On the Suitability of Different Features for Anomaly Detection in Wire Ropes

Esther-Sabrina Platzer[1], Herbert Süße[1], Josef Nägele[2], Karl-Heinz Wehking[2]
and Joachim Denzler[1]

[1] Chair for Computer Vision, Friedrich Schiller University of Jena
{Esther.Platzer,Joachim.Denzler,Herbert.Suesse}@uni-jena.de
http://www.inf-cv.uni-jena.de
[2] Institute of Mechanical Handling and Logistics, University Stuttgart
{Naegele,Karl-Heinz.Wehking}@ift.uni-stuttgart.de
http://www.uni-stuttgart.de/ift

Abstract. Automatic visual inspection of wire ropes is an important but challenging task, as anomalies in the rope are usually unobtrusive. Certainly, a reliable anomaly detection is essential to assure the safety of the ropes. A one-class classification approach for the automatic detection of anomalies in wire ropes is presented. Furthermore, the performance of different well-established features from the field of textural defect detection are compared with respect to this task. The faultless rope structure is thereby modeled by a Gaussian mixture model and outliers are regarded as anomaly. To prove the practical applicability, a careful evaluation of the presented approach is performed on real-life rope data. In doing so, a special interest was put on the robustness of the model with respect to unintentional outliers in the training and on its generalization ability given further data from an identically constructed rope. The results prove good recognition rates accompanied by a high generalization ability and robustness to outliers in the training set.

1 INTRODUCTION

Wire ropes are used in many fields of logistics. They are deployed as load cable for bridges, elevators and ropeways. This implies a high strain by external powers every day. Unfortunately, this can lead to structural anomalies or even defects in the rope formation. A defective rope bears a high risk for human life. This motivates the strict rules summarized in the European norm [1], which instruct a regular inspection of wire ropes.
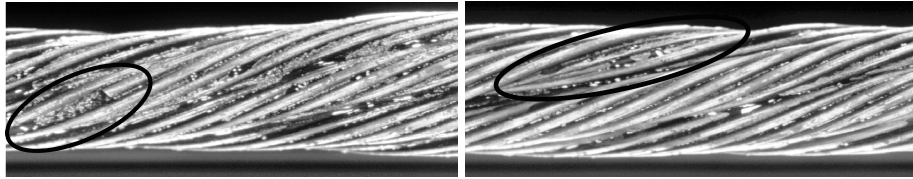
Fig. 1. Rope defects: in the left image you can see a wire fraction and in the right image a wire is missing.

Risky defects, prominent in wire ropes, are small wire fractions, missing wires, and damaged rope material due to lightening strokes. Furthermore, structural anomalies caused by interweavement of the rope ends or a reduced stress are also in the focus of interest. In Fig. 1, two exemplary defects are marked in the rope. Visual inspection of wire ropes is a difficult and dangerous task. Besides, the inspection speed is quite high (on average 0.5 meters/second) which makes it a hard effort, to concentrate on the passing rope without missing small defective rope regions.

A prototypic acquisition system was developed to overcome these limitations [2]. Four line cameras record the passing rope and yield four different views. By this, the rope can be inspected in the office without time pressure. The images in Fig. 1 were acquired with this system.

Defects and anomalies in wire ropes are unimposing and small. The image quality is deranged by mud, powder, grease or water and the lighting conditions change frequently. Therefore, the choice of features for the detection task is important. Recent approaches for defect or anomaly identification focus on fault detection in material-surfaces. In [3] we introduced a one-class classification approach for anomaly detection in wire ropes using linear prediction (LP) coefficients as features and a Gaussian mixture for model learning. This former work is extended now by two main aspects. First of all, we compare the performance of LP features to that of well-established features from the field of textural defect detection as well as to that of features based on histograms of oriented gradients (HOG). Secondly, the robustness to outliers in the training set as well as the generalization ability of the presented approach are carefully evaluated. The last point is of particular interest for the practical relevance of the method.

