unsmoothed data even when we set the smoothing parameter through cross validation. We are not claiming smoothing has no worth. It has an important role in exploratory analysis and helps with specific classification problems where domain knowledge can be exploited. What we observe is that the automatic application does not help and that we cannot explain the improvement of other time series classification algorithms over the baseline classifiers simply as a function of the absence of smoothing.

1 Introduction

Time Series Classification (TSC) is differentiated from standard classification by the fact that the ordering of the attributes may be important in finding discriminatory features. Standard vector classifiers such as rotation forest and standard time dependent approaches such dynamic time warping with 1-NN are strong benchmark algorithms to compare against the range of bespoke TSC algorithms that have been proposed in recent years. Some of these achieve impressive performance and are significantly better than the benchmarks. Nevertheless, there has always been a suspicion that sensible standard preprocessing of the data would perhaps increase the accuracy of benchmark classifiers and that would make at least some of the bespoke algorithms redundant [1]. Broadly speaking, there are four types of preprocessing that may improve classifier performance: normalisation; smoothing; dimensionality reduction; and discretization. We address the question of whether smoothing series can significantly improve the accuracy of benchmark classifiers. Smoothing is the process of reducing the noise in the series to make patterns in the data more apparent and is generally used as part of an exploratory analysis.

It is important to stress we are only concerned with class independent noise, since class dependent noise is possibly useful as a discriminatory feature. This is where we diverge from the majority of signal processing research into noise modeling and reduction. We are not necessarily trying to “clean up” a signal. Instead, we are trying to remove artifacts that may confound the classifier.

We test whether six smoothing algorithms improve three base classifiers. These are described in detail in Section 2. It is clearly important to set the parameters of the algorithm when smoothing so that it is relevant to a specific problem. Because we are attempting to smooth to improve classification, we set parameters through cross validation on the train data using the base classifier we are testing. The experimental design is described in Section 3. Our experiments address the following two questions.

1. Does smoothing with default parameters increase the accuracy of benchmark classifiers?
2. Can we learn smoothing parameters on the train data to significantly improve benchmark TSC algorithms?

A priori, we believed it unlikely that systematic smoothing would improve accuracy over the diverse data sets in the archive, since many of the series have very little noise. However, we thought that supervised smoothing, where no smoothing was an option, would improve performance albeit at the large computational cost of the parameter search. Our results, presented in Section 3 show that in fact smoothing makes very little difference, even when supervised. We discuss these results in Section 5.

2 Background

2.1 Time Series Classification

A large number of new classification problems have been proposed in the last ten years. While not exhaustive by itself, it is important to evaluate new algorithms against sensible benchmark classifiers on standard test problems in order to ascertain the usefulness of new research. The UEA-UCR archive is a widely used archive of test problems [5]. The archive is a continually growing collection of real valued TSC datasets¹ which come from a range of different domains and have a range of characteristics, in terms of size, number of classes, imbalances, etc. Most TSC publications benchmark against a 1-NN classifier using either Euclidean distance (ED) or Dynamic Time Warping (DTW) distance. DTW compensates for potential misalignments amongst series of the same class. DTW has a single parameter, the maximum warping window, and DTW performs significantly better when this parameter is set through cross validation. A recent comparative study [2] found that the classifier rotation forest [9] was also a strong benchmark. It is able to discover relationships in time through the internal principle component transformation it uses, and is not significantly worse than

¹ <http://www.timeseriesclassification.com>

DTW with window set through cross validation (DTWCV henceforth). The same study compared 22 TSC algorithms on 85 of the UEA-UCR archive data and found that just nine out of twenty two TSC algorithms were significantly more accurate than both a rotation forest and DTW classifier. Some of these TSC algorithms are highly complex and both memory and computationally expensive. A case was made that the superior algorithms achieved higher accuracy because the representation they use allows for the detection of discriminatory features that the benchmarks cannot find. This was further demonstrated on the archive and through data simulation [7]. We wish to test whether simple preprocessing can significantly improve the benchmarks and hence narrow the gap between DTW and rotation forest and the nine significantly better TSC algorithms.

