Correlation analysis between the vibroacoustic behavior of steering gear and ball nut assemblies in the automotive industry

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

The increase in quality standards in the automotive industry requires specifications to be propagated across the supply chain, a challenge exacerbated in domains where the quality is subjective. In the daily operations of *ThyssenKrupp Presta AG*, requirements imposed on the vibroacoustic quality of steering gear need to be passed down to their subcomponents. We quantify the influence of ball nut assemblies on the steering gear by finding optimal encodings of their respective vibroacoustic signals, iteratively maximizing the correlation of their order spectra under orthogonality constraints. We compare the performance of linear and non-linear variants of this approach known as Canonical Correlation Analysis and establish the superiority of the neural network based variant in terms of attainable correlation. The practical relevance of our findings is guaranteed, since the visualization of the weights enables the identification of core influence areas.

1 Introduction

The development and incorporation of novel technologies in the automotive industry such as autonomous driving or steer-by-wire systems have accelerated technological change, increasing the system complexity significantly and posing new engineering challenges. The simultaneous increase in quality standards requires strict specifications to be propagated across the supply chain, from the carmaker companies down to the automotive component suppliers. Due to the complex interactions, it is often unclear how benchmarking thresholds for

*Corresponding author: pabucur@edu.aau.at

assembled components can be translated to quality requirements for their subcomponents. This challenge is exacerbated in domains where the quality of a product can be subjective, such as in automotive acoustics. With continuous new developments, the cumulative noise level inside a vehicle is being steadily reduced, rendering the contributions of auxiliary systems more obtrusive. In particular, the design and production of steering systems thus ceases to be a matter of simple functionality, with considerable efforts dedicated to the understanding of the vibroacoustic interaction of steering components. In the current work, we direct our attention to such a situation encountered in the daily operations of one of the world's leading steering system suppliers, *ThyssenKrupp Presta AG*. Strict quality requirements imposed on the vibroacoustic quality of the steering gear need to be passed down to its subcomponents, but the sheer number of vibroacoustic contributors and the inherent complexity make the building of a physical model impossible. 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. Known for its precision in the transformation of rotational to linear motion, the BNA features a helical raceway for the balls around the shaft axis, acting thus as a feed screw. The rotational movement of the balls is intrinsically paired with noise, rendering the BNA one of the major sources of undesired acoustic effects in the steering gear. The final result of each vibroacoustic test is the order spectrum of the recorded vibroacoustic signal. In order to minimize costs along the production chain, it is imperative to ensure that for each steering gear which fails its vibroacoustic test due to the built-in BNA, it would have been possible to recognize the cause in the BNA's own vibroacoustic test, avoiding the situation altogether in the future by setting quality thresholds on the BNA order spectra. Tracing back steering gear root faults to the BNA is not an easy task, since the identified patterns at steering gear level may manifest themselves differently and at different orders at BNA level. Furthermore, each BNA and steering gear order spectrum still contains 7510 values after the data preprocessing steps, which include heavy data aggregation. The optimization problem to be solved is the identification of an optimal data mapping from the high-dimensional feature space of the order spectra to a lower-dimensional set of principal variables, while maximizing a certain function under special contraints. While techniques such as the Principal Component Analysis aim to maximize the data variance in the lower-dimensional representation, we take an interest in encodings for the BNA and steering gear order spectra which are simultaneously the most predictive of and by each other. We achieve this kind of encoding by iteratively maximizing the correlation between the projections of BNA and steering gear order spectra under orthogonality constraints. The underlying technique is known as Canonical Correlation Analysis (CCA) [1], which is a subset of multi-view problems [2] and also closely related to mutual information problems [5]. The CCA estimate on the attainable correlation can be updated using non-linear methods, such as Kernel CCA (KCCA), which uses the kernel trick to solve the non-linear KCCA optimization problem [4]. Due to their high suitability as non-linear function approximators, deep neural networks can also be used to find new, maximally

correlated representations which, unlike the KCCA encodings, are not restricted to reproducing kernel Hilbert spaces with the corresponding kernels [3]. To the best of our knowledge, the canonical correlation analysis setting has not yet been employed for the identification of core influence areas in the vibroacoustic interactions between the BNA and steering gear in the automotive domain, an extension which the current work provides.

