

CLASSIFYING DEFECTS IN TOPOGRAPHY IMAGES OF SILICON WAFERS

Corinna Kofler
Gunter Spöck

Department of Statistics
Alpen-Adria-Universität Klagenfurt
Universitätsstraße 65-67
9020 Klagenfurt am Wörthersee, AUSTRIA

Robert Muhr

Infineon Technologies Austria AG
Siemensstraße 2
9500 Villach, AUSTRIA

ABSTRACT

In this work, we demonstrate that automatically classifying defects in topography images of silicon wafers is feasible. We process topography images of a set of sample wafers with controlled induced defects in their wafer back surfaces. We group these induced defects into three classes: cavities, cracks, and star cracks. With this sample set, we train and test selected classifiers with suitable feature vectors extracted from their wafer back surface topography images. A comparison reveals, that training and testing linear and quadratic classifiers with two Fisher scores as features, yield the best classification performances. We correctly classify all cavities and can separate them from the critical cracks and star cracks, which show a sufficient signal in the topography images.

1 INTRODUCTION

This paper presents our approach, demonstrating a way how automatically classifying defects in topography images of silicon wafers is feasible. These topography images are generated by a measurement method based on deflectometry. This deflectometric measurement method is applied for inspecting specular surfaces since years (Höfer 2011, pp.201-212), and is now entering the semiconductor market. It allows measuring high reflective surfaces which are virtually invisible. Thus, three-dimensional defects in these surfaces can only be observed by the distortions they cause in mirrored images of structured patterns. Such distortions in patterns with a defined structure allow determining the surface curvature. This principle is applied in the deflectometric measurement method, which provides topography images representing intensity coded curvature values. Hence, deflectometric measurements are utilized to control the quality of specular surfaces in industrial productions, allowing to find local defects. Fuller descriptions of the deflectometric measurement principle can be found in (Werling, Balzer, and Beyerer 2007), (Hsakou 2006), and (Muhr, Schutte, and Vincze 2010).

Currently, available deflectometric wafer inspection equipment provide promising data, allowing to assess the wafer surface quality. However, only little research to automatically classify the defects visible in these topography images of silicon wafers is available. Hence, we develop such a pattern recognition system. Recently, pattern recognition tasks are addressed with deep learning algorithms (Goodfellow, Bengio, and Courville 2016). However, it is challenging to interpret and understand these models, since they have kind of a black box-like characteristic. Furthermore, we want to gain knowledge about the features describing our defects. Therefore, we chose the statistical pattern recognition approach for our work. Our pattern recognition system comprises four stages: image segmentation, feature extraction, dimensionality reduction, and classification. For training and testing that pattern recognition system we need a sufficient sample set size. Thus, we have to produce a set of sample wafers with controlled induced defects in their wafer back surfaces. We chose to induce cracks and star cracks, since these are critical defects and affected wafers must

be sorted out. By artificially inducing these cracks and star cracks we additionally get regions containing continuous cavities. We consider these continuous cavities as uncritical defects in this work. Hence, we group these induced defects into three classes: cavities, cracks, and star cracks. We analyze these defects in Section 2. Then, we process topography images of our sample wafer back surfaces with a deflectometric measurement equipment. To reduce the amount of data we start with the defect segmentation as described in (Gonzalez 2009, pp.689ff), (Pernkopf and O’Leary 2002), (Narkhede 2013), and (Kaur and Singh 2002). We describe this defect segmentation, which is the first step in our 4-stage pattern recognition system, in detail in Section 3. In Section 4, we continue with the feature extraction. The goal of feature extraction is to find distinguishing features describing the objects to be classified. The feature values of objects belonging to the same class should be very similar. Whereas feature values of objects belonging to different classes are ideally very different from each other. Furthermore, it is important that the features are invariant to transformations, such as rotation, translation, or scale. (Duda, Hart, and Stork 2012). In Section 5, we explain the dimensionality reduction. Dimensionality reduction removes irrelevant and redundant data and thus increases the subsequent classification performance. Information about dimensionality reduction can be found in (Webb 2003), (Brownlee 2015), (Guyon and Elisseeff 2003), (Design 2015), and (PRTTools 2014). The last step in our pattern recognition system is the classification. Classification is the procedure of assigning objects based on their features, also called input variables, to one of a set of predefined target variables, also called classes. Hence, classification is a supervised learning method. Building a suitable model comprises training the classifier and testing the classifier. For this procedure, a set of sample data is required, where each data object consists of its input features and its corresponding class label. This sample set must be divided into a training set and a testing set. Training the classifier means determining the classifier based on the input features and the labels of the training set. Testing the classifier means applying the classifier to the testing set. The classes assigned by the classifier are then compared to the true class labels of the sample set. We use cross-validation to test the classifier and we measure the classifier performance based on the classification error as described in (Bishop 2006), (Webb 2003), and (Perner 2009). We evaluate classifiers and compare their performances in Section 6. Finally, we draw a conclusion and present an outlook in Section 7.

