Predicting the vibroacoustic quality of steering gears

Paul Alexandru Bucur *,[1], Klaus Frick[2], and Philipp Hungerländer[1]

[1]Department of Mathematics, Alpen-Adria-Universität Klagenfurt, Universitätsstraße 65-67, Klagenfurt 9020, Austria
[2]Institute of Computational Engineering, NTB, Werdenbergstraße 4, Buchs 9471, Switzerland

Abstract

In the daily operations of ThyssenKrupp Presta AG, ball nut assemblies (BNA) undergo a vibroacoustical quality test and are binary classified based on their order spectra. In this work we formulate a multiple change point problem and derive optimal quality intervals and thresholds for the order spectra that minimize the number of incorrectly classified BNA. We pursue a multiobjective goal: the first objective function maximizes the Cohen Kappa metric, while the second objective function reduces the number of employed order intervals. The proposed approach is based on a genetic algorithm and incorporates prior information on the correlation structure of BNA and steering gear vibroacoustics, gained via canonical correlation analysis. The computational experiments show a reduction of both the number of employed order intervals and the costs arising from falsely classified BNA parts with respect to the current production setting, ensuring thus a high practical relevance of our suggested approach.

1 Introduction

The propagation of quality specifications across the supply chain is challenging in domains where quality can be subjective, such as in automotive acoustics. In the current work, we direct our attention to a situation encountered in the daily operations of one of the world's leading steering system suppliers, ThyssenKrupp Presta AG, where requirements imposed on the vibroacoustic quality of the steering gear need to be passed down to its subcomponents. Furthermore, only one subcomponent, the ball nut assembly (BNA), is subject to an own vibroacoustical quality test equivalent to the one of the steering gear.

*Corresponding author: pabucur@edu.aau.at

In the current production setting, the vibrational signals of the BNA are transformed to the frequency domain and analyzed as order spectra. For each order spectrum curve, acoustic domain experts determine a set of order intervals and corresponding quality thresholds. Between the orders corresponding to the left and right boundaries of each interval, the maximum value of the order spectrum curve must lie below the threshold, otherwise the BNA is marked as being qualitatively not ok and thus destined to be scrapped or reworked (see Figure 2 for an intuitive visualization).

This approach operates under the assumption that single orders in the steering gear order spectra are influenced by single orders in the BNA order spectra. If any steering gear order violates its threshold in its own test, then, in the case of a BNA fault, the corresponding BNA order must be itself subject to a stricter threshold. In the present work, we propose a more flexible methodology, which uses as input only the BNA order spectra and the information whether the steering gear failed its test due to the BNA, classifying thus the BNA in *OK* and *NOK*. An advantage of this relaxation is that the boundaries and their respective thresholds can be optimized and learned from the data, considering also possible non-linear effects in the vibroacoustical transfer path between the BNA and steering gear. Furthermore, the restriction on a limited number of spectral intervals, each equipped with a constant intensity threshold, can be considered as a regularization strategy against overfitting. As such, deriving the classifier amounts to optimizing an appropriate accuracy measure over the set of piecewise constant functions, with the design variables of the optimization problems being the interval boundaries, the individual threshold values and the number of intervals.

Similar optimization problems have been extensively considered in the context of piecewise constant regression, where a given time series is approximated by piecewise constant functions. This problem is often referred to as *multiple change point problem* [1, 2, 3]. An essential shortcoming of these approaches for our application is the assumption that the cost function $E(f)$ can be expressed as the sum of interval costs, i.e. $E(f) = \sum_{j=0}^{k} E\left(f_{|I_j}\right)$, where f is a piecewise constant function with k change points and intervals I_j. For this situation, efficient optimization algorithms such as dynamic programming [4] or binary segmentation [5] have been proposed. In our case, these approaches are not applicable, due to the fact that if the spectral threshold of any interval is violated, the cost is already maximal.

Our main contribution is two-fold: first, we solve a real-world multiple change point problem by developing an algorithm which hybridizes a custom genetic algorithm and the Nelder-Mead downhill simplex method. This results in a better performance w.r.t. the Cohen Kappa metric than the currently employed method and in the reduction of the solution complexity, due to the usage of less change points. Second, we speed up computations by guiding the genetic mutations using a weight function gained via deep canonical correlation analysis between the ball nut assemblies and the corresponding steering gear.

The remainder of the paper is organized as follows: in Section 2 we describe the novel problem formulation, followed by the description of the hybrid algorithm and its use of the canonical correlation weights in Section 3. In Section 4 we discuss the computational experiments and their results, concluding the work in Section 5 by mentioning possible directions for further research on the topic.

2 Problem Formulation

We shall henceforth assume N BNA order spectra as input, sampled at K common, equidistant order indices. The spectral intensities are thus given as a matrix $X \in \mathbb{R}^{N \times K}$. Each row of the matrix corresponds to an order spectrum and bears an additional label, *OK* or *NOK*, indicating whether the assembled steering gear fails the vibroacoustic test due to the BNA or not.

Mathematically speaking, the current quality examination method described in Section 1 amounts to binary classification of spectral signals. A signal is classified *OK* (positive class) if it is majorized by a predefined piecewise constant function and *NOK* else. Let $s_1, \dots, s_K$ be the grid of spectral orders at which the signals are sampled. We call f a *piecewise constant function with k change-points* if there exist indices $1 = i_0 < i_2 < \cdots < i_{k+1} = s_K$ and values $f_0, \dots, f_k$ such that $f(s) = f_j$ if $s \in (s_{i_j}, s_{i_{j+1}}]$. We further denote by $\mathcal{P}_k$ the set of all piecewise constant functions with k change-points.

In this work we aim to find an optimal function $f \in \mathcal{P}_k$, where optimality is expressed by means of two fitness measures. Due to our highly imbalanced dataset, in which the *OK* class has a prevalence of $p = 0.965$, we choose *Cohen's Kappa* [6] as an accuracy metric and thus as the first fitness function, accounting for the solution quality. In practice, the quantity $\kappa(f)$ is estimated from the training data. We refer to this estimated value as $\kappa_{\text{train}}(f)$. We further take an interest in the business aspect of the problem, since a faulty BNA assembled into a steering gear is ten times as expensive as a BNA that is mistakenly destroyed. To this effect, we introduce $E(f) = 10\ \text{FPR}(f) + \text{FNR}(f)$, where $\text{FPR}(f)$ and $\text{FNR}(f)$ are the false positive and false negative rate of the classification method defined by f; $E(f)$ is only used to evaluate the business impact, but not for optimization. The second objective function aims to reduce the total number of employed intervals.

With these definitions at hand, we formulate the multiple change point optimization problem:

$$\kappa^*(k) = \kappa_{\text{train}}(f^*(k)) = \max_{f \in \mathcal{P}_k} \kappa_{\text{train}}(f), \quad k \in \mathbb{N}, \tag{1}$$

where the optimal number of change points is denoted by k^* and the corresponding optimal piecewise constant function by $f^* \in \mathcal{P}_{k^*}$.

3 Proposed Algorithm

In order to solve problem (1) we maximize κ_{train} for a fixed k w.r.t. the threshold values and the change-point positions. We pursue an alternate direction approach [7], iteratively maximizing w.r.t. the thresholds while keeping the change-points fixed and vice versa. The maximization w.r.t. to the change-points is carried out by a genetic algorithm and the maximization w.r.t. to the threshold values by means of the Nelder-Mead downhill simplex method, with the complete algorithm sketched in Algorithm 1.

```
input  : N order spectrum curves containing N_OK curves with label OK
         number of change points k
         maximal number of generations B
output: Optimal piecewise constant threshold function f* ∈ P_k

Initialize population π[0] of 100 randomly generated functions in P_k;
for i from 1 to B do
    Update current population
    for f in π do
        Update threshold values for f such that κ_train(f) is maximized
        while keeping the change points constant.
    end
    Breed next generation
    π[i] = evolve(π[i − 1])
end
```

Algorithm 1: Genetic algorithm for finding an optimal thresholding function.

This procedure is repeated for $k = 1, 2, \ldots$, with the optimal k^* value manually chosen so as to optimize the trade-off between solution quality (*Cohen's Kappa*) and solution complexity (number of change points); the corresponding $f^* \in \mathcal{P}_{k^*}$ is the corresponding proposed solution. In the following we describe the individual substeps in Algorithm 1 in more detail.

In the **update** step a given function $f \in \mathcal{P}_k$ is modified by adapting its threshold values $(f_0, \ldots, f_k)$ while keeping the change point positions constant, such that $\kappa_{\text{train}}(f)$ is maximized. Since we deal with a data-driven objective function, the mapping $f \mapsto \kappa_{\text{train}}(f)$ is not differentiable w.r.t to the threshold values. We thus employ the Nelder-Mead downhill simplex method, a gradient-free maximization technique. We stop the iteration when the increments in κ_{train} and f_i both are less than 10^{-4}.

The **evolution** step follows classical approaches in genetic optimization [8]. In each iteration we retain those 20% of the population individuals with the highest fitness value for reproduction and add 5% randomly generated elements in $\mathcal{P}_k$. Pairs of step functions are chosen from these individuals at random with replacement and crossed-over to generate offspring, which is appended to the population until the original population size is reached. A proportion of 15% of the offspring is mutated after generation.

The **crossover** method is rather straightforward: The k change points of the two parent functions are merged and an agglomerative cluster algorithm is run on the $2k$ positions. The resulting dendrogram is then cut such that k clusters result, with the centroids of these clusters chosen as the offspring change points. The values between the new change points are obtained by averaging the function values of the parent functions over the new intervals.

In the **mutation** step the change point locations are randomly shifted with a probability of 0.15, under the constraint of preservation of order. The maximal allowed deviation is additionally constrained depending on the order. The results of previous work by means of deep canonical correlation analysis [9] showed that some parts of the BNA order spectrum are more influential on the acoustic property of the stearing gear than others. We use the correlation coefficients as weight functions in the mutation routine, introducing larger change-point variation in relevant parts of the spectrum.

4 Computational experiments

The dataset is composed of a total of 13257 (9279 training and 3978 testing) BNA, from which 454 (305/149) are *NOK*. The population size for the genetic algorithm is fixed at 100, with the maximal number of generations B fixed at 500. In Figure 1 the κ-values κ_{train} and κ_{test} are plotted against the number of change-points. Together with the company and considering the quality-complexity tradeoff on the κ_{train} values, we identified 9 as an optimal number of change points instead of the 14 currently employed, with the optimal piecewise constant function f^* depicted in Figure 2. The usage of f^* clearly increases the classification accuracy w.r.t. the Cohen Kappa metric, from the $\kappa_{\text{test}}(f^\dagger) = 0$ value resulting from the usage of the currently employed piecewise constant function $f^\dagger$ to a value of $\kappa_{\text{test}}(f^*) = 0.283$. While $\kappa = 0$ indicates only coincidental agreement, values in the range $(0.2, 0.4]$ represent already fair agreement. At the same time, the business costs on the test data are reduced by 24.16% from an initial value of $E(f^\dagger) = 1490$ to $E(f^*) = 1130$.

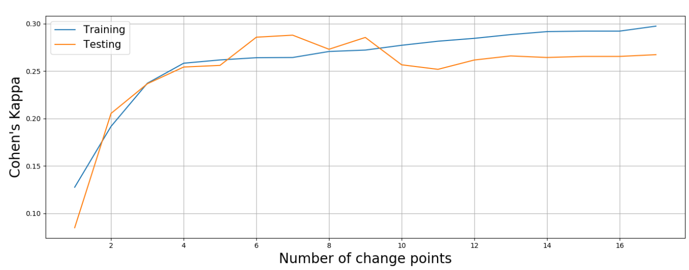

Figure 1: Optimal Cohen's Kappa in dependence of the number of change points k. The blue line corresponds to the optimal value of $\kappa_{\text{train}}(f^*(k))$ and the orange line to $\kappa_{\text{test}}(f^*(k))$.

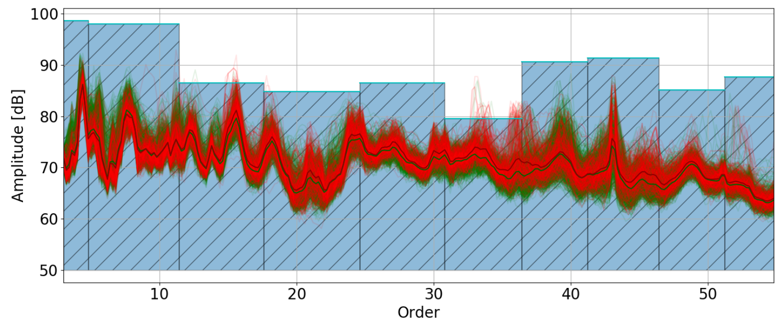

Figure 2: Best piecewise constant function $f^* \in \mathcal{P}_9$ together with *OK* (green) and *NOK* (red) BNA order spectra.

5 Conclusion

The cooperation with ThyssenKrupp Presta AG showed the potential of mathematical optimization techniques for real-world industrial problems. The combination of genetic algorithms and canonical correlation analysis allowed the optimization of the currently employed BNA quality examination method, increasing substantially its performance w.r.t. the Cohen Kappa metric and reducing at the same time its number of used change points and thus its ultimate complexity. The developed method is currently being evaluated for introduction into worldwide production systems, with the plausibility of the canonical weights already confirmed by acoustic domain experts. In an extended version of this work we aim to present more complex methods which are potentially able to further improve the results.

References

[1] Niu, Y.S., Hao, N., Zhang, H.: Multiple Change-Point Detection: A Selective Overview. In: Statistical Science, vol. 31, nr. 4, pp. 611-623, The Institute of Mathematical Statistics (2016).

[2] Frick, K., Munk, A., Sieling, H.: Multiscale change point inference. In: Journal of the Royal Statistical Society: Series B (Statistical Methodology), vol. 76, nr. 3, pp. 495-580 (2014).

[3] Truoung, C., Oudre, L., Vayatis, N.: A review of change point detection methods. In: CoRR (arXiv), vol. abs/1801.00718, (2018).

[4] Friedrich, F., Kempe, A., Liebscher, V., Winkler, G.: Complexity Penalized M-Estimation. In: Journal of Computational and Graphical Statistics, vol. 17, nr. 1, pp. 201-224, Taylor&Francis (2008).

[5] Fryzlewicz, P.: Wild binary segmentation for multiple change-point detection. In: The Annals of Statistics, vol. 42, nr. 6, pp. 2243-2281, The Institute of Mathematical Statistics (2014).

[6] Ben-David, A.: Comparison of classification accuracy using Cohen's Weighted Kappa. In: Expert system with Applications, vol. 34, nr. 2, pp. 825-832 (2008).

[7] Niesenm U., Shah, D., Wornell, G.: Adaptive Alternating Minimization Algorithms. In: 2007 IEEE International Symposium on Information Theory, pp. 1641-1645 (2007).

[8] Mitchell, M.: An Introduction to Genetic Algorithms. MIT Press, Cambridge MA USA, vol. 1 (1998).

[9] Bucur, P.A., Frick, K., Hungerländer, P.: Correlation analysis between the vibroacoustic behavior of steering gear and ball nut assemblies in the automotive industry, (2018). Preprint available at https://tinyurl.com/bucurp1.