In our work we tackle different engineering challenges: first, we extend the two-view representation learning setting to a situation where the second view does not originate from the same product, but from it assembled together with hundreds of other components into a larger product. By analyzing the weights of the new encodings, we additionally show that the influence of the BNA on the assembled product, the steering gear, can be clearly reduced to core physical influence areas in their respective order spectra, which has also been validated by current domain expert knowledge.

The remainder of this contribution is organized as follows: in Section 2 we describe the novel problem formulation, followed by the description of the methodology in Section 3. In Section 4 we discuss the maximal correlation achieved by the different approaches, the optimization algorithm used for the tuning of the hyperparameter and the engineering interpretation of the weight curves generated by the optimization methods. Section 5 concludes the work offering suggestions for future research questions.

2 Problem formulation

The vibroacoustic test, both at BNA and at steering gear level, consists of steering the product left, then right, at two different mean rotational speeds: $300°/s$ and $500°/s$. The test is initiated with a speed of 97.5% of the mean rotational speed, increasing it uniformly until reaching $102,5\%$. The originating vibrations and their magnitude are recorded by two different accelerometers, positioned perpendicularly with respect to each other. The first accelerometer is positioned in the direction of the BNA axis, while the second records the lateral vibrations. For each tested product, the sensors are attached to it by industrial robots and measure the vibrations with a sampling rate of 25.6 kHz. The signals are then transformed to the frequency domain via the Fast Fourier Transform with a FFT length of 4096, Hann FFT-window and 75% overlap, where they are normed and rescaled to obtain an order spectrum. The resolution of the order spectrum is relatively high for an industrial context, with a stepsize of 0.2 and thus 5 samples per order. Finally, each of the 8 spectra resulting from the different combinations of accelerometers, steering direction and mean rotational speed was further standardized across the dataset to have zero mean and unit variance.

In the current production setting, acoustic domain experts average the order spectra for each single mean rotational speed over both accelerometers and steering directions, then define a set of boundary windows for each of the two averaged spectra. Between the left and right boundaries of each window, the

maximum value of the respective order spectrum must lie between the upper and lower bounds of the window. The averaged order spectra of each produced steering gear are thus checked to not violate the boundaries of any window, otherwise the product is scrapped or mechanically reworked and retested. The boundary windows of the steering gear may not directly apply to the BNA due to different reasons. First, it is important to notice that in the steering gear spectra we observe the interaction of the BNA with multiple other components. Since we lack data for the other components, it is not possible to determine how much each component influences a specific steering gear order. Furthermore, due to the presence of damping mechanisms in the steering gear, some order intervals of the BNA do not have a significant influence on the final steering gear acoustic behavior. Blindly setting the same upper and lower boundaries for the BNA as for the steering gear may thus cause unnecessary scrap on the BNA side, without ameliorating the problem. Finally, due to the complex interactions in the assembled steering gear, it it not guaranteed that a negatively prominent BNA order will manifest itself in negative patterns at the same order in the steering gear spectrum. An overview of the order spectra of both BNA and steering gear is depicted in Figure 1 for both qualitatively 'OK' and 'NOK' products. Due to the aforementioned reasons, it is imperative to identify the areas of the BNA and steering gear order spectra which are correlated. It is only for these areas of the BNA order spectrum that quality boundaries must be set. Furthermore, due to the high dimensionality of the feature space represented by the order spectra, a new lower dimensional BNA encoding must be found, which at the same time is most predictive of the BNA behavior in the assembled steering gear. Such encodings can be achieved by canonical correlation analysis methods, which iteratively maximize the correlation between the projections of BNA and steering gear order spectra under orthogonality constraints. This can be understood as extracting the representation of the underlying semantics present in two views of the same entity [4]. In the current setting, we do not posses two views of the same object. We propose however to consider the steering gear order spectra as another view of the BNA order spectra, since the BNA is one of the main vibroacoustic contributors to the steering gear. The mutual information in both spectra shall be the amplitudes of the BNA orders which also influence the steering gear vibroacoustic behavior. In this proposition, any contribution to steering gear vibroacoustics which can not be in average explained by the BNA is considered as noise, rendering the complete steering gear spectrum as a BNA spectrum under heavy noise. A first advantage of our proposition is that if any interaction with other components can in average be explained by the BNA, it will be present in the new encodings. The remaining noise is assumed to be uncorrelated between both production sites and thus filtered out in the new maximally correlated representations.