2 SAMPLE SET

In this section, we analyze the defects in detail to describe their physical appearances. Figure 1 shows one example of each defect class. These controlled induced defects are visible on the wafer back surfaces. Image (a) shows a cavity with a star crack, image (b) a cavity with a crack, and image (c) a continuous cavity. A measuring bar is displayed in the bottom right corner of each image. These images are captured with an optical microscope with a magnification of $2.5\times$. The field of view is $2.9\text{mm} \times 2.2\text{mm}$. This field of view only shows the center areas of the defects. The overall physical dimensions of the cracks and the star cracks in our sample set range from 5mm up to 15mm . The cavities have a diameter of about 10mm .

We further analyze one center of a cavity with a star crack by a white light interferometer measurement (Figure 2). The white light interferometer allows producing a 3D measurement of the wafer back surface (upper image). It shows a $580\mu\text{m} \times 380\mu\text{m}$ section of the wafer back surface, containing the center area of a star crack. Right beside this 3D measurement, a color bar indicates the heights. Its z-values range from $z = -1\mu\text{m}$ to $z = -3.5\mu\text{m}$. The horizontal blue line indicates the position of the profile depicted in the lower part of the illustration. This profile shows a height and a slope discontinuity at the crack. Such slope discontinuities cause disturbances within the cavity signals in the topography images.




Figure 1: Controlled induced defects in wafer back surfaces.




Figure 2: 3D measurement of a cavity with a star crack.

Figure 3 illustrates such a topography image of a silicon wafer as an example. It shows the intensity coded topography values of the wafer back surface in pseudo color representation. The green colored intensity value 128 represents the intact area of our wafer back surface having curvature values of zero or curvature values close to zero. The intensity values from 0 to 127 and from 129 to 255 represent areas of distorted stripe patterns originating from local controlled induced defects in the wafer back surface. These local defects are colored in the range from blue to red in our topography images. The yellow pixels in these images form the border areas of the cavities having low negative curvature values. The blue colored pixels represent the center areas of the cavities having positive curvature values. Some of the discontinuous cavities have red and dark blue pixels in their center areas. They arise from the curvature calculation at the position of the crack, where the surface has slope discontinuities. The green pixels between the border areas and the center areas form the transition regions of the cavities. The curvature values in these transition regions are close to zero too. These local defects are our regions of interest. They are marked with a blue rectangle and numbered in ascending order for the following defect segmentation step.

3 DEFECT SEGMENTATION

Analyzing the image data of the wafer back surfaces reveals, that thresholding is a promising image segmentation method for our task. It allows to identify the center point of possible defects and thus to precisely segment the regions of interest. Figure 4 shows an overview of the implemented image segmentation routine based on thresholding. The input image is the topography image of the wafer back surface. The outputs are the individual regions of interest containing possible defects, saved as sub-images of the wafer. The developed image segmentation algorithm can be divided into three major steps. In the first step, we derive the thresholds. The thresholds are required for binarizing the image of the wafer back surface in the second step. The result is a binary image representing the center areas of possible defects as

