Robust machine fault detection with Independent Component Analysis and Support Vector Data Description

Alexander Ypma, David M. J. Tax and Robert P.W. Duin

Pattern Recognition Group
Dept. of Applied Physics, Faculty of Applied Sciences, Delft University of Technology
Lorentzweg 1, 2628 CJ Delft, The Netherlands
Tel: +31-15-2781845. Fax: +31-15-2786740
{ypma,davidt}@ph.tn.tudelft.nl

Keywords: Independent Component Analysis, blind signal separation, outlier detection, Support Vector Machines, machine diagnostics, fault detection

Abstract

We propose a novel approach to fault detection in rotating mechanical machines: fusion of multichannel measurements of machine vibration using Independent Component Analysis, followed by a description of the admissible domain (part of the feature space indicative of normal machine operation) with a Support Vector Domain Description (SVDD) method. The SVDD-method enables the determination of an arbitrary shaped region that comprises a target class of a dataset. In this particular application, it provides a way to quantify the compactness of the admissible class in relation to data preprocessing. Application to monitoring of a submersible pump indicates that combination of measurement channels with ICA gives improved results in fault detection, without requiring detailed prior knowledge on origin and type of the failure.

1 Introduction

Monitoring of mechanical machinery involves the analysis of an ensemble of vibration measurements, yielding a combined view of the periodic force variations within the machine. Fault development or wear is usually accompanied by increased harmonics of the machine running frequency, increased (amount of) sidebands around structure-related fault frequencies and a shift in the amplitude distribution of the measurement signals [CT93]. Several authors proposed the use of higher-order statistics in this problem ([BPZ95, BKH93, MP97]) to exploit nongaussian and nonlinear behaviour in a machine (e.g. in the presence of bearing and gearbox failure [Oeh96] and cyclostationarity [MN98]).

In this paper we use Independent Component Analysis for the combination of several sensor measurements in the time domain (*sensor fusion*). The aim is to find the manifestation of a vibrating (fault) source by using the spatial redundancy in the multichannel measurement time series. This could provide a fault detection procedure that is less sensitive to careful sensor choice and specific details of the sensor mounting process.

In machine condition monitoring it is not possible to collect an exhaustive set of measurements from all possible failure scenarios. Moreover, since a machine may be used under very different operating modes (running speed, load, type of lubricants) and environmental conditions (indoor or outdoor, placed near interfering machinery), the admissible behaviour (usually called: the machine fingerprint) actually comprises a *set* of fingerprints. In order to perform reliable fault detection, significant deviations from this normal description should be detected. We propose a new method for description of a dataset (the target class), that enables to determine the rejection rate of the target class a priori. The method is based on the support vector method proposed by Vapnik [Vap95].

2 Sensor fusion with Independent Component Analysis

In Independent Component Analysis (ICA) one tries to retrieve an ensemble of independent source signals $\mathbf{s}$ out of an ensemble of linear mixtures $\mathbf{x}$ by determining the inverse of the mixing matrix A (up to permutation and scaling) [HO97]

$$\mathbf{y} = W\mathbf{x}, \quad W = \Lambda P A^{-1} \tag{1}$$

Several solutions to this *blind signal separation* problem have been developed, but all of them are based on the notion of minimization of the mutual information between the reconstructed components $\mathbf{y} = \hat{\mathbf{s}}$ (which is sufficient for obtaining independency) [LGBS98]. There are a few assumptions involved here: the model assumes linear instantaneous mixing (which is expected to be correct for the pump under investigation), whereas only one of the sources is allowed to be Gaussian and noise is not explicitly taken into account. Lately, attempts have been made to take temporal coherence into account ([BAMCM97, Hyv98, Ama99, YP99]), leading to algorithms in which second-order statistics suffice for separation (if sources have different spectral characteristics).