Features based on local binary patterns (LBP) were first introduced by Ojala [4] for texture classification. Recently, they were used for defect detection in fabrics [5] and for real-time surface inspection [6]. Textural features, extracted from co-occurrence matrices, were proposed by Harlick in the early 70's [7] and are frequently used for texture description [8]. Iivarinen [9] compares two histogram-based methods for surface defect detection using LBP and co-occurrence matrices. Rautkorpi et al. [10] used shaped-based co-occurrence matrices for the classification of metal surface defects. Vartiainen et al. [11] focus on the detection of irregularities in regular, periodic patterns. They separate the image data in a

regular and an irregular part. Based on the resulting irregularities, we compute local histograms, which serve as features. In addition to our previously published work [12] on this topic, also features based on histograms of oriented gradients (HOG) are considered. They were originally introduced for human detection by Dalal and Triggs [13] and have gained in importance in the last years. As it is obvious, the regular rope structure results in some eminent gradient orientations. For this reason, HOG features seem to be a suitable choice for the problem of rope anomaly detection.

Another important category of features for texture analysis and textural defect detection are wavelet-based features. Kumar and Pang [14] for example use Gabor features for the detection of defects in textured material. However, the computation of these features requires large filter banks and high computational costs. Due to the huge size of rope data sets (20-30 GB) the time-consuming computation of Gabor features seems to be not the best choice. In [15] the authors state, that similar results to that obtained by the usage of wavelet features can be resolved with help of joint neighborhood distributions and less computational effort.

The one-class classification strategy proposed in [3] was chosen due to a lack of defective training samples for a supervised classification. In contrast, it is no problem to design a huge sample set of faultless training samples. With this faultfree training set a model of the intact rope structure can be learned. In the detection step outliers with regard to this model are classified as defect. However, the only available ground truth information about this training data is the labeling of the human expert. In the following, there remains a small uncertainty of underdiagnosed defects in the training set. For this reason, the robustness of the proposed method to outliers in the training set is evaluated. Results obtained by learning from a faultless training set are compared to those, obtained by learning from a training set with intentionally added, faulty samples. The generalization ability of a learned model is a further important point, especially for the practical relevance of the presented method. There exist only a limited number of different construction types for wire ropes. The differences between them are mainly a different number of wires and strands, different thickness of single wires, the length of twist and the diameter. If just one model for every possible rope type would have to be learned in advance, this would save a lot of computational effort. However, the rope data from different ropes differs significantly due to the changing acquisition conditions and a different mounting of the ropes. Nevertheless, it is desirable to have just one model for every construction type and to overcome the challenges of a changing acquisition environment. Therefore, the generalization ability of the learned models is evaluated by learning and testing on different rope data from nearly identical constructed ropes.

The paper is structured as follows: in Sect. 2 we briefly review the feature extraction using linear prediction. A short description of the used textural and HOG features and their extraction is given. The one-class classification of wire rope data is shortly summarized in Sect. 3. Experiments, revealing the usability

and robustness of our approach, have been performed on real-life rope data and results are presented in Sect. 4. A conclusion and a discussion about future work is given in Sect. 5.

2   FEATURE EXTRACTION

In this section, the different features are briefly reviewed. Their extraction from the underlying rope data is described, as it differs for the LP features in contrast to the remaining ones. Furthermore, we will shortly motivate the choice of every feature with respect to the rope analysis task.

Local binary patterns (LBP) code the local graylevel-structure of a pixel neighborhood. Histograms based on the resulting codes lead to a local feature distribution. Since local binary patterns incorporate contextual information from a local neighborhood, a comparison of their performance with that of the LP features is of particular interest.

Harlick [7] introduced a set of 14 different textural features computed from co-occurrence matrices. They reveal the spatial distribution of gray-levels and though seem to be an interesting choice for structures with a certain regularity.

The detection of irregularities, proposed in [11] focuses on anomalies in regular, periodic patterns. Since the structure of wire ropes is not perfectly periodic, but offers some regular periodicities, we used the detected irregularities for the computation of local, histogram-based features.

Finally, HOG features [13] were used as they are based on gradient orientations. The regular structure of the rope features articulated gradients within a certain orientation range, related to the twist direction. Gradients with a perpendicular orientation can be considered as noise or anomaly.

2.1   Linear Prediction Based Features

Linear prediction can be seen as one key technique in speech recognition [16]. It is used to compute parameters determining the spectral characteristics of the underlying speech signal.