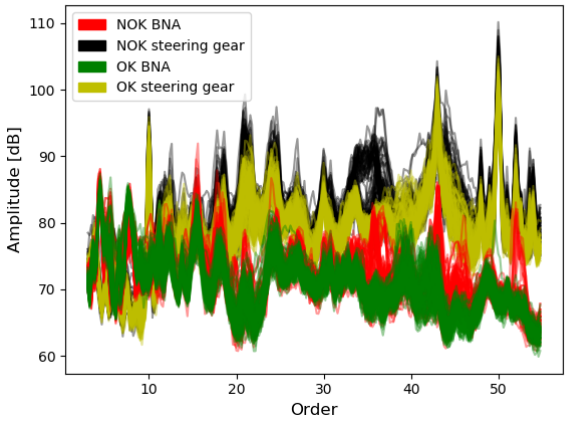

Figure 1: Overview of the connection between the BNA and the steering gear by means of the 300°/s averaged order spectrum plot for 70 OK and 70 NOK pairs. The 'NOK' BNA label refers to the case when the corresponding steering gear was faulty due to the BNA, even though the BNA had passed its own vibrational test.

3 Methodology

We assume the existence of two data matrices, $X = (\vec{x}_1, \ldots, \vec{x}_n)$ and $Y = (\vec{y}_1, \ldots, \vec{y}_n)$, whose entries represent paired observations $(\vec{x}_1, \vec{y}_1), \ldots, (\vec{x}_n, \vec{y}_n)$, where $\vec{x}_i$ is the i-th BNA order spectrum and $\vec{y}_i$ the corresponding spectrum of the steering gear. Each BNA spectrum $\vec{x}_i$ and steering gear spectrum $\vec{y}_i$ can be considered to be outcomes of two random vectors, $\hat{X}$ and $\hat{Y}$, respectively. To be more precise, the i-th observation consists of discrete order spectrum values $\vec{x}_i = (x_{i1}, \ldots, x_{im})^{'}$, where we assume that the discrete measurements $x_{1j}, \ldots, x_{nj}$ are sampled at the same spectral order s_j for all observations (no alignment or interpolation of the values is needed). Formally and following [3], given the two random vectors $(\hat{X}, \hat{Y}) \in \mathbb{R}^m \times \mathbb{R}^m$ with covariances $(\Sigma_{11}, \Sigma_{22})$ and cross-covariance Σ_{12}, the CCA algorithm [1] allows the computation of orthogonal sets of latent scores $(w_1^{'}\hat{X}, w_2^{'}\hat{Y})$. These represent maximally correlated linear combinations of the $\hat{X}$ and $\hat{Y}$ vectors, obtained using:

$$(w_1^*, w_2^*) = \arg\max_{w_1, w_2} \frac{w_1^{'} \Sigma_{12} w_2}{\sqrt{w_1^{'} \Sigma_{11} w_1 w_2^{'} \Sigma_{22} w_2}}.$$

Additional constraints require the projections to have unit variance and subsequent projections to be uncorrelated with the previous ones. The resulting optimization problem for the computation of the top $k \leq m$ projection pairs can then be formalized in the following way:

$$\begin{aligned} &\text{maximize} && \operatorname{Tr}(A_1^{'}\Sigma_{12}A_2) \\ &\text{subject to} && A_1^{'}\Sigma_{11}A_1 = A_2^{'}\Sigma_{22}A_2 = I, \end{aligned}$$

where the columns of the matrix $A_1 \in \mathbb{R}^{m\times k}$ are the top k projection weight vectors w_1^i, with A_2 constructed analogously. One solution to this optimization problem, assuming the non-singularity of the covariance matrices, is introduced in [6]. Defining

$$T \coloneqq \Sigma_{11}^{-\frac{1}{2}}\Sigma_{12}\Sigma_{22}^{-\frac{1}{2}}, \tag{1}$$

the optimal objective value is obtained by summing the top k singular values of T. This optimal value is encountered at $(A_1^*, A_2^*) = (\Sigma_{11}^{-\frac{1}{2}}U_k, \Sigma_{22}^{-\frac{1}{2}}V_k)$, where U_k, V_k are the matrices of the first k left-, respectively right- singular vectors of T. For practical purposes, we use the sample data matrices X and Y for the estimation of the covariances and the cross-covariance, which also has the benefit of allowing regularization.