The linear approach of (1) often leads to algorithms that maximize nongaussianity of the recovered components, i.e. the distribution of the projection of the data onto a vector in the ICA-basis should be as far from Gaussian as possible. If it cannot be assumed that independent sources underly a dataset, ICA may therefore be regarded as a form of exploratory projection pursuit [Fri87]. In the application at hand (machine diagnostics) it is known that incipient faults often manifest themselves as impulses in the time signal, e.g. it has been reported that bearing failure and loose foundation give rise to nongaussianity (increased kurtosis) of vibration measurements.

We performed the ICA-based sensor fusion in the experiments of section 4 using the *fast-ica*-algorithm developed by Hyvärinen ([HO97]). This algorithm works with both sub- and supergaussian sources and shows cubic convergence. More specifically, we used the deflation approach with cubic nonlinearity. If a component couldn't be found because the algorithm didn't converge in 100 cycles, the multichannel time series at hand was disregarded in feature extraction (see section 4).

3 Data domain description with support vectors

For description of the domain of a dataset one can try to capture it with a sphere with minimum volume. Inspired by the Support Vector Method by Vapnik (see [Vap95], or for a more simple introduction [TdRD97]) one can extend this idea to determine an arbitrary shaped region in the original feature space. We will refer to the method as the Support Vector Data Description (SVDD) method [TD99].

Assume we a have data set containing N data objects, $\{x_i, i = 1, .., N\}$ and the sphere is described by center a and radius R. We now try to minimize an error function containing the volume of the sphere. The constraints that objects are within the sphere are imposed by applying Lagrange multipliers:

$$L(R, a, \alpha_i) = R^2 - \sum_i \alpha_i \{R^2 - (x_i^2 - 2ax_i + a^2)\} \tag{2}$$

with Lagrange multipliers $\alpha_i \geq 0$. This function has to be minimized with respect to R and a and maximized with respect to α_i.

Setting the partial derivatives of L to R and a to zero, gives:

$$\begin{aligned} \sum_i \alpha_i &= 1 \\ a &= \frac{\sum_i \alpha_i x_i}{\sum_i \alpha_i} = \sum_i \alpha_i x_i \end{aligned} \tag{3}$$

This shows that the center of the sphere a is a linear combination of the data objects x_i.

Resubstituting these values in the Lagrangian gives to maximize with respect to α_i:

$$L = \sum_i \alpha_i (x_i \cdot x_i) - \sum_{i,j} \alpha_i \alpha_j (x_i \cdot x_j) \tag{4}$$

with $\alpha_i \geq 0$, $\sum_i \alpha_i = 1$.

This function should be maximized with respect to α_i. In practice this means that a large fraction of the α_i become zero. For a small fraction $\alpha_i > 0$ and the corresponding objects are called Support Objects. We see that the center of the sphere depends just on the few support objects, objects with $\alpha_i = 0$ can be disregarded.

Object z is accepted when:

$$\begin{aligned} &(z-a)(z-a)^T = (z - \sum_i \alpha_i x_i)(z - \sum_i \alpha_i x_i) \\ &= (z \cdot z) - 2\sum_i \alpha_i (z \cdot x_i) + \sum_{i,j} \alpha_i \alpha_j (x_i \cdot x_j) \\ &\leq R^2 \end{aligned} \tag{5}$$

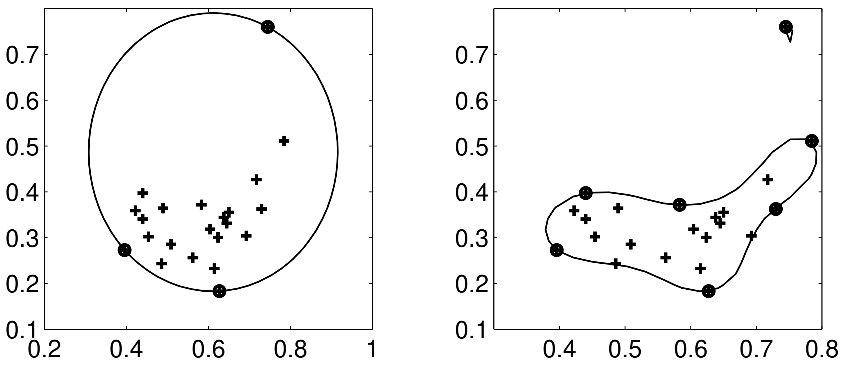

Figure 1: Data description of a small data set, (left) normal spherical description, (right) description using a Gaussian kernel.