The behavior of the underlying signal is modelled by forecasting the value $x(t)$ of the signal x at time t by a linear combination of the p past values $x(t-k)$ with $k = 1, \ldots, p$, where p is the order of the autoregressive process. The prediction $\hat{x}(t)$ of a 1-D signal can be written as

$$\hat{x}(t) = -\sum_{k=1}^{p} \alpha_k x(t-k) \tag{1}$$

with the following prediction error

$$e(t) = x(t) - \hat{x}(t) = x(t) + \sum_{k=1}^{p} \alpha_k x(t-k) \ . \tag{2}$$
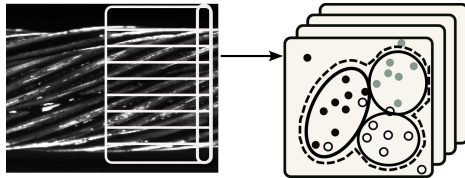
Fig. 2. Multichannel version of the classification model. For every channel (horizontal white boxes) a feature is extracted and examined in a separate feature space. The vertical white box marks the signal values, which are actually predicted.

This motivates the choice of linear prediction for feature extraction. For the prediction of the actual value the contextual information of the past values is used and is implicitly incorporated in the resulting feature.

Based on a least-squares formulation, the optimal parameters $\boldsymbol{\alpha} = (1, \alpha_1 \ldots \alpha_p)$ can be obtained by solving the normal equations [17]. The optimal coefficients are derived by use of the auto-correlation method and the Levinson-Durbin recursion [17, 16]. Free parameters of this method are framesize and the order of the process. In experiments the optimal framesize was found to be 20 camera lines with an incremental overlap of 10 lines. Best results were achieved for order $p = 8$.

Rope data, obtained from the acquisition system, can be seen as a sequence of 2-d images. Thus, with 1-d linear prediction it is not possible to analyze the 2-d signal. To overcome this, the rope data is considered as a multichannel time series. The signal $\boldsymbol{x}$ consists of c channels $\boldsymbol{x} = (\boldsymbol{x}_1 \, \boldsymbol{x}_2 \ldots \boldsymbol{x}_c)^T$ and every channel represents a 1-dimensional time series $\boldsymbol{x}_i(t) = (x_i(1), \ldots, x_i(t))$. For every channel i of this signal an individual 1-d linear prediction is performed, leading to the estimate $\hat{x}_i(t)$, the squared prediction error for the whole frame, and the coefficient vector $\boldsymbol{\alpha}_i$. These components are used as corresponding feature for the actual frame and the channel i. Best results were obtained with a combined feature vector, including prediction coefficients and the squared error. In the training step, a separate model for every channel is learned. This is schematically depicted in Fig. 2. By this, the different appearance of the rope at different positions in the images is taken into consideration.

2.2 Local Binary Pattern

For the local binary pattern (LBP) a texture region is seen as a joint distribution of $P + 1$ pixel-graylevels in a predefined neighborhood. Often a circular neighborhood with radius R and P equally spaced samples is chosen. The center pixel grayvalue g_c serves as threshold for the binarization of the neighborhood pixels g_p, $p = 1, \ldots, P$. The local binary pattern operator can be summarized as

follows:

$$LBP_{P,R}(g_c) = \sum_{p=1}^{P} s(g_p - g_c)2^{p-1} \quad \text{with} \quad s(x) = \begin{cases} 1, & x \geq 0 \\ 0, & x < 0 \end{cases} . \qquad (3)$$

Transforming the binary vector into a decimal number (3) results in a pixel label, based on the neighborhood information. Ojala et al. [18] developed a rotational invariant and uniform extension of the local binary pattern. For the anomaly detection task there is no need for rotational invariance due to the constant rope orientation. The uniformity of the pattern is defined based on the number of $0/1$ transitions U in the binary vector. The resulting LBP code is computed as follows:

$$LBP_{P,R}^{u}(g_c) = \begin{cases} \sum_{p=1}^{P} s(g_p - g_c), & U \leq 2 \\ P + 1, & \text{otherwise} \end{cases} . \qquad (4)$$

A histogram with a predefined number of bins is built from the underlying code distribution and serves as feature. The optimal parameters P and R and the number of quantization levels for the local histograms were determined in extensive experiments. We found the optimal parameter setting to be $R = 1$, $P = 8$ with 16 quantization levels for the histogram. As already mentioned, defects usually have just a small elongation. Hence, the histogram computation resulting in the feature vector is done for a small detection window (20×20 pixels), which moves over the underlying frame of rope data. By this, more than one feature is obtained for every frame.

