To appear in Proceedings of "Process and Materials Characterization and Diagnostics in IC Manufacturing" SPIE Vol. 5041.

Sidewall structure estimation from CD-SEM for lithographic process control*

Philip R. Bingham[†a], Jeffery R. Price[a], Kenneth W. Tobin[a], Thomas P. Karnowski[a], Marylyn H. Bennett[b], Hal Bogardus[b], and Michael Bishop[b]

[a]Oak Ridge National Laboratory, Oak Ridge, Tennessee, 37831-6010
[b]International SEMATECH, Austin, Texas, 78741-6499

ABSTRACT

In semiconductor device manufacturing, critical dimension (CD) metrology provides a measurement for precise line-width control during the lithographic process. Currently scanning electron microscope (SEM) tools are typically used for this measurement, because the resolution requirements for the CD measurements are outside the range of optical microscopes. While CD has been a good feedback control for the lithographic process, line-widths continue to shrink and a more precise measurement of the printed lines is needed. With decreasing line widths, the entire sidewall structure must be monitored for precise process control. Sidewall structure is typically acquired by performing a destructive cross sectioning of the device, which is then imaged with a SEM tool. Since cross sectioning is destructive and slow, this is an undesirable method for testing product wafers and only a small sampling of the wafers can be tested. We have developed a technique in which historical cross section/top down image pairs are used to predict sidewall shape from top down SEM images. Features extracted from a new top down SEM image are used to locate similar top downs within the historical database and the corresponding cross sections in the database are combined to create a sidewall estimate for the new top down. Testing with field test data has shown the feasibility of this approach and that the approach will allow CD SEM tools to provide cross section estimates with no change in hardware or complex modeling.

Keywords: Sidewall Structure, Semiconductor Manufacturing, content-based image retrieval, critical dimension

1. INTRODUCTION

Critical Dimension (CD) is key measure used in semiconductor lithography metrology. CD is a measure of the line width for printed lines on a wafer and is typically measured using CD scanning electron microscope (SEM) tools. These CD SEM tools use a top down image to measure line width of printed lines on the wafer. The yield management team uses the line width measurements to monitor the lithographic process and make corrections to keep the lines within the operation region. For process setup, test wafers with a printed Focus Exposure Matrix (FEM) are cross-sectioned to determine the optimal lithographic setting for a particular process. During this characterization step, top-down and cross-section SEM images are collected. As line widths shrink, the CD measurement will no longer provide sufficient information for process control, and a more complete measure of sidewall shape will be required. Figure 1 shows two 100nm lines that have different cross sections (one is under cut and the other is over cut). The top-down SEM images of these same lines have very similar width and could produce the same CD measurement. Cross sectioning of the wafer will allow monitoring of sidewall shape; however, this is a time consuming operation that results in destruction of wafers being tested. Since CD SEM tools that capture top down images are currently being used for measurement of CD, the next question is whether these same tools provide information that can be used to determine sidewall shape.

The Image Science and Machine Vision group of the Engineering Science and Technology Division at Oak Ridge National Laboratory has been involved with the used of Content Based Image Retrieval (CBIR) techniques to make use of image databases collected in the manufacturing environment [1,2]. Now we have applied the CBIR techniques to the estimation of sidewall shape by using the FEM data (top down/ cross section pairs) collected as a process database that can be searched with the top down imagery taken from a CD SEM tool during production to find images with similar

---

* Prepared by the Oak Ridge National Laboratory, Oak Ridge, Tennessee, 37831-6285, operated by UT-Battelle, LLC for the U.S. Department of Energy under contract DE-AC05-00OR22725.