In general this does not give a very tight description. Analogous to the method of Vapnik [Vap95], we can replace the inner products $(x \cdot y)$ in equations (4) and in (5) by kernel functions $K(x, y)$ which gives a much more flexible method. When we replace the inner products by Gaussian kernels for instance, we obtain:

$$(x \cdot y) \rightarrow K(x, y) = \exp(-(x-y)^2/s^2) \qquad (6)$$

Equation (4) now changes into:

$$L = 1 - \sum_i {\alpha_i}^2 - \sum_{i \neq j} \alpha_i \alpha_j K(x_i, x_j) \qquad (7)$$

and and the formula to check if a new object z is within the sphere (equation (5)) becomes:

$$1 - 2\sum_i \alpha_i K(z, x_i) + \sum_{i,j} \alpha_i \alpha_j K(x_i, x_j) \leq R^2 \quad (8)$$

We obtain a more flexible description than the rigid sphere description. In figure 1 both methods are shown applied on the same two dimensional data set. The sphere description on the left includes all objects, but is by no means very tight. It includes large areas of the feature space where no target patterns are present. In the right figure the data description using Gaussian kernels is shown, and it clearly gives a superior description. No empty areas are included, what minimizes the change of accepting outlier patterns.

This Gaussian kernel contains one extra free parameter, the width parameter s in the kernel (equation (6)). As shown in [TD99] this parameter can be set by setting a priori the maximal allowed rejection rate of the target set, i.e. the error on the target set. This error can be estimated by the number of support vectors:

$$E[P(\text{error})] = \frac{\#SV}{N} \qquad (9)$$

where $\#SV$ is the number of support vectors. We can regulate the number of support vectors by changing the width parameter s and therefore we can also set the error on the target set on the prespecified value.

Note we cannot set a priori restrictions on the error on the outlier class. In general we only have a good representation of the target class and the outlier class is per definition everything else.

4 Experiments

An in-house test rig (figure 2) consists of a small submersible pump in a water basin, which can be made to run at several speeds (from 46 to 54 Hz) and several loads (by closing a membrane controlling the water flow). A number of faults were induced to this pump (loose foundation, imbalance, failure in the outer race of the uppermost ball bearing), and vibration was measured with three accelerometers (yielding 5 measurement channels, one of the sensors is triaxial). Both normal and faulty behaviour was measured at several speeds and loads.

4.1 Feature extraction

We compared several methods for feature extraction from vibration data. It is well known that faults in rotating machines will be visible in the acceleration spectrum as increased harmonics of running speed or presence of sidebands around characteristic (structure-related) frequencies. Due to overlap in series of harmonic components and noise, high spectral resolution may be required for adequate fault identification. This may lead to difficulties because of the curse of dimensionality: one needs large sample sizes in high-dimensional spaces in order to avoid overfitting