2.3   Co-occurrence Features

Features for texture classification based on co-occurrence matrices were first introduced by [7]. A co-occurrence matrix is defined with respect to a certain displacement vector $\boldsymbol{d} = (d_x, d_y)$ and results in the joint distribution of co-occurring grayvalues. The relative frequency p_{ij}, which defines the co-occurrence of two neighboring grayvalues (with respect to $\boldsymbol{d}$) i and j, is defined as

$$p_{ij}(\boldsymbol{d}) = \lambda | \{(x,y) : I(x,y) = i, I(x + d_x, y + d_y) = j\} | \qquad (5)$$

with $i, j \in \{0 \ldots G - 1\}$ and G the number of gray levels. I represents an image of size $M \times N$ and λ is a normalization factor such that $\sum_{ij} p_{ij}(\boldsymbol{d}) = 1$.

Harlick introduced 14 different textural features [7]. Experiments for the determination of the most discriminative ones were performed. As the information theoretic texture features named difference entropy, information measure one, information measure two and the maximum correlation coefficient lead to the best results, a combination of these four features is used. Furthermore, a parameter evaluation resulted in an optimal displacement vector of 2 pixels length with an angle of 90 degrees. As co-occurrence matrices lead to a global representation of the underlying texture, they are usually computed for a local region of interest. For the detection of small, regional anomalies in the rope structure this is important, as small defects will not be recognized with global features. Again, a detection window of 20×20 pixels was used for the feature computation.

2.4 Features Based on Pattern Irregularity

Vartianinen et al. [11] describe an approach for irregularity detection in regular patterns based on the Fourier transform. By filtering out the distinct frequency peaks of a regular pattern in the Fourier domain, followed by an inverse transformation a perfectly regular pattern can be obtained. On the other hand, it is possible to substract this regular part from a unit function in the frequency domain, which results in the irregular part of the pattern:

$$I(x,y) = \mathcal{F}^{-1}(I(u,v)) \tag{6}$$
$$= \mathcal{F}^{-1}((\mathbf{1} - \mathcal{M}(u,v) + \mathcal{M}(u,v))I(u,v)) \tag{7}$$
$$= \underbrace{\mathcal{F}^{-1}(\mathcal{M}(u,v)I(u,v))}_{\text{regular part}} + \underbrace{\mathcal{F}^{-1}((\mathbf{1} - \mathcal{M}(u,v))I(u,v))}_{\text{irregular part}} \ .$$

$I(x,y)$ is the input image, $I(u,v)$ is the Fourier transformed image, $\mathcal{F}$ is the Fourier transform and $\mathcal{M}(u,v)$ is the filter function in the frequency domain. $\mathbf{1}$ represents the unit function. Without prior knowledge about the pattern structure a reasonable filter function is self-filtering [19]. Filtering is performed with the magnitude of the Fourier spectrum $\mathcal{M}(u,v) = |I(u,v)|$. As the rope consists of regular structures, filtering is done with regard to the irregular part of the data. For the computation of the local histograms again a detection window of size 20×20 is used. Experimental evaluation has led to an optimum of 16 quantization levels for histogram computation.

2.5 HOG features

Descriptors based on histograms of oriented gradients can be computed from gradient images. For predefined portions of the input image, the cells, an evenly spread gradient orientation histogram is computed given a predefined number N of histogram bins. Thereby, the histogram entries are weighted by the gradient magnitude. In a last step a descriptor or feature vector is formed, by concatenating and normalizing the cell histograms given a larger block consisting of $m \times n$ cells. In our experiments in analogy to the detection windows cell sizes of 20×20 pixels are used, with a block dimension of $m \times h$. h is the height (in pixels) of the segmented rope and m was chosen to be of 20 pixels width. We found out, that histograms with just 4 bins performed best in our application, as the distinct number of gradient orientations in the rope data is limited. As it is important for rope analysis, to discriminate between noise (caused by dirt or reflections) and actual defects, we further add the entropy given the gradient orientations of each HOG cell to the feature vector, which results in a dimension $d = \frac{h}{n}(N+1)$. In our case d is 35.