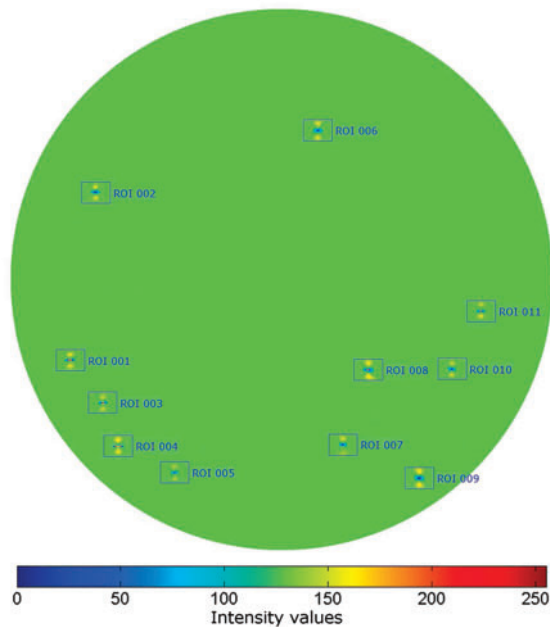


Figure 3: Topography image of a wafer back surface with numbered ROIs.

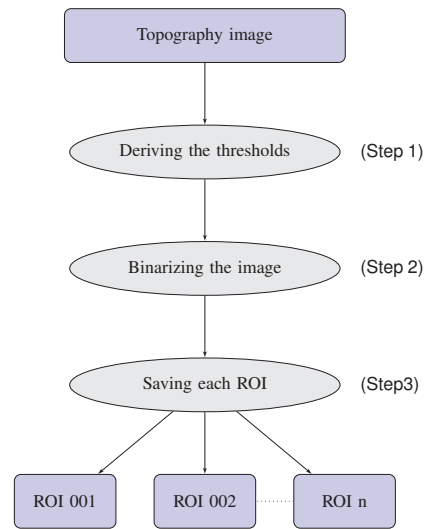


Figure 4: Segmenting the regions of interest.

connected components. In step three, we determine the areas of these connected components to identify possible defects. We save these possible defects as sub-images of the wafer. These are our regions of interest. All three process steps are described in detail in the following paragraphs.

In the first step, we derive the thresholds manually from the histograms of multiple regions of interest containing possible defects. The histogram (Figure 6) of the region of interest ROI001 (Figure 5) is a typical one of our sample set. It shows the number of pixels for each intensity value of the wafer back surface. We crop the y-axis at a value of 60 pixels to particularly visualize the number of pixels at low and high intensity values. The two thresholds $T_1 = 100$ and $T_2 = 180$, indicated by the red lines, separate the center area from the background. These two thresholds divide the histogram into three modes. We use the first and the third mode to select the center area of our regions of interest. We select them through binarizing the input image of the wafer back surface using the derived thresholds T_1 and T_2 in the next step.

The second step is converting the intensity image, which represents the topography of the wafer back surface, into a binary image. To binarize the image we use the previously derived thresholds from the histogram of the wafer back surface, $T_1 = 100$ and $T_2 = 180$. A pixel $p_{(x,y)}$ belongs to the foreground, if $I(x,y) < T_1$ or $I(x,y) > T_2$. A pixel $p_{(x,y)}$ belongs to the background, if $T_1 \leq I(x,y) \leq T_2$. Foreground pixels are set to the value '1' representing the objects, and background pixels are set to the value '0'. The result is a binary image derived from the topography image of the wafer back surface. Based on this binary image we identify the regions of interest in the next final step.

In this final third step, we identify possible defects in the binary image and save them as sub-images of the wafer. We extract the area of each object in the binary image of the wafer back surface to identify possible defects. We define possible defects as connected components with an area above 170 pixels. Then we calculate a rectangle with a size of $16\text{mm} \times 12\text{mm}$ centered around possible defects and crop each of these regions of interest. Finally, we save each region of interest as a numbered sub-image of the wafer. We use these sub-images for the feature extraction in the next section.




Figure 5: ROI001.