† Correspondence: Philip Bingham, binghampr@ornl.gov; (865) 574-5680.

cross sections. The cross sections returned from the database query are then combined to estimate the cross section of the current product line. This effort required development of algorithms for extraction of sidewall shape from the cross section images, extraction of features from the top down images that distinguish between various sidewall shapes, and combining multiple cross sections returned from a query

The Sidewall Structure Estimator (SSE) has been developed into a software package, and testing of the SSE has been performed with focus exposure matrix image datasets and product test data provided by International SEMATECH (ISMT). Results from testing with the ISMT datasets have shown that (1) Top down SEM images encapsulate three-dimensional (3D) structural information similar to that which is made apparent through physical cross-sectioning; (2) The unique feature description developed for our feasibility study effectively captures and describes this 3D structure, and; (3) An image retrieval-based approach can be used to bound the problem of sidewall structural estimation and prediction by making use of a historical data repository. The following sections provide some background on CBIR and CD SEM, a detailed description of the sidewall structure estimator (SSE), and results of SSE field tests. Finally, conclusions and paths for further improvement of the SSE are presented.

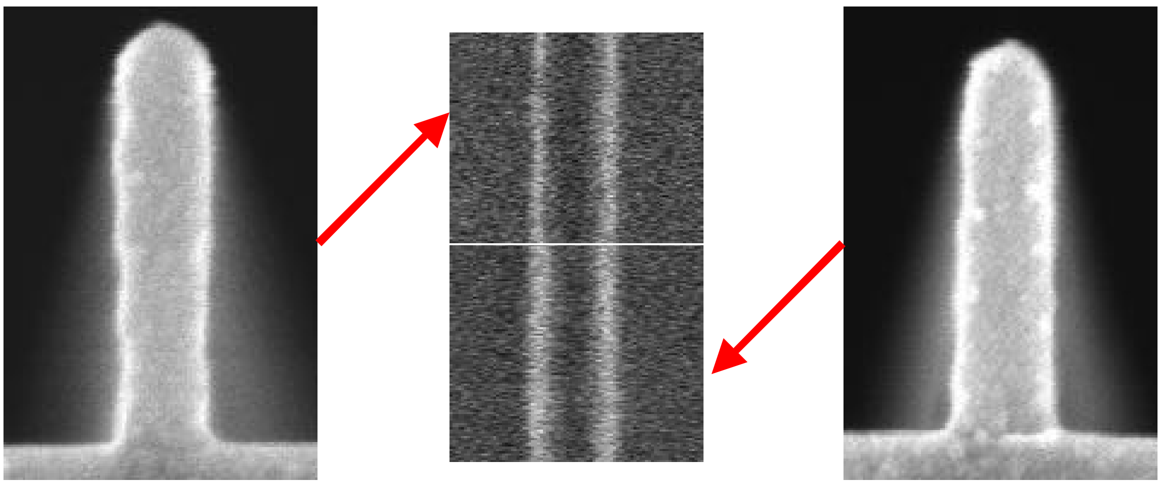

Figure 1. Two different cross sections with similar top down line widths.

2. BACKGROUND

In the development of an image from any sensor, many factors determine the response of the sensor to various surface types and shapes. For example, the illumination angle can greatly affect the image captured by a traditional camera. Top down SEM images are influenced by many factors of both the inspection surface topology and the SEM setup and geometry. Figure 2 shows an example of an electron beam striking the surface of three different line types from over cut to under cut with the electron detector positioned at an angle. A very basic interaction of the electron beam with the line material is shown and clearly indicates that the number of scattered electrons escaping from the material increases at the beam location as the line moves from over cut (left) to under cut (right). However, the backscattered electrons may not reach the detector depending on the detector position.

Figure 2 shows only a simple model of the electron beam interaction, but this example quickly shows the complexity of the SEM response to various shapes. The shape and properties of the surface material as well as the electron beam shape and power affect the interaction between the electron beam and the surface. A list of variables affecting SEM image formation follows as:

1. Electron beam energy,
2. Material properties,
3. Geometry of feature (width, height, wall angle, edge sharpness, etc.),
4. Relationship to surrounding features (isolated or dense region),
5. Charging effects (scanning speed, beam current, and all the above factors), and
6. Detector position, energy selectivity, and amplifier characteristics.

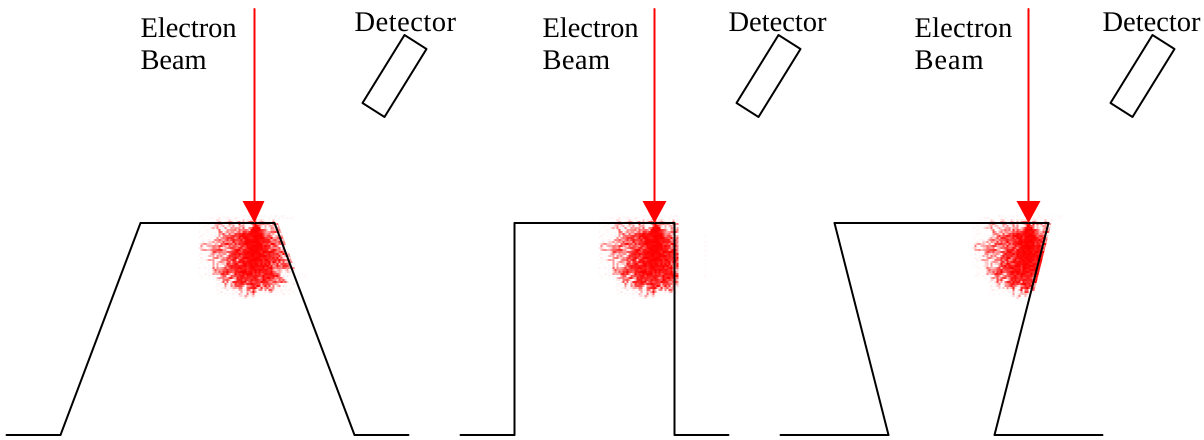

Figure 2. Electron beam interaction with various line shapes.

Based on variation in SEM interaction with different line shapes, the good news is that the sidewall shapes should affect the top down SEM image. However, interpretation of sidewall shape from a top down SEM image is a very difficult modeling problem due to the many factors involved (i.e. material properties, beam power, detector location, surface features, etc.). For this reason, we have taken a learning approach to determining sidewall structure from top down SEM imagery. A learning approach takes into account all of the parameters of the SEM and surface materials by using training data of known results taken from the same SEM setup and product material.

Training data is commonly available for this approach since FEMs are often used to characterize the process before production. Figure 3, depicts the setup for a FEM and provides an example of the data acquired from the FEM. At each of the locations in the matrix, a top down SEM image is acquired. Next, the FEM wafer is cleaved at each of the locations to be measured in the matrix and a cross section SEM image is captured. Thus, the FEM provides top down and cross section data over a range of focus and exposure settings, which results in a good sampling of the types of sidewall shapes that may be encountered during production. As long as the matrix provides ample representation of all possible sidewall shapes during production, a learning system is possible. Typically, the entire FEM is not cross-sectioned. Possible operating regions are selected using the top down SEM images and the operation region is further reduced by cross sectioning the regions selected from the top down imagery. Therefore, using the FEM as a training set may require more effort during process setup, but the gain of being able to accurately predict sidewall shape from top down SEM images should be well worth the additional effort. We expect that historical FEM data can be used for training of the system to reduce the dependence on acquiring a new FEM training set for each new product setup, but further testing is needed. For the development and testing of the SSE, International SEMATECH has provided ORNL with a variety of FEM datasets that span several design rules, line pitches, and CD SEM tools.

Since the sidewall estimation system uses an image database as the training set, Content Based Image Retrieval (CBIR) provides a mechanism for developing this image based learning approach. CBIR refers to techniques used to index and retrieve images from databases based on their pictorial content. Pictorial content is typically defined by a set of features extracted from an image that describe the color, texture and/or shape of the entire image or of specific image regions. This feature description is used in CBIR to index a database through various means such as distance-based techniques, approximate nearest-neighbor searching, rule-based decision-making, and fuzzy inferencing [3,4]. In the semiconductor yield management arena CBIR addresses a problem created by the growing proliferation of automated defect review and Automatic Defect Classification (ADC) technologies; i.e., the management and reuse of the large amounts of image data collected during review. For semiconductor yield management applications we have denoted CBIR technology as Automated Image Retrieval (AIR) [1,2]. *The fundamental premise of the AIR technology is that a similar process or phenomena likely generates images that are visually similar*. With this premise in mind, the top down SEM images taken of lines with similar sidewall shapes on a particular CD SEM tool using a consistent setup will be visually similar.

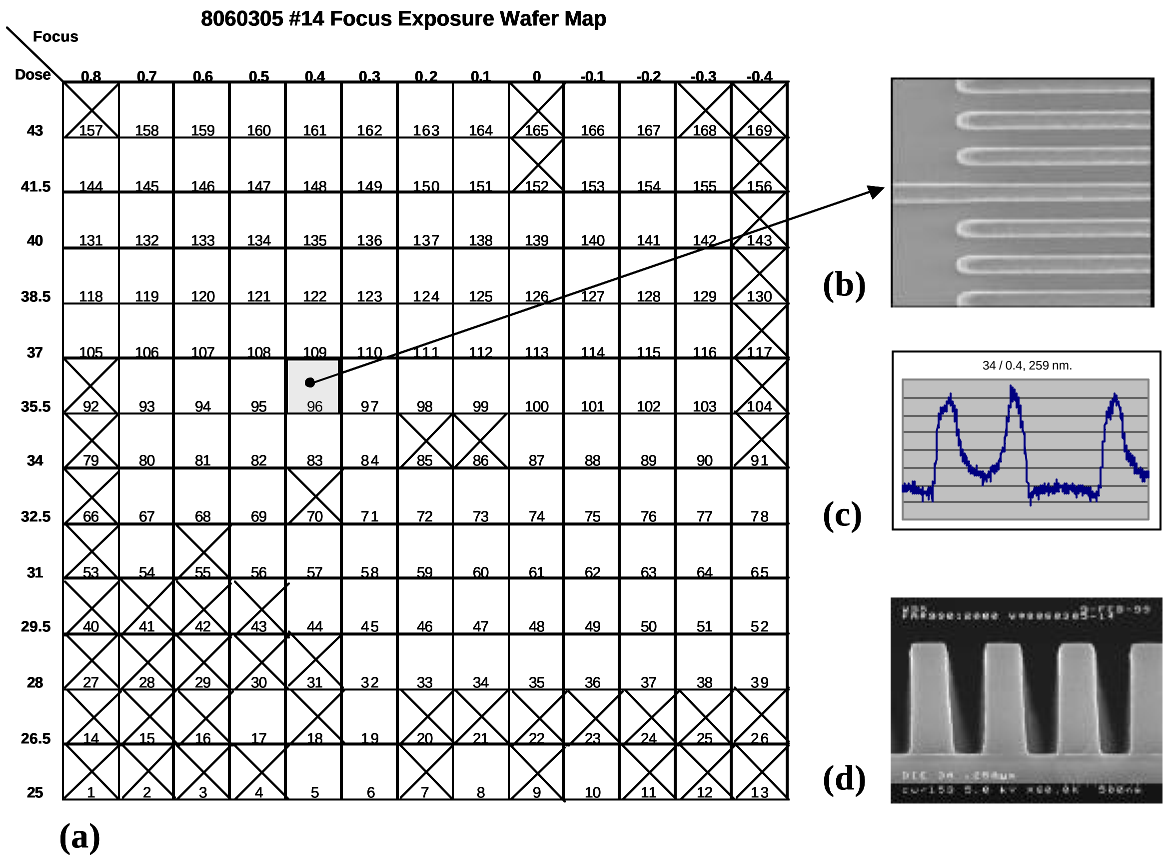

Figure 3. FEM data (a)Matrix describing focus exposure levels (b)top down SEM image, (c) linescan of top down SEM image, (d)SEM image of cross section.

3. SIDEWALL STRUCTURE ESTIMATOR

3.1 Overview

The Sidewall Structure Estimator (SSE) is the learning based system developed to estimate the sidewall shape of lines based on top down SEM imagery collected by CD SEM tools. The core of this system has been developed in C++ to test the concept and facilitate easy integration into CD SEM tool environments. Figure 4 illustrates the building blocks of the SSE. As in any learning system, a dataset representative of the types of cross sections that will be encountered is required. For the SSE, this dataset consists of top down SEM images and their corresponding cross section SEM images. In addition to a database that holds training data, the two major database interaction components are the build and query components.

The build component is used to enter cross section/top down pair training data to the system and consists of algorithms to extract information (features) from the top down images that distinguish between various cross section types and algorithms for extraction of sidewall shape parameters from the corresponding cross section images, which are used for sidewall shape estimation. The query component is responsible for submission of top down images for sidewall estimation, searching the database for similar top downs, and returning a sidewall estimate. Features must be extracted from a top down image submitted for sidewall estimation, so the query component uses the same algorithms for extraction of features from top down images used in the build component. Additionally, the query component contains algorithms for searching and combining similar cross sections to give a sidewall shape estimate. Cross section extraction, top down feature extraction, and sidewall estimation are the key modules of the SSE system. Descriptions of the interface and algorithms for these three modules are provided in the following subsections.

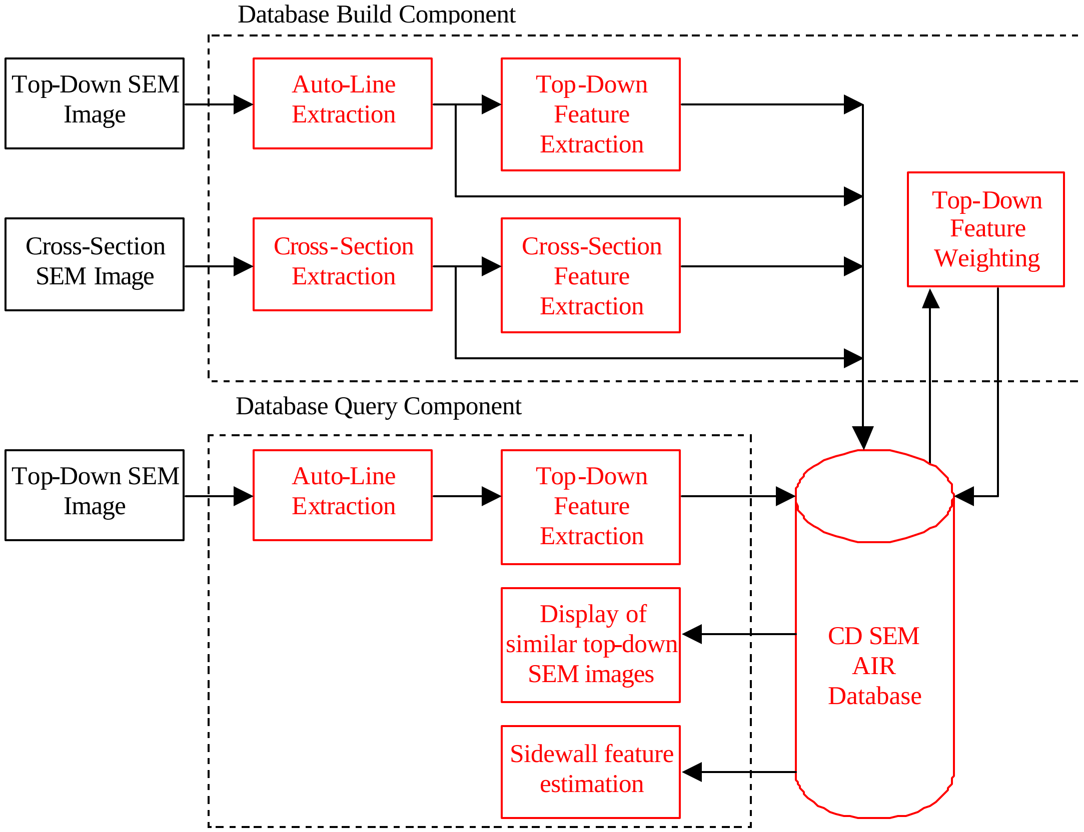

Figure 4. SSE block diagram.

3.2 Cross section extraction

Once a wafer has been cleaved and a SEM cross-section image acquired, the sidewall shape must be extracted from the cross section image. Initial attempts were to automate the extraction, but even with the small number of cross-section images we have encountered (<500 images), the variety in the images has made automation difficult. For this initial implementation, we have developed a semi-automatic algorithm that has been successful. Full automation of the sidewall extraction is not as important as for the top down feature extraction, because sidewall extraction only needs to be performed during the training process and not during sidewall estimation. In the current software version, the sidewall extraction process requires the user to measure the scale on the image, rotate the image such that the substrate is on the horizon, choose a line in the cross section image to be extracted, and adjust two threshold levels to properly extract the sidewall. When placed in a more constrained environment (same cross section SEM tool and setup), the sidewall extraction should be easily automated. As mentioned earlier, the cross section extractor need only be used during the training cycle. Therefore, everyone using the tool to estimate sidewalls need not learn how to extract the sidewall shape from a cross section image. This allows the tool to be easily used in the fab once a so-called expert has entered the training data.

Figure 5 is a screen shot of the sidewall extraction interface. The original cross section image appears on the left of the screen, and the white box indicates the subimage selected for edge extraction. Once the subimage is selected, histogram normalization is used to alleviate contrast variations experienced in the SEM images and scaling is performed to ensure cross section sidewall parameters in the training set are on the same scale. The resulting subimage is displayed in the top right of the interface. Edges are then extracted using the algorithm outlined in Figure 6.

A Gaussian filter is used to remove roughness on the cross sectioned line material. A first pass at edge extraction uses a Sobel edge filter followed by an edge threshold set by the user. This intermediate binary edge image is multiplied by the original subimage to place the intensity values of the cross section in the edge image. Since many of the line edges are bright, an intensity threshold set by the user allows further isolation of the cross section edges from other edges in the image. Finally clustering is performed and clusters smaller than 100 pixels are filtered out removing any spurious edges along the base of the line and those resulting from large rough features within the cleaved line. Notice that two thresholds must be supplied to this edge extraction algorithm. The two sliders on the interface control these two thresholds. With adjustment of these two thresholds, we have been able to use this sidewall extraction algorithm for all the cross sections we have experienced.

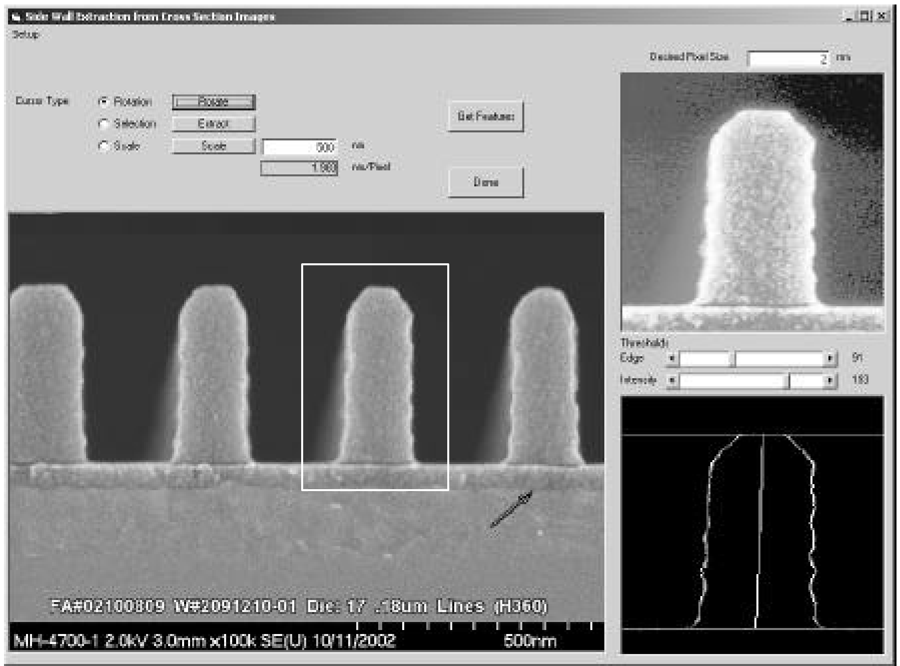

Figure 5. Cross section sidewall extraction interface.