3 ONE-CLASS CLASSIFICATION

In order to exclude as many rope meters as possible from a further inspection, the theory of one-class classification seems to be a good choice. A separation between faultless and faulty samples is required. In this case, the faultless samples

represent the target class ω_T and the defects are considered as outliers ω_O. As it is no problem to construct a large training set of defect-free feature samples, a representation for this target density $p(\boldsymbol{x} \mid \omega_T)$ has to be found without any knowledge about the outlier density $p(\boldsymbol{x} \mid \omega_O)$ [20]. Here, $\boldsymbol{x}$ is the feature vector. For one-class classification problems the false negative rate (FNR) is the only rate which can be measured directly from the training data. The false positive rate (FPR) is the most important measure for defect detection, but cannot be obtained without a sample set containing a sufficient number of defective samples. In case of a uniform distributed outlier density, however, a minimization of the FNR in combination with a minimization of the descriptive volume of the target density $p(\boldsymbol{x} \mid \omega_T)$ results in a minimization of the FNR and FPR [20].

There exist many different methods for one-class classification (also called novelty detection) [21, 22]. In [3] two approaches, namely the K-means clustering and a Gaussian mixture model (GMM), were compared. In contrast to our former work, where the training sample set contained only faultless samples, the learning step is now modified. Model learning is performed on a sample set with intentionally included samples from defective rope regions. The aim is to evaluate the robustness of the method against outliers in the training set. This would reduce the need of a human inspector, who determines an optimal, faultfree rope region for model learning.

3.1 Decision Making

For a classification into target class and outliers, a threshold is defined on the density. This threshold is based on the mean and the minimal probability reached in the training. It is stated, that an optimal threshold should be within the range of mean and minimum probability. As the training samples are all considered as defect-free, the minimum probability gives the lower bound for the likelihood of faultless samples. To account for possible outliers, the threshold is varied in this range and the evaluation is done by means of receiver operating characteristics (ROC). Since anomaly detection is a security relevant application, it is important not to miss any defect. As a consequence, the optimal threshold maximizes the TPR (number of samples correctly classified to the target class) while keeping the FPR zero.

Due to feature extraction with a detection window (or in case of the LP features based on one single channel), a rope frame consists of more than one feature. Accordingly, the decision for the overall frame is based on the decisions for the single windows/channels. In case of feature extraction by a detection window, the frame is classified as outlier, if one of the corresponding windowsor blocks (for HOG features) is rejected as outlier. For the channel-based LP features a further pre-processing is necessary. As the channels have no spatial extension like the detection window, one single channel is prone to noise. Therefore, a local channel-neighborhood consisting of 15 channels is scanned for potential outliers and only if the number of channel-votes exceeds five channels, the whole frame is rejected as defect.
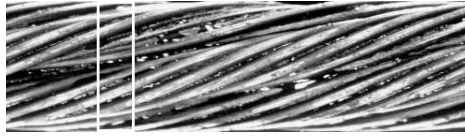
Fig. 3. In this rope region a wire is missing. The window marks the frame, which was detected as an outlier by the described system.

Since a defect usually lasts over several frames, the whole defect is regarded as detected, if one frame in this range is rejected as outlier. Consequently, defects are detected but not localized at the moment. Figure 3 displays one defect, detected by the described system. The borders of the frame, detected as outlier are depicted by the window.

4 EXPERIMENTS AND RESULTS

In the following section, experiments and their outcomes are presented. All experiments were performed on authentic rope data, acquired from real ropeways. In the generalization experiment (Subsect. 4.3) the data used for model learning was acquired in a controlled environment, but testing was again performed on data from a real ropeway. Model learning was done with a Gaussian mixture composed of five mixture components and rope data belonging to one of the four views. Testing was performed on all four views and the resulting ROC curves were averaged over the different views. Interference between the views was not yet considered. The length of the used rope regions in all experiments is given by the number of camera lines, followed by the corresponding length in meters put into brackets. Learning on 20.000 camera lines (2m rope) of one view takes between 25 seconds and one minute on a Intel Pentium 4 with 3,4GHz, according to the choice of features. Surely, the LP model learning needs the most time due to the separate computation of one model for every channel. In testing, we reach an average detection speed of 10-25 seconds per meter of rope (10.000 camera lines).