The computation of the maximally correlated views has been extended from the linear setting to a kernel variant, the kernel CCA [4], and to a variant which uses neural networks to find the optimal representations, the so-called Deep CCA [3]. In the latter, the goal is to train two neural networks which learn the maximally correlated representations of $\hat{X}$ and $\hat{Y}$, respectively. Assuming a training size of n_t and a final encoding length of o in the final layer of each neural network, we denote by $H_1 \in \mathbb{R}^{o\times n_t}, H_2 \in \mathbb{R}^{o\times n_t}$ the matrices resulting from the final representation of the neural networks for the training data. The respective centered data matrices $\bar{H}_1 = H_1 - \frac{1}{n_t}H_1\mathbf{1}$ and $\bar{H}_2$ (calculated analogously), can then be used to estimate the covariance matrices $\hat{\Sigma}_{11} = \frac{1}{n_t-1}\bar{H}_1\bar{H}_1^{'} + r_1 I$, $\hat{\Sigma}_{22}$ and cross-covariance $\hat{\Sigma}_{12} = \frac{1}{n_t-1}\bar{H}_1\bar{H}_2^{'}$, where $r_1 > 0, r_2 > 0$ denote regularization constants. The estimated covariance and cross-covariance matrices are then employed to compute the correlation by summing the top k singular values of T in Equation 1, which in the special case when the respective encoding length o of each neural network equals k reduces to $\operatorname{corr}(H_1, H_2) = \operatorname{Tr}(T^{'}T)^{\frac{1}{2}}$.

4 Computational experiments

The dataset corresponding to the experimental framework described in Section 2 is composed of a total of 13257 BNA-steering gear pairs, which were split to include 70% of the pairs in the training set. A percentage of 20% of randomly chosen samples from the training set was further used to build the validation set, which remained identical across all experiments. A single data sample consists of the order spectra of the corresponding BNA-steering gear pair, concatenated for each of the two products to form an one-dimensional vector. We disregard the performance of the different algorithms on the training dataset and use the achieved correlation on the validation and test datasets as a performance metric,

with the gap between the validation and test correlation used as a measure for the inability of the model to generalize to unknown data. First experiments with the linear CCA algorithm showed that that upon training on the training data, the explanatory power of ≈ 50 CCA components was sufficient in order to explain almost all patterns leading to high correlation.

Research contributions such as [7] have repeatedly shown that the optimization of the hyperparameter in machine learning methods is crucial, with fine-tuned, comparatively simple methods offering competitive performance in image classification challenges. The tuning can occur manually or automatically; in the case of the latter, it is possible to perform grid- or random search over the hyperparameter space or to choose a model-based hyperparameter tuning strategy [8]. Such models describe the hyperparameter search as an optimization problem, where the cost consists of error made on the validation set by using a specific set of hyperparameter. The algorithm employed in the current work is the Tree-Structure Parzen Estimator (TPE), a Sequential Model-Based Global Optimization technique introduced in [9]. For each of the CCA algorithms introduced in Section 3, we set for the TPE algorithm a total of 200 iterations. During each iteration, we allowed the respective algorithm to train on the training set, then computed the achieved correlation on the validation set. After all iterations, the model with the hyperparameter set which featured the smallest validation set error was returned as the best candidate for each algorithm.

In the case of the linear CCA, the regularization parameter used for the estimation of the covariance matrix for each view were tuned in the range $[10^{-10}, 1000]$, with optimal values achieved at 4.99 for the BNA view and at 275.20 for the steering gear view. The KCCA algorithm follows the incremental SVD-based approach proposed in [10]. The kernel chosen for the KCCA implementation was the radial basis function kernel $k(x_i, x_j) = e^{-\frac{\|x_i - x_j\|^2}{2\sigma^2}}$, where σ denotes the bandwith parameter. For each view, the respective bandwidth was tuned in the range $[0.01, 10000]$, with optimal values found at 5659 for the BNA- and at 7085 for the steering gear kernel, respectively. The number of random samples for each view was tuned in the range $[10, 1000]$, with the regularization parameter again chosen between $[10^{-10}, 1000]$. The BNA kernel reached optimal performance at a value of 89 random samples and a regularization parameter of 346.79, compared to the values of 950 and 374.96 in the case of the second kernel. For both linear and kernel CCA, upon identification of the optimal hyperparameter, the respective algorithm was rerun on the complete training data, including validation data, with the best hyperparameter configuration.