2.2 Time Series Smoothing

Given a time series $T = \langle t_0, \dots, t_{m-1} \rangle$, a smoothing function produces a new series $S = \langle s_0, \dots, s_{p-1} \rangle$, where $p \leq m$ (we index from zero to make the equations simpler). Most algorithms employ a sliding window, of length w , along the series, resulting in a series of length $p = m - w$. The simplest form of smoothing is to take the **moving average (MA)**.

$$s_j = \frac{\sum_{i=j-w}^j t_i}{w} \quad \text{for } j = w \dots m - 1,$$

where w is the single parameter, window size. **Exponential smoothing (EXP)** is a generalisation of moving average smoothing that assigns a decaying weight to each element rather than averaging over a window.

$$s_0 = t_0 \quad \text{and} \quad s_j = \alpha \cdot t_j + (1 - \alpha) \cdot t_{j-1}$$

where $0 \leq \alpha \leq 1$. For consistency with other smoothing algorithms, EXP is often given a window size w , then the decay weight is set as $\alpha = \frac{2}{w+1}$.

A **Gaussian filter (GF)** applies a fixed convolution over a window

$$s_j = \sum_{i=j-w}^j t_i \cdot c_i,$$

where the convolution values c_i are derived from a standard normal distribution over the window w , the single parameter.

Like GF, the **Savitzky-Golay (SG)** filtering method is a convolutional method of smoothing. Instead of using a fixed convolution, it estimates a different convolution on each window based on local least-squares polynomial approximation.

$$s_j = \sum_{i=j-w}^j t_i \cdot c_{i,j}$$

Since its initial introduction [10], it has been used successfully and pervasively across many signal processing domains for different purposes, particularly in chemometrics [4, 6]. SG has two parameters, window size w and polynomial order n . For accessible explanations of how the polynomial coefficients are calculated, we refer the reader to [11].

Discrete Fourier Approximation (DFT) smooths the series by first transforming into the frequency domain, discarding the high frequency terms, then transforming back to the time domain. DFT has a single parameter, r , the proportion of Fourier terms to retain.

The **Recursive Median Sieve (SIV)** is a one-dimensional recursive median filter [3] that filters the data by removing extrema of specific scales.

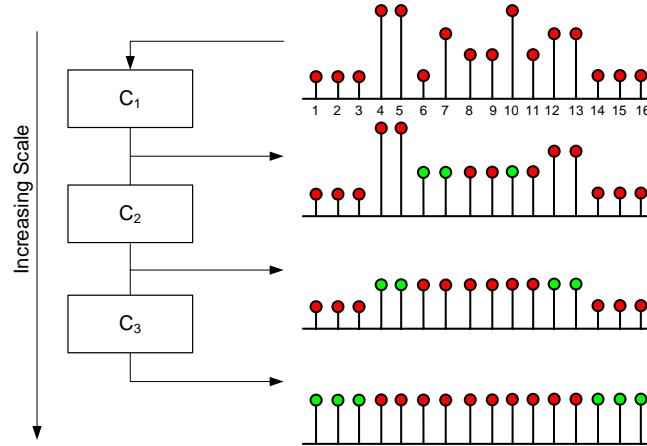


Fig. 1. An example sieve decomposition of a 1D signal. Green vertices are the vertices affected at each scale level.

The sieve performs a decomposition removing extrema (both maxima and minima) at different scales as shown in Figure 1. At the scale c_1 the maxima and minima at points 6,7 and 10 are smoothed to equal the nearest value of the neighbours. At scale c_2 the pairs at (4,5) and (12,13) are smoothed. At the highest scale, the series is uniform. The sieve takes in a single parameter, c , which is the scale to smooth the signal to.

For each of pair of filter+classifier combination, we perform 10 stratified random resamples of each data set and report the average results across those resamples. The first resample, fold 0, is always the exact train/test split published on the UCR-UEA archive, to allow for easier comparison. To avoid ambiguity, we stress that in all cases the training of a classifier, including any parameter tuning and model selection required, is performed independently on the train set of a given fold, and the trained classifier is evaluated exactly once on the corresponding

Table 1. The parameter spaces searched for each filtering method over the course of our experiments (default value in bold). m is the series length. For Savitzky-Golay (SG), all combinations of w, n are searched where $w > 2n$.