4.1 Comparison of Features

To compare the performance of the different features, model learning was done for every feature on the same training set, consisting of 20.000 lines (2m rope) of defect-free rope data from a real ropeway. Experiments were performed on a connected region of rope data, containing 600.000 camera lines (60m rope) and covering all known defects in the rope. The receiver operating characteristics in Fig. 4 point out, that the HOG features outperform all other features, since the goal is to maximize the TPR for an FPR= 0. However, for all used features respectable results were obtained, whereas the context-sensitive features seem to be the best choice. Their overall characteristics show a more robust behavior.
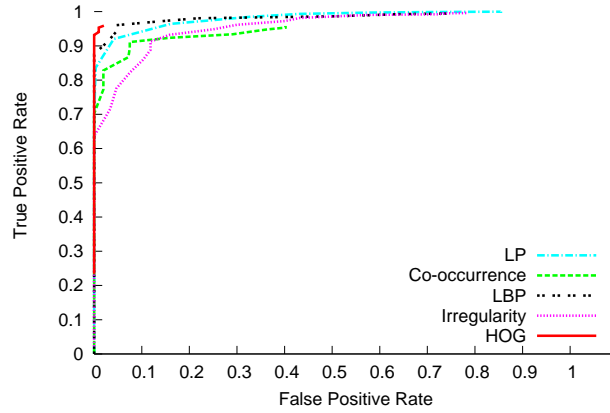
Fig. 4. Comparison of the ROC curves for the different choice of features.

Features based on detected irregularities perform the worst. A lot of noise is contained in the rope raw data and the structure is not perfectly regular, so that a certain amount of irregularities are detected in every frame. This results in a less discriminative behavior. Table 1 summarizes the maximum TPR for every view and feature, which was reached for a FPR of zero. For the HOG features sometimes a threshold equal to the minimum probability obtained in training was not enough to result in an FPR greater zero. This emphasizes the discriminative ability of these features. The reason for the decreased TPR of the LP features in view 4 is just one underdiagnosed error. A manual inspection depicted that this defect is spread over more than one view and was discovered correctly in the remaining views. Accordingly, results could be improved by incorporating interference between the different views. In summary, these results reveal the importance of context-sensitive features for the challenging task of defect detection in wire ropes.

Table 1. Comparison of the maximum TPR for a FPR of zero for all features and all views.

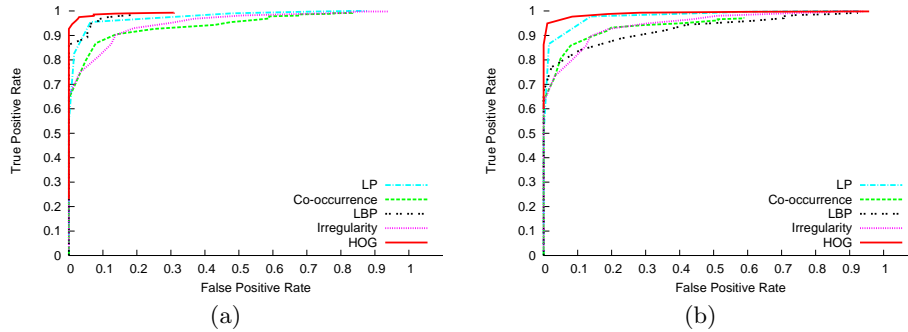| view#	LP	CoOccurrence	LBP	Irregularity	HOG
view1	0.96	0.78	0.96	0.62	0.95
view2	0.93	0.77	0.82	0.89	> 0.97
view3	0.94	0.77	0.88	0.71	> 0.95
view4	0.62	0.88	0.93	0.78	0.90

Fig. 5. Comparison of ROC curves for learning with a defect-free training set (a) and a training set including defects (b).