According to the 'Practical Methodology' Chapter in [8], in the case of supervised learning with fixed-size vectors acting as input, it is recommended to use feedforward neural networks with fully connected layers, also known as Multilayer Perceptrons (MLP). A characteristic of the MLP is the fully connectedness of the layers, which implies that each neuron in a specific layer is connected to all neurons in the previous and next layer. While this certainly increases the explanatory power of the network, it is also prone to overfitting. We ameliorate this problematic by making use of diverse regularization mechan-

isms, such as Early Stopping, Dropout and Batch Normalization. Specifically, our implementation of Early Stopping treats the number of training epochs as another tuneable hyperparameter and allows model training for a maximum of 200 epochs: if after a certain number of epochs the reached correlation on the validation set has not increased, the training is stopped in order to avoid overfitting. The number of epochs used for monitoring purposes is also called the patience parameter, which was set to 3 epochs in our case; the minimal increase in terms of correlation is set to 0.01 after the 3 epochs. In case the training needs to be stopped, the model parameters which were in use before the last 3 epochs are returned. Finally, the networks are further trained on the data from the validation dataset for that specific number of epochs. A last regularization mechanism employed provided synthetic data augmentation by means of adding zero-centered Gaussian Noise to the input spectra, then normalizing the input vectors in the corresponding training batch. The variance of the added noise was treated as a hyperparameter and mostly found to be optimal at approximately 1.19 for both networks.

In the case of the MLP implementation for each of the two subnetworks, the number of fully connected layers, as well as their number of neurons and activations, were treated as hyperparameter. The network operating on the BNA spectra featured 5 fully connected layers after training, while the second network operating on the steering gear spectra employed 7 from a total of 15 maximally allowed layers. The number of neurons per layer was tuned in the range $[5, 5000]$, with all optimal values found in the $[1500, 2000]$ range. The activation functions were chosen by the optimization algorithm in each layer between *ReLu, SELU, sigmoid* and tanh. As a regularization mechanism, each fully connected layer was followed by a Batch Normalization and a Dropout layer. The optimal batch size was found to be 1710, with an optimal learning rate value of $6.27 \cdot 10^{-5}$ for the Adam optimizer. The optimal number of training epochs was encountered at 6, while most Dropout layers set a fraction rate of approximately 0.26 of the input units to 0 for the training phase.

4.1 Results in terms of achievable correlation

The performance of the different CCA algorithms was measured on the validation and test datasets with respect to the achieved correlation between the concatenated order spectra of the BNA, on one side, and of the steering gear, on the other side. The resulting correlation values are presented in Table 1. A noticeable property is the absence of a gap between the correlation achieved on the validation and the test data. This is a strong indication that the employed regularization was adequate, which specially in the case of the neural networks often constitutes a challenge. A further interesting result is the performance of the kernel variant of the CCA, since it did not manage to match the correlation achieved by the linear CCA. Furthermore, the very large values found to be optimal for the kernel bandwidths also suggest that the KCCA is not suitable for our specific data setting. The gap in achieved correlation between both methods and the DCCA by means of the MLP architecture is striking, with

	LCCA	KCCA	MLP
Validation	16.68	16.06	37.49
Test	15.97	15.24	37.43

Table 1: The performance of the linear, kernel and deep CCA compared on the validation and test datasets with respect to the achieved correlation between the concatenated order spectra of the BNA and of the steering gear.

it clearly outperforming them despite the modest architecture and the strong regularization mechanisms.