Method	Parameters and default values in bold
Moving Average (MA)	$w \in \{2, 3, \mathbf{5}, 10, 25, 50, 100, \sqrt{m}, \log_2(m)\}$
Exponential Smoothing (EXP)	$w \in \{2, 3, \mathbf{5}, 10, 25, 50, 100, \sqrt{m}, \log_2(m)\}$
Gaussian Filtering (GF)	$w \in \{2, 3, \mathbf{5}, 10, 25, 50, 100, \sqrt{m}, \log_2(m)\}$
Savitzky-Golay (SG)	$w \in \{\mathbf{5}, 9, 17, 33, 65\}$ $n \in \{\mathbf{2}, 3, 4, 8, 16, 32\}$
Fourier Approximation (DFT)	$r \in \{0.01, 0.05, \mathbf{0.1}, 0.25, 0.5, \log_2(m)/m\}$
Sieve (SIV)	$c \in \{\frac{1}{15} \cdot \log_{10}(m), \dots, \frac{5}{15} \cdot \log_{10}(m), \dots, \log_{10}(m)\}$

test set. We conduct 10 resamples on 76 of the 85 UCR archive TSC problems. We have omitted the largest problems due to time constraints. We average test accuracy over the 10 resamples, then present results in critical difference diagrams, which display the average ranks of the classifiers over all problems and group classifiers into cliques, within which there is no significant difference. For each resample, we perform a 10 fold cross validation (CV) on that resamples’ train data to find smoothing parameters. Our code² reproduces the splits used in this evaluation exactly, and full, reproducible results are available³. For all smoothing algorithms except the sieve, we used the standard MATLAB implementations and performed the smoothing and classification in separate stages. The default parameters given in Table 1 are those of the Matlab implementations. The sieve is implemented in C and was similarly isolated from the classification stage.

We use three baseline classifiers. 1-NN with Euclidean distance is a weak baseline, but it is still frequently used in research. 1-NN with DTW is the most common benchmark, although it is important to set the window through cross validation [8]. This is computationally expensive, although we use the DTW version described in [12] which speeds up the calculation by orders of magnitude. All the UEA-UCR data are normalised. For consistency, we renormalise each series after smoothing.

3 Results

In Figure 2 we present results for three baseline classifiers with both default smoothing and tuned smoothing. For all three classifiers, smoothing of any kind provides no benefit. Tuning provides no benefit over using default values, and in many cases makes things worse due to overfitting.

For all six experiments, the classifiers built on unsmoothed data are in the top clique. For four of the experiments, the unsmoothed classifier is the highest ranked. Setting the parameter through cross validation is if anything worse than using a default parameter. Given the order of magnitude more computation

² <https://bitbucket.org/TonyBagnall/time-series-classification>

³ <http://research.cmp.uea.ac.uk/SmoothingAALTD18/>

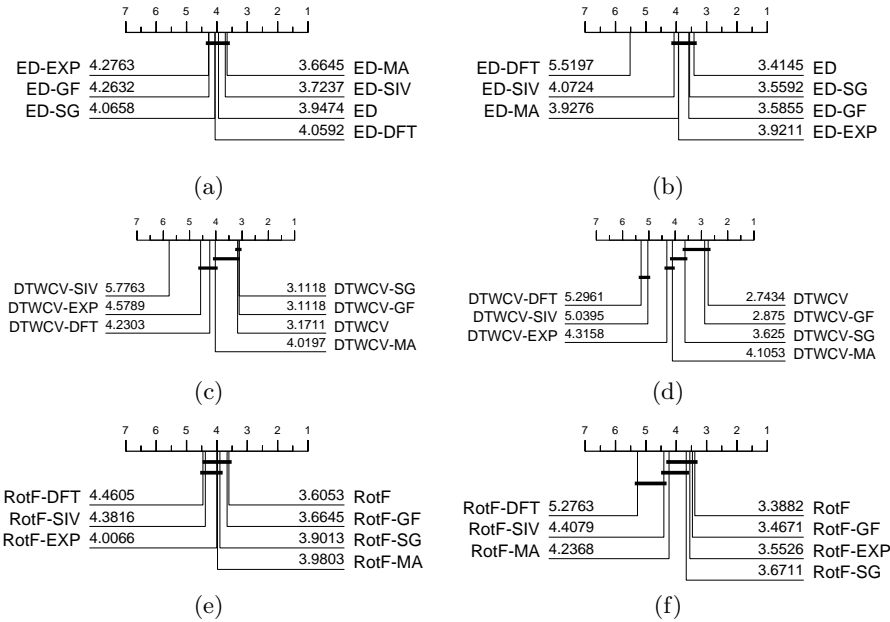


Fig. 2. Average ranks on 76 UEA-UCR problems without smoothing (left) and with six types of smoothing described in Section 2 (right), using the default parameters given in Table 1 and tuned over the parameter range given in Table 1 for Euclidean distance ((a) and (b)), dynamic time warping ((c) and (d)) and rotation forest ((e) and (f)).

required to tune these parameters, this is surprising, particularly as *no smoothing* was one of the options. Further analysis shows that *no smoothing* was selected approximately 25% of the time. This could be an indication that the archive data are simply not suited to smoothing.