4.2   Robustness to Outliers

To evaluate the robustness to outliers in the training set, model learning was performed on a training set containing 200.000 (20m rope) lines of rope data. For learning, the view containing the most defects (9 defects) was chosen. Testing was performed on the remaining three views, also containing each at least seven defects. For comparison, the same experiment was performed with a model, learned from 200.000 defect-free camera lines (20m rope). The resulting ROC curves are compared in Fig. 5. The ROC curve in Fig. 5(a) gives the averaged ROC for the model, learned on defect-free training data. Figure 5(b) visualizes the results obtained with a model, learned from a training set including outliers. Obviously, the method is robust to few outliers in the training set, as the results differ not significantly from each other. Where especially the HOG and LP features show the most robust behavior, the LBP features seem to be error-prone if outliers are contained in the training set. The size of the training set was increased in the experiment, to incorporate as many defects as possible.

4.3   Generalization Ability

For the evaluation of the generalization ability, learning was performed on a real, faultless rope, acquired in a controlled environment. Testing on the other hand was performed on different rope data from a real ropeway containing defective regions. Both ropes belong to the same construction type and they only differ in their diameter by 10 pixels. In Fig. 6 the results are depicted by the corresponding ROC curves, averaged over all views. Figure 6(a) is generated by learning and testing on the same rope from the ropeway and Fig. 6(b) shows the result for learning in the controlled setup and testing on real-life rope data. In both cases the size of the learning set was 20.000 camera lines (2m rope). The HOG
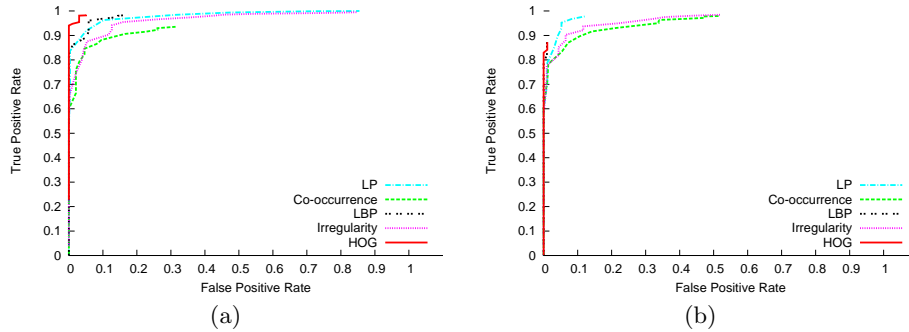
Fig. 6. Comparison of the ROC curves for learning and testing on the same data set (a) and learning and testing on two different, identically constructed rope data sets (b).

features still outperform all other features, but their performance reduces distinctly. The same can be stated for the LP features, whereas the global features like the co-occurrence features and the features based on irregularities seem not to suffer from a big loss of performance. Especially, for context-sensitive features this performance loss can be explained by the modified context, which is the fact in case of learning and testing on different datasets. But nevertheless, these results indicate a quite good generalization ability of the overall approach.

5   DISCUSSION AND OUTLOOK

As an important and meaningful extension of our former work [3, 12] we extended the one-class classification approach for anomaly detection in wire ropes by a performance comparison of five well-established and interesting features. The results emphasize the necessity of context-sensitive features for this challenging task and especially reveal the suitability of HOG features for this application domain. With the presented approach about 90 percent of the defect-free rope can be excluded from a reinspection by a human expert and only a region of 10 cm around a detected defect has to be re-examined again. Furthermore, experiments emphasizing the robustness and generalization ability of the approach were presented. They pointed out that a perfect, faultless training set is not essential for model learning. Especially from the practical point of view, this is an important insight, precisely because one cannot assure a completely defect-free training set. With regard to the generalization ability it was shown, that learning and testing on different datasets of identical constructed ropes still results in good recognition rates and therefore implies a good generalization ability of the learned models. Concerning the practical applicability, this is a remarkable

finding, as it is a difficult and time consuming task to learn an individual model of the respective rope previous to every detection run.

In the meantime our research already revealed a significant improvement of defect detection by the usage of context-based classification methods like Hidden Markov Models (HMM). The permit defect localization instead of pure detection [23]. Hence, future work will be focused on a fusion of all available rope views to make a robust decision for a whole rope section.