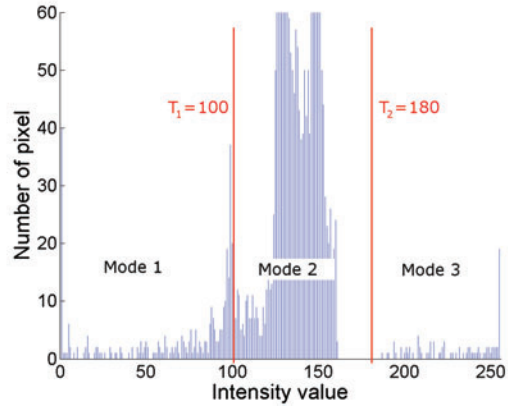


Figure 6: Histogram of ROI001.




Figure 7: Cavity with star crack.




Figure 8: Cavity with crack.




Figure 9: Continuous cavity.

4 FEATURE EXTRACTION

Figure 7 to Figure 9 show intensity coded topography signals of our regions of interest (left images). Analyzing them reveals that the major differences between the classes are in the center areas of the cavities. Hence, we calculate our features out of circular-shaped center areas with a diameter of 40 pixels, as depicted in the right images. The intensity values forming the center areas of cavities with star cracks typically range from intensity value 0 to 255. The center areas of continuous cavities typically range from intensity value 1 to 120 and the same applies to cavities containing cracks. Thus, some distinguishing statistical features out of these circular-shaped center areas must be extracted in order to properly classify them.

We extract the summary statistics, the histogram divided into 16 intensity bins, and the empirical isotropic semivariogram out of these circular-shaped center areas. To extract these features, we reshape our two-dimensional image data into a one-dimensional vector format. We calculate the summary statistics comprising: The count, arithmetic mean, median, mode, range, standard deviation, skewness, kurtosis, 10% - percentile, and the 90% - percentile of the one-dimensional intensity vector.

Furthermore, we extract the histogram divided into 16 bins

$$H(k) = \text{card}\{X \mid B_k \leq I(X) < B_{k+1}\},$$

$$k = \{1, 2, 3, \dots, 16\}, \quad B = \{B_1 = 0, B_2 = 16, B_3 = 32, \dots, B_{17} = 256\},$$

where k is the histogram bin, X is the column number of the intensity vector, $I(X)$ is the intensity value at column X , and B is the break vector.

Moreover, we calculate the empirical isotropic semivariogram

$$\gamma(h) = \frac{1}{2N(h)} \sum_{(s_i, s_j) \in N(h)} (z(s_i) - z(s_j))^2, \quad h = s_i - s_j,$$




Figure 10: Circular-shaped region of interest containing a cavity with a star crack, its histogram divided into 16 bins, and its empirical isotropic semivariogram.

where, $\gamma(h)$ is the semivariogram value for a specified distance range h , $N(h)$ is the number of data pairs falling within the specified distance range, s_i, s_j are the two locations of data points, and $z(s_i), z(s_j)$ are the intensity values of the two data points.

Figure 10 shows, from the left to the right, a circular-shaped region of interest, its extracted binned histogram, and its calculated empirical isotropic semivariogram as an example. The intensity histogram provides information about the intensity distribution (Gonzalez 2009, p. 120ff) and the empirical isotropic semivariogram describes the roughness of our regions of interest as described in (Esri 2015) and (Waller and Gotway 2004). To estimate the occurrence probability of an intensity value, we additionally normalize the histogram. To compare the shapes of different empirical isotropic semivariograms we additionally scale them to maximum. We add all extracted features, the summary statistics, the histogram bins, and the empirical isotropic semivariogram values to one feature vector containing in total 82 values for each circular-shaped region of interest.

Finally, we perform a univariate data analysis of our multivariate data as suggested in (Greenacre and Primicerio 2014). This analysis shows us four characteristics of our sample set. The first one is the class distribution of our sample set. We have in total 360 samples consisting of 172 samples of cavities, 115 samples of cracks, and 73 samples of star cracks. The second characteristic is, that we have significant differences in the class means at 43 features and several outliers in our data. The third characteristic is, that our feature vectors are not normally distributed over the classes nor within the classes. The fourth characteristic of our data set is, that the feature vectors have different scales. Thus, to compare the feature vectors, we standardize them using z-scores

$$Z = \frac{X - \mu}{\sigma},$$

where X is the vector containing the data values, μ is the arithmetic mean, and σ the standard deviation of the data vector X . With these standardized data, we can now start with the dimensionality reduction.