4.2 Visualization of the weight functions

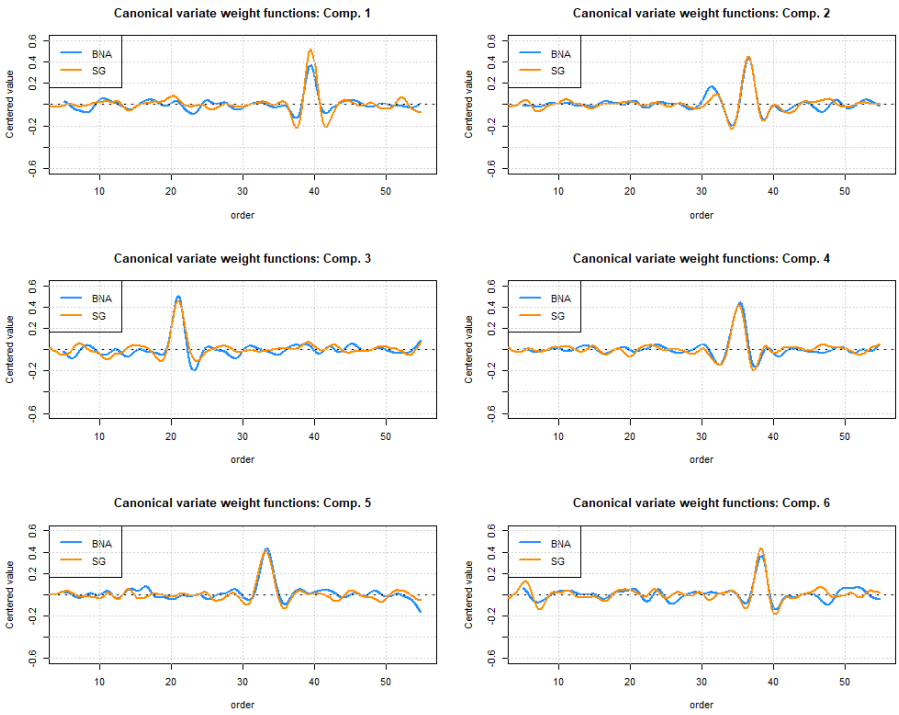

Figure 2: Centered values of the weights imposed by the linear CCA algorithm on the BNA and steering gear (SG) $300°/s$ order spectrum.

While the achievable correlation values are indicative of the validity of the assumptions and the new learned lower dimensional representations can be used for further quality improvement efforts, it is also highly interesting to visualize

the weights imposed on the original order spectra. In the case of the linear CCA, the weights for the $300°/s$ mean rotational speed can be visualized in Figure 2, with the $500°/s$ weights presenting a very similar picture. We encounter specific intervals which are heavily weighted, concentrated around order 20 on the one hand and between orders 30 and 43 on the other hand. Also remarkable are the almost identical amplitudes and form of the weight curves of both BNA and steering gear across the complete order spectra. All significant oscillations in the weight function could be traced back to root issues in the whirling or grinding manufacturing processes or to other manufacturing issues related to the ball nut.

5 Conclusion

In this paper we presented a solution to an engineering problem encountered in the supply chain of an automotive company, *ThyssenKrupp Presta AG*, whose specific problem setting is substantially different from the nearest research domain, the field of multi-view problems. An extended version of this publication will present a metaheuristic solution which makes use of the CCA weights in order to direct the optimization of the boundary windows towards more promising areas of the search space, thus speeding up computations.

References

[1] Hotelling, H.: Relations between two sets of variates. In: Biometrika, vol. 28, nr. 3/4, pp. 321-377, JSTOR (1936).

[2] Zhao, J., Xie, X., Xu, X., Sun, S.: Multi-view learning overview: Recent progress and new challenges. In: Information Fusion, vol 38, pp. 43-54 (2017).

[3] Galen, A., Arora, R., Bilmes, J., Livescu, K.: Deep canonical correlation analysis. In: International Conference on Machine Learning, pp. 1247-1255 (2013).

[4] Hardoon, D.R., Szedmak, S., Shawe-Taylor, J.: Canonical correlation analysis: An overview with application to learning methods. In: Neural computation, vol. 16, nr. 12, pp. 2639-2664, MIT Press (2004).

[5] Borga, M.: Learning multidimensional signal processing (PhD Thesis). Linköping University Electronic Press (1998).

[6] Bibby, J.M., Kent, J.T., Mardia, K.V.: Multivariate analysis. Academic Press, London (1979).

[7] Coates, A., Ng, A.Y.: The importance of encoding versus training with sparse coding and vector quantization. In: Proceedings of the 28th International Conference on Machine Learning (ICML-11), pp. 921-928 (2011).

[8] Goodfellow, I., Bengio, Y., Courville, A.: Deep Learning. MIT Press, Cambridge, vol. 1 (2016).

[9] Bergstra, J., Bardenet, R., Bengio, Y., Kégl, B.: Algorithms for Hyper-parameter Optimization. In: Proceedings of the 24th International Conference on Neural Information Processing Systems, pp. 2546-2554. Curran Associates Inc., USA (2011).

[10] Arora, R., Livescu, K.: Kernel CCA for multi-view learning of acoustic features using articulatory measurements In: Symposium on Machine Learning in Speech and Language Processing (2012).