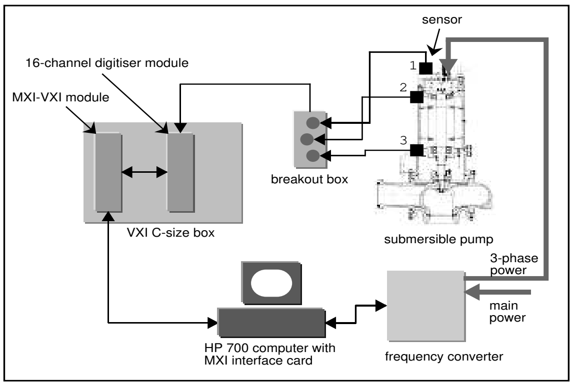

Figure 2: Experimental setup with submersible pump

of the train set. Hence we focused on relatively low feature dimensionality (64 spectral bins). A previous study [TYD] showed that features based on autoregressive modelling (order 64) and MUSIC spectrum estimation (64 spectral bins) performed very well in this application. We compared this to feature vectors obtained with a 64-bins envelope spectrum, i.e. the spectrum of an envelope detected (or *demodulated*) time signal. No bandpass-filtering was employed here. As a reference, a high-resolution (512 spectral bins) normalized power spectrum was obtained from the measurements. For further details on experimental setup and feature extraction for machine diagnostics, see [TYD, YP99].

We compared our ICA combination approach with taking a linear combination of multichannel measurements based on variance maximization (Principal Component Analysis). Alternatively, we used the SVDD with datasets that were generated with feature extraction from measurement signals without applying ICA or PCA. Here, features from all measurement channels were included into one dataset.

4.2 Detection Results

The SVDD is applied to several datasets, differing in sensor combination method and feature extraction method. Since test objects from the outlier class are available (i.e. the fault class comprising imbalance, loose foundation and bearing failure), the rejection performance on the outlier set can also be measured.

In all experiments we have used the SVDD with a Gaussian kernel. For each of the feature sets we have optimized the width parameter s in the SVDD such that $1\%, 5\%, 10\%, 25\%$ and 50% of the target objects will be rejected, so for each data set and each target error another width parameter s is obtained. For each feature set this gives a acceptance-rejection curve for the target and the outlier class.

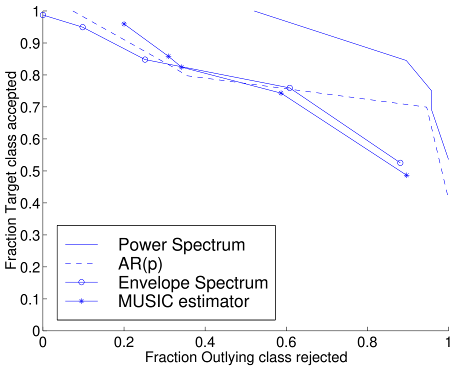

Figure 3: Acceptance/rejection performance of the SVDD on all types of features with all sensor measurements collected.

In figure 3 the performances of the different features are compared with the 512 bin power spectrum with no ICA applied. The high resolution power spectrum clearly outperforms all other methods. When all of the target class objects is accepted, over 50% of the outlier class is rejected. When about 10% of the target class is rejected, almost 80% of the outlier class can be rejected. This shows that a large fraction of the abnormal patterns can be distinguished from the normal class, but some overlap exist. Performance of the SVDD on the AR-model features, envelope spectrum features and the MUSIC estimator are clearly inferior. Especially when large target acceptance rates are required, large fractions of the outlier class are accepted. Only the AR-model features approximates the performance of the 512 bin power spectrum for smaller target acceptance rates.

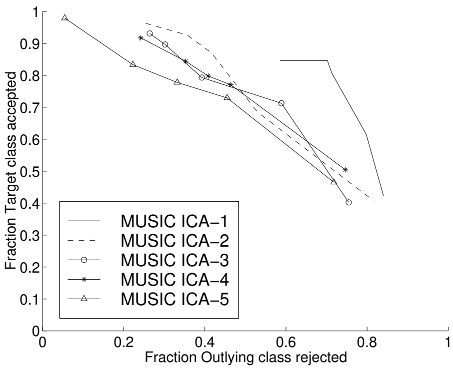

Figure 4: Acceptance/rejection performance of the SVDD on the MUSIC features with different number of components used.

In figure 4 the SVDD is applied to the MUSIC estimator data on which an ICA is performed. Results vary with the number of independent components that are included into the dataset. Performance is optimal when just one component is used, using more than one component results in (comparable) worsened performance for all situations. Inclusion of all 5 independent components gives the worst results. Clearly the first ICA-component contains almost all useful information to distinguish between normal and abnormal behaviour.

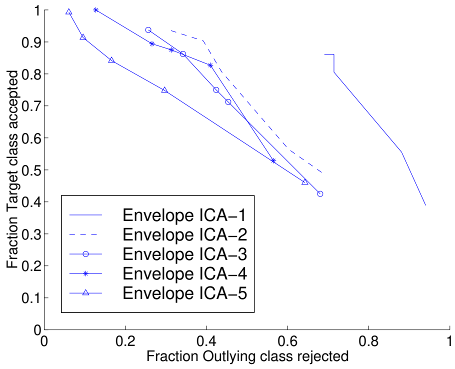

Figure 5: Acceptance/rejection performance of the SVDD on the envelope features with different number of components used.

In figure 5 can observe the same trend for the envelope features. Again, the first independent component yields most information for fault detection, including multiple components gives significantly worse results. The envelope features also perform slightly better than the MUSIC frequency features, but still do not reach the performance obtained with the power spectrum.

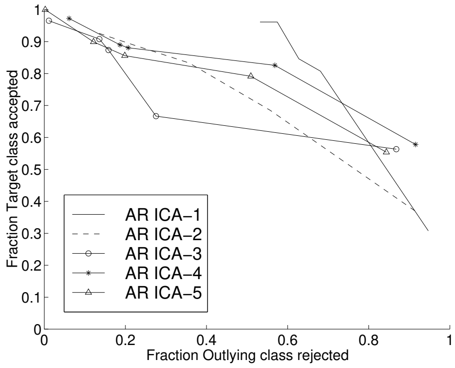

Figure 6: Acceptance/rejection performance of the SVDD on the envelope features with different number of components used.

In figure 6 the performance is shown for the AR-model features. Again the first component gives best performance, but for smaller target class acceptance rates, inclusion of multiple components is useful. In this case it is less clear how much useful information is retained when taking only the first component into account. Although in the non-ICA case the AR model features better distinguish between target and outlier class (see figure 3) than envelope spectrum features, in the ICA case the AR features perform worse than envelope features (compare figures 5 and 6).

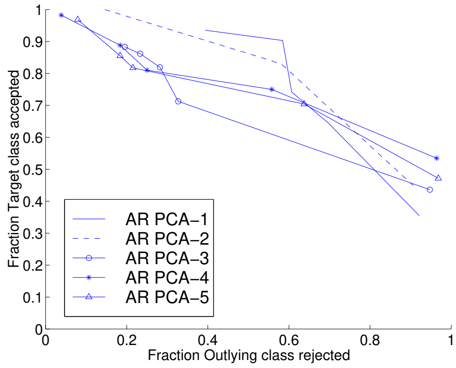

Figure 7: Acceptance/rejection performance of the SVDD on the envelope features with different number of components used.

Last, AR-model features were computed on data with PCA preprocessing (figure 7). Taking just the first principal component is not enough, only for very large target acceptance rates it outperforms the sets using more components. When larger fractions of the outlier class should be rejected, information from more components should be used.

The final observation is that performance is best using envelope spectrum features on ICA-combined measurement channels.

5 Conclusion

For robust fault detection in rotating mechanical machinery it may be useful to use spatial redundancy in measurements that are obtained simultaneously with multiple spatially distributed vibration sensors. In order to quantify this effect, we used a data domain description method to capture the normal behaviour of a machine and analyzed the trade-off between acceptance of data from normal behaviour and rejection of data from fault situations. For three typical features (AR-modelling, MUSIC spectrum estimation and envelope spectrum estimation) performance on ICA-combined measurement channels was usually better than linear combination using Principal Components Analysis. Moreover, if we fuse all channels into one combined channel with ICA, detection results are superior to the results obtained with datasets from all original (not fused) measurement channels. Previous results [TYD] showed that careful choice of sensor channel to include in a dataset may enhance performance, but this requires prior knowledge on the nature and spatial origin of a failure. Using ICA for sensor fusion may overcome this limitation, while retaining high detection accuracy.