ACKNOWLEDGMENTS

References

1. EN 12927-7: Safety requirments for cableways installations designed to carry persons. ropes. inspection, repair and maintenance. European Norm: EN 12927-7:2004 (2004)
2. Moll, D.: Innovative procedure for visual rope inspection. Lift Report **29**(3) (2003) 10–14
3. E.-S. Platzer and J. Denzler and H. Süße and J. Nägele and K.-H. Wehking: Challenging anomaly detection in wire ropes using linear prediction combined with one-class classification. In: Proceedings of the 13th International Fall Workshop Vision, Modeling and Visualization. (2008) 343–352
4. Ojala, T., Pietikäinen, M., Harwood, D.: A comparative study of texture measures with classification based on featured distributions. Pattern Recognition **29**(1) (1996) 51–59
5. Tajeripour, F., Kabir, E., Sheikhi, A.: Fabric Defect Detection Using Modified Local Binary Patterns. EURASIP Journal on Advances in Signal Processing **8**(1) (2008) 1–12
6. Mäenpää, T., Turtinen, M., Pietikäinen, M.: Real-time surface inspection by texture. Real-Time Imaging **9**(5) (2003) 289–296
7. Harlick, R.M., Shanmugam, K., Dinstein, I.: Textural Features for Image Classification. IEEE Transactions on Systems, Man and Cybernetics **3**(6) (1973) 610–621
8. Chen, C.H., Pau, L.F., Wang, P.S.P., eds.: The Handbook of Pattern Recognition and Computer Vision (2nd Edition). World Scientific Publishing Co. (1998)
9. Iivarinen, J.: Surface Defect Detection with Histogram-Based Texture Features. Volume 4197 of Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series. (2000) 140–145
10. Rautkorpi, R., Iivarinen, J.: Shape-Based Co-occurrence Matrices for Defect Classification. In: Proceedings of the 14th Scandinaian Converence on Image Analysis (SCIA). (2005) 588–597
11. J. Vartiainen and A. Sadovnikov and J.-K. Kamarainen and L. Lensu and H. Kälviäinen: Detection of irregularities in regular patterns. Machine Vision and Applications **19**(4) (2008) 249–259
12. E.-S. Platzer and J. Denzler and H. Süße and J. Nägele and K.-H. Wehking: Robustness of Different Features for One-class Classification and Anomaly Detection in Wire Ropes. In: Proceedings of the 4th International Conference on Computer Vision Theory and Applications (VISAPP). Volume 1. (2009) 171–178

13. Dalal, N., Triggs, B.: Histograms of oriented gradients for human detection. In: Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR). (2005) 886–893
14. Kumar, A., Pang, K.H.: Defect Detection in Textured Materials Using Gabor Filters. IEEE Transactions on Industry Applications **38**(2) (2002) 425–440
15. Varma, M., Zisserman, A.: Texture Classification: Are Filter Banks Necessary? In: Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR). Volume 2. (2003) 691–698
16. L. Rabiner and B.-H. Juang: Fundementals of speech recognition. Prentice Hall PTR (1993)
17. Makhoul, J.: Linear Prediction: A Tutorial Review. Proceedings of the IEEE **63**(4) (1975) 561– 580
18. Ojala, T., Pietikäinen, M., Mäenpää, T.: Gray Scale and Rotation Invariant Texture Classification with Local Binary Patterns. In: Proceedings of the European Conference on Computer Vision (ECCV). (2000) 404–420
19. Bailey, D.: Frequency domain self-filtering for pattern detection. In: Proceedings of the first New Zealand Conference on Image and Vision Computing. (1993) 237–243
20. Tax, D.M.J.: One-Class classification: Concept-learning in the absence of counter-examples. Phd thesis, Delft University of Technology (2001)
21. Hodge, V.J., Austin, J.: A Survey of Outlier Detection Methodologies. Artificial Intelligence Review **22**(2) (2004) 85–126
22. Markou, M., Singh, S.: Novelty detection: a review - part 1: statistical approaches. Signal Processing **83**(12) (2003) 2481 – 2497
23. E.-S. Platzer and J. Nägele and K.-H. Wehking and J. Denzler: HMM-based Defect Localization in Wire Ropes - A new Approach to Unusual Subsequence Recognition. In: Proceedings of the 31st Annual Symposium of the German Association for Pattern Recognition (DAGM2009). (2009) to appear