5 DIMENSIONALITY REDUCTION

We apply feature selection and feature extraction methods to reduce the dimensionality of our data set. Feature selection routines seek for n meaningful features that contribute to discrimination out of a set of m features. Features which do not contribute to a class separation are neglected with these feature selection methods. The feature extraction methods transform the m features to an n -dimensional feature space. Thus, all initial m features are used to create a reduced subset of n underlying features describing the data.

For feature selection, we apply to our sample set the sequential forward feature selection, the sequential backward feature selection, and the plus l - take away r feature selection routines of the PRTools. We run




Figure 11: Plus 1 - take away r feature selection.

each of the routines once using the sum of the squared Euclidean distances as separation criteria and once using the minimum of the squared Euclidean distances as separation criteria. Figure 11 shows the feature size curves of the plus 1 - take away r feature selection routine as an example. The feature size curve plots the averaged classification error on the y-axis against the corresponding number of used features on the x-axis. In this case, the number of features displayed on the x-axis refers to the feature position of the tested feature vector. All the feature vectors are ordered in decreasing discrimination performance. Hence, feature 1 is the best-performing feature of the feature vector, feature 2 the second best-performing feature and so forth. Analyzing all the feature size curves reveals, that the classification error decreases the more features we use for the classification. Overall, the major classification error reduction can be achieved by the top ten features. By adding further features, the classification improvements are comparatively minor.

Analyzing these top ten features of each feature vector shows, that they consist of the same features independent from the used feature selection routine. However, using the sum of the squared Euclidean distances as separation criteria returns other features than using the minimum of the squared Euclidean distances. Using the feature selection routine with the sum of the squared Euclidean distances as separation criteria, the top ten selected features consist of summary statistic features as well as empirical isotropic semivariogram values. Using the feature selection routine with the minimum of the squared Euclidean distances as separation criteria, the top ten selected features consist of empirical isotropic semivariogram values as well as histogram bins and normalized histogram bins.

For the feature extraction, we apply the principal component analysis (Webb 2003) and the Fisher linear discriminant analysis (Welling 2005). For the principal component analysis, we define the number of dimensions to be retained with ten. The retained ten dimensions are sorted according to their eigenvalues in decreasing order. Figure 12 presents this eigenvalue plot. The x-axis depicts the first ten dimensions and the y-axis reveals the corresponding eigenvalue of each dimension. This plot shows, that the eigenvalues are decreasing approximately following an exponential function. According to this plot we compare the averaged classification errors of selected classifiers applying the first two principal components 'PCAM 2D', the first three principal components 'PCAM 3D', the first five principal components 'PCAM 5D', and the first ten principal components 'PCAM 10D' as feature vectors (Figure 13). Using the first five principal components the K-NN(3), the K-NN(5), and the SVC (polynomial kernel 2^{nd} order) classifiers can yield an averaged classification error rate below 20%. Applying the first ten principal components the K-NN(3) and the K-NN(5) classifiers can yield an averaged classification error rate below 20%. Overall, the K-NN(3) using the first ten principal components as features reaches the lowest averaged classification error rate in this comparison. Hence, we apply the first ten principal components in the overall comparison.




Figure 12: Eigenvalues of the first ten dimensions retained by the PCA.




Figure 13: Performance of classifiers applying the PCA scores as features.

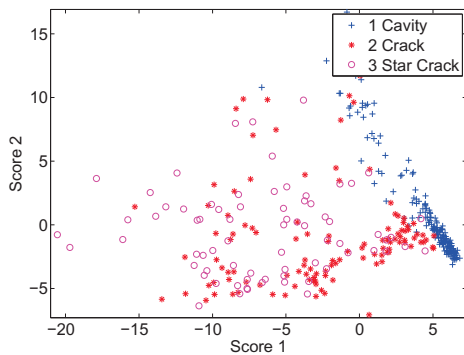


Figure 14: Scatter plot of the first two dimensions of the PCA scores.




Figure 15: Scatter plot of the two dimensions of the Fisher scores.

Figure 14 shows the scatterplot of the first two dimensions of our sample set retained by the principal component analysis. The samples are represented with different symbols according to their class. Samples belonging to class '1 Cavity' are depicted as blue crosses. The red stars represent samples of class '2 Crack' and samples of class '3 Star Crack' are drawn as pink circles. Examining this sample set reveals, that the cavities are reasonably separated from the cracks and the star cracks. Yet, the cracks and the star cracks are not separated in this two-dimensional scatterplot.

For the Fisher linear discriminant analysis, the number of mapped dimensions N is smaller than the number of classes C in the data set. Hence, we can only retain two dimensions of our three-class data set. Figure 15 shows the scatter plot of the two Fisher scores retained from our data set. Examining this scatter plot reveals, that the cavities are well separated from the cracks and the star cracks. Though, the cracks and the star cracks are not separated either.

Thus, we use these retained feature vectors from feature selection and feature extraction methods to train and test selected classifiers in the next section.

6 CLASSIFICATION

There are numerous classifier algorithms available. The difficulty is to preselect which classifiers are best suited for which kind of data set. A so-called cheat sheet for choosing the right classifier, in general, can

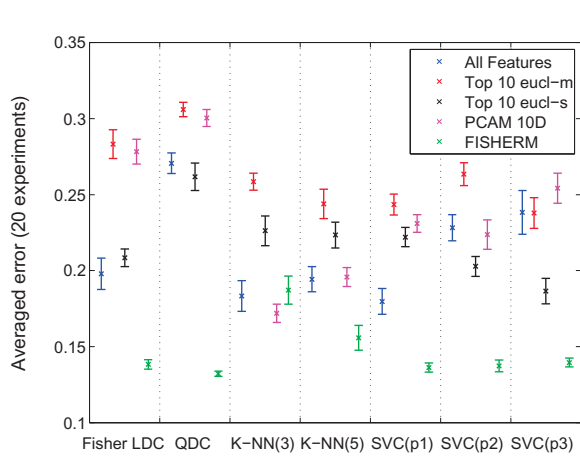


Figure 16: Performance comparison of selected classifiers applying different feature vectors.




Figure 17: Scatter plot of Fisher LDC.

be found in (scikit-learn developers 2014). This cheat sheet suggests us to start with the support vector classifier (SVC) using a linear kernel. If this is not sufficiently working, we should continue with testing the k-nearest neighbor classifier. If the k-nearest neighbor classifier (KNNC) is not working as demanded either, we should try it with support vector classifiers using other kernels. To broaden our test scenario, we also include the Fisher linear discriminant classifier (Fisher LDC) and the quadratic classifier (QDC).

Figure 16 shows the performances of the selected classifiers applying different feature vectors. We analyze the performance of the different classifiers separately on each feature set using the k-fold cross-validation with $k = 10$. Hence, our data set of 360 samples is divided into ten almost equally sized parts. Each time 9 folds are used for training the classifiers and the remaining fold is used for testing the classifier. We repeat this 10-fold cross-validation 20 times and compare the averaged classification error

$$error_{avg} = \frac{1}{N} \sum_{j=1}^N \sum_{i \in c} w_i \frac{n_{EL \neq i}}{n_{TL=i}}$$

and its standard deviation of the different classifiers. Where $error_{avg}$ is the averaged classification error based on error counts weighted by class priors, w_i is the class prior of class i , c are the classes, $n_{EL \neq i}$ is the number of wrongly classified samples for class i , $n_{TL=i}$ is the total number of samples belonging to the class i , and N is the number of experiments. This evaluation shows, that the classification performance is the best for most of the classifiers using the Fisher scores as features. The classification errors are below 15% for the trained Fisher LDC, the QDC, and all SVCs.

Figure 17 shows the scatter plot including the decision boundaries of the trained Fisher LDC of the 7th fold as an example. This example has the largest false positive rate

$$FPR = \frac{FP}{FP + TN}$$

of the class '1 Cavity'. Where FP is the number of false positives and TN is the number of true negatives. The figure title shows the classification error and the false positive rate of class '1 Cavity' of this 7th fold. The x-axis shows the first Fisher score and the y-axis the second Fisher score. The test samples are plotted according to their true class labels as blue crosses, red stars, and pink circles. The decision boundary of the trained Fisher LDC is drawn as a black line in the figure. Close to this decision




Figure 18: ROI004




Figure 19: ROI007.

boundary are the two false positives of class '1 Cavity'. Their topography images are included in the figure. The topography image of 'ROI004' only has a weak signal and the one of 'ROI007' is not extracted from the center area of the crack. The regarding microscopic images reveal, that 'ROI004' is a very weak crack (Figure 18) and that the 'ROI007' is extracted from the border region of the defect instead its center (Figure 19). This ROI007 demonstrates, that the image segmentation algorithm still needs to be improved for this region of interest. The left blue colored circle marks the center area of the crack, which should have been selected by the image segmentation routine. Instead, the right area marked by the red colored circle was segmented as the region of interest.

7 CONCLUSION & OUTLOOK

With this work, we demonstrated a way how automatic defect classification in topography images of silicon wafers is feasible. To realize this, we developed a pattern recognition system comprising image segmentation, feature extraction, dimensionality reduction, and classification. Our sample set consists of topography images of wafers with controlled induced defects. We grouped these defects into three classes: cavities, cracks, and star cracks. These defects induce strong signals in the topography images. In a first step, we segmented the defect regions by thresholding. In a second step, we extracted features out of the circular-shaped regions of interest containing the defects. The extracted features comprise the summary statistics, histogram values, and empirical isotropic semivariogram values. In the third step, we reduced the dimensionality of our feature vectors applying feature selection and feature extraction methods. In the fourth and last step, we trained and tested selected classifiers on our different feature vectors. A comparison revealed, that applying two Fisher scores as features perform best on most classifiers, reaching classification errors below 15%. We evaluated the classifier performances based on the averaged classification error of 20 experiments. Additionally, we thoroughly analyzed all classification results and can conclude that we are able to correctly classify all cavities. Furthermore, we can separate the cavities from the cracks and star cracks, which show a strong signal in the topography images. The fact, that we are not able to distinguish between cracks and star cracks is acceptable since either one must be sorted out anyway. The two wrongly classified circular-shaped regions of interest containing cracks are caused by a weak signal in the topography image and by a wrongly segmented region of interest, respectively. Hence, we can improve the classification performance by three influencing factors. Firstly, we can seek for further even more discriminating features. Secondly, we can improve the defect segmentation algorithm by using more complex segmentation methods compared to thresholding. Thirdly, we can improve the classification performance by evaluating and setting a classification threshold to reduce wrong classifications.

Another open and important issue is to verify whether cracks or star cracks can be detected if they do not lie in a cavity. Cavities create strong signals in the topography images, due to their curvatures. Cracks or star cracks cause discontinuities in such cavities and hence produce disturbances in the topography signals of these cavities. In this work, we classify these cracks or star cracks based on these disturbances they cause in the topography signals. However, the cavities are only a side effect of the controlled induced defects. Hence, a method to detect cracks and star cracks not lying in a cavity, as it is the case for productive wafers, is under development.

ACKNOWLEDGMENTS

The work has been funded by Infineon Technologies Austria AG and supported by funding from Austria under COMET K1 Zentrum.

REFERENCES

- Bishop, C. M. 2006. *Pattern recognition and machine learning*. Springer.
- Jason Brownlee 2015. “Machine Learning Mastery: An Introduction to Feature Selection”. Last visited April 12, 2017. <http://machinelearningmastery.com/an-introduction-to-feature-selection>.
- PR Sys Design 2006-2015. “perClass: Dimensionality reduction and data representation”. Last visited April 12, 2017. http://perclass.com/doc/guide/dimensionality_reduction.html.
- Duda, R. O., P. E. Hart, and D. G. Stork. 2012. *Pattern classification*. John Wiley & Sons.
- Esri 1995/2015. “ArcGIS Help 10.1: Semivariogram and covariance functions”. Last visited April 12, 2017. http://resources.arcgis.com/en/help/main/10.1/index.html#/Semivariogram_and_covariance_functions/003100000036000000/.
- Gonzalez, R. C. 2009. *Digital image processing*. Pearson Education India.
- Goodfellow, I., Y. Bengio, and A. Courville. 2016. *Deep Learning*. MIT Press. <http://www.deeplearningbook.org>.
- Greenacre, M., and R. Primicerio. 2014. *Multivariate analysis of ecological data*. Fundacion BBVA.
- Guyon, I., and A. Elisseeff. 2003, March. “An Introduction to Variable and Feature Selection”. *J. Mach. Learn. Res.* 3:1157–1182.
- Höfer, S. 2011. “Developments in the Field of Deflectometry”. *Proceedings of the 2010 Joint Workshop of Fraunhofer IOSB and Institute for Anthropomatics, Vision and Fusion Laboratory* 7:201–212.
- Hsakou, R. 2006. “Curvature: the relevant criterion for class-a surface quality”. *Ondulo, Visual technologies* 23:105–108.
- Kaur, M., and E. N. Singh. 2002. “Image Segmentation Techniques: An Overview”. *Issues* 2:y2.
- Muhr, R., G. Schutte, and M. Vincze. 2010. “A triangulation method for 3D-measurement of specular surfaces”. *International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences* 38 (5): 466–471.
- Narkhede, H. 2013. “Review of Image Segmentation Techniques”. *International Journal of Science and Modern Engineering (IJISME) ISSN* 1:2319–6386.
- Perner, P. 2009. *Machine Learning and Data Mining In Pattern Recognition: 6Th International Conference, Mldm 2009, Leipzig, Germany, July 23-25, 2009: Proceedings*. LNCS sublibrary: Artificial intelligence. Springer.
- Pernkopf, F., and P. O’Leary. 2002. “Visual inspection of machined metallic high-precision surfaces”. *EURASIP Journal on Applied Signal Processing* 2002 (1): 667–678.
- PRTTools 2014. “PRTTools Guide”. Last visited April 12, 2017. <http://www.37steps.com/prhtml/prtools.html>.
- scikit-learn developers 2010 - 2014. “Choosing the right estimator”. Last visited April 12, 2017. http://scikit-learn.org/stable/tutorial/machine_learning_map/.
- Waller, L., and C. Gotway. 2004. *Applied Spatial Statistics for Public Health Data*. Wiley Series in Probability and Statistics. Wiley.
- Webb, A. 2003. *Statistical Pattern Recognition*. Wiley InterScience electronic collection. Wiley.
- Welling, M. 2005. “Fisher linear discriminant analysis”. Technical report, Department of Computer Science, University of Toronto.
- Werling, S., J. Balzer, and J. Beyerer. 2007. “A New Approach for Specular Surface Reconstruction Using Deflectometric Methods”. *Computer* 2:S1.

AUTHOR BIOGRAPHIES

CORINNA KOFLER is currently pursuing the Ph.D. degree in the field of pattern recognition and image processing at the Department of Statistics at the Alpen-Adria-Universität Klagenfurt, Austria. She holds a Dipl.Ing. degree in Information Technology from the Alpen-Adria-Universität Klagenfurt, Austria and received her BSc degree in Systems Engineering from the Carinthia University of Applied Sciences (CUAS), Villach, Austria. Her e-mail address is corinna.kofler@aau.at.

GUNTER SPÖCK is Associate Professor of Statistics at the Alpen-Adria-Universität Klagenfurt, Austria. His research interests are in the fields of computational statistics, Bayesian statistics, stochastic differential equations, spatial statistics, sensor networks and spatial sampling design, and pattern recognition. His email address is gunter.spoeck@aau.at.

ROBERT MUHR is expert in the field of computer vision at Infineon Technologies Austria AG, Austria. He holds a Dipl.Ing. degree in Automation and Control from the Technische Universität Wien, Austria. His research interests include deflectometry, image processing, and pattern recognition. His email address is robert.muhr@infineon.com.