6 Acknowledgments

This work was partly supported by the Foundation for Applied Sciences (STW) and the Dutch Organisation for Scientific Research (NWO). We would like to thank Ronald Ligteringen, SKF Condition Monitoring B. V. and Landustrie Sneek B. V. for cooperation in measurements and setup of the test rig.

References

[Ama99] S.-I. Amari. ICA of temporally correlated signals - learning algorithm. In *Proceedings of the First International Workshop on Independent Component Analysis and Signal Separation, ICA'99*, pages 13 – 18, 1999.

[BAMCM97] A. Belouchrani, K. Abed-Meraim, J.-F. Cardoso, and E. Moulines. A blind source separation technique using second-order statistics. *IEEE Transactions on signal processing*, 45(2):434 – 444, 1997.

[BKH93] R. W. Barker, G. Klutke, and M.J. Hinich. Monitoring rotating tool wear using higher-order spectral features. *Journal of engineering for industry*, 115:23 – 29, February 1993.

[BPZ95] B. Boashash, E. J. Powers, and A. M. Zoubir, editors. *Higher-Order Statistical Signal Processing*, Melbourne, 1995. LONGMAN (Wiley).

[CT93] C. M. Campbell and M. S. Tavakoli. Comparative signature analysis using needle bearings. *ASME Journal on Vibration, shock, damage and identification of mechanical systems*, 64:55 – 64, 1993.

[Fri87] J. H. Friedman. Exploratory projection pursuit. *Journal of the American Statistical Association*, 82(397):249 – 266, March 1987.

[HO97] A. Hyvärinen and E. Oja. A fast fixed-point algorithm for independent component analysis. *Neural computation*, 9(7):1483–1492, 1997.

[Hyv98] A. Hyvärinen. Independent component analysis of time-dependent stochastic processes. In *Proceedings of the 8th International Conference on Artificial Neural Networks (ICANN98)*, volume 1, pages 135–140. Springer-Verlag, 1998.

[LGBS98] T.-W. Lee, M. Girolami, A. J. Bell, and T. J. Sejnowski. A unifying information-theoretic framework for independent component analysis. *International journal on mathematical and computer modeling*, 1998.

[MN98] A. C. McCormick and A. K. Nandi. Cyclostationarity in rotating machine vibrations. *Mechanical Systems and Signal Processing*, pages 225–242, March 1998.

[MP97] A. Murray and J. Penman. Extracting useful higher order features for condition monitoring using artificial neural networks. *IEEE Transactions on signal processing*, 45(11):2821 – 2828, November 1997.

[Oeh96] H. Oehlmann. *Analyse temps-fréquence de signaux vibratoires de boîtes de vitesses*. PhD thesis, Université Henri Poincaré, Centre de Recherche en Automatique de Nancy, Nancy, (Fr), 1996.

[TD99] D.M.J. Tax and R.P.W Duin. Data domain description using support vectors. In Verleysen, M., editor, *Proceedings of the European Symposium on Artificial Neural Networks 1999*, pages 251 – 256, Brussels, April 1999. D-Facto.

[TdRD97] D.M.J. Tax, D. de Ridder, and R.P.W. Duin. Support vector classifiers: a first look. In *Proceedings of ASCI'97*. ASCI, 1997.

[TYD] D. M.J. Tax, A. Ypma, and R. P. W. Duin. Support vector data description applied to machine vibration analysis. To appear in the Proceedings of ASCI'99.

[Vap95] V. Vapnik. *The nature of statistical learning theory*. Springer-Verlag, New York, 1995.

[YP99] A. Ypma and P. Pajunen. Rotating machine vibration analysis with second-order independent component analysis. In *Proceedings of the First International Workshop on Independent Component Analysis and Signal Separation, ICA'99*, pages 37 – 42, January 1999.