4 Analysis

We examine whether there are any characteristics of the data that could help determine whether any of the six types of smoothing would improve performance. We would expect that smoothing might be more useful for longer series. Figures 3 and 4 show the scatter plot of length against classifier rank for DTWCV and rotation forest. There is no obvious relationship between the performance of the unsmoothed classifier and series length.

5 Conclusion

It has long been a suspicion of many researchers in this field that much of the improvement seen in complex TSC algorithms could equally be achieved with




Fig. 3. Ranks on 76 UEA-UCR problems of DTWCV compared to six untuned smoothing versions plotted against series length.




Fig. 4. Ranks on 76 UEA-UCR problems of rotation forest compared to six untuned smoothing versions plotted against series length.

comparatively simple preprocessing. Our experiments indicate for the case of smoothing, this is not true. We have taken six very popular smoothing algorithms and applied them using sensible default parameters and using extensive extra computation to discover optimal parameters through cross validation. We have found no significant difference between smoothed and unsmoothed classification with three benchmarks. The nature of the UCR data may explain this to a degree:

the data from problems such as image processing will have less noise than, for example, financial data. We are not claiming that smoothing has no role to play in TSC, merely that the application of smoothing does not on average improve the performance of baselines and that the absence of smoothing cannot explain the performance of algorithms that outperform the baselines.

Acknowledgement. This work is supported by the UK Engineering and Physical Sciences Research Council (EPSRC) [grant number EP/M015807/1] and Biotechnology and Biological Sciences Research Council [grant number BB/M011216/1]. The experiments were carried out on the High Performance Computing Cluster supported by the Research and Specialist Computing Support service at the University of East Anglia and using a Titan X Pascal donated by the NVIDIA Corporation.

References

1. Y. Chen B. Hu and E. Keogh. Time series classification under more realistic assumption. In *Proc. 13th SIAM International Conference on Data Mining*, 2013.
2. A. Bagnall, J. Lines, A. Bostrom, J. Large, and E. Keogh. The great time series classification bake off: a review and experimental evaluation of recent algorithmic advances. *Data Mining and Knowledge Discovery*, 31(3):606–660, 2017.
3. J. Bangham. Data-sieving hydrophobicity plots. *Analytical biochemistry*, 174(1):142–145, 1988.
4. G. Betta, D. Capriglione, G. Cerro, L. Ferrigno, and G. Miele. The effectiveness of Savitzky-Golay smoothing method for spectrum sensing in cognitive radios. *Proceedings of the 2015 18th AISEM Annual Conference*, (4):1–4, 2015.
5. Y. Chen, E. Keogh, B. Hu, N. Begum, A. Bagnall, A. Mueen, and G. Batista. The UEA-UCR time series classification archive. http://www.cs.ucr.edu/~eamonn/time_series_data/, 2015.
6. B. Fernandes, G. Colletta, L. Ferreira, and O. Dutra. Utilization of Savitzky-Golay filter for power line interference cancellation in an embedded electrocardiographic monitoring platform. *Proc. IEEE International Symposium on Medical Measurements and Applications*, pages 7–12, 2017.
7. J. Lines, S. Taylor, and A. Bagnall. HIVE-COTE: The hierarchical vote collective of transformation-based ensembles for time series classification. In *Proc. IEEE International Conference on Data Mining*, 2016.
8. C. Ratanamahatana and E. Keogh. Three myths about dynamic time warping data mining. In *Proc. 5th SIAM International Conference on Data Mining*, 2005.
9. J. Rodriguez, L. Kuncheva, and C. Alonso. Rotation forest: A new classifier ensemble method. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(10):1619–1630, 2006.
10. A. Savitzky and M. Golay. Smoothing and differentiation of data by simplified least squares procedures. *Analytical Chemistry*, 36(8):1627–1639, 1964.
11. R. Schafer. What is a Savitzky-Golay Filter? *IEEE Signal Processing Magazine*, 28(4):111–117, 2011.
12. C. Tan, M. Herrman, G. Forestier, G. Webb, and F. Petitjean. Efficient search of the best warping window for dynamic time warping. In *Proc. 18th SIAM International Conference on Data Mining*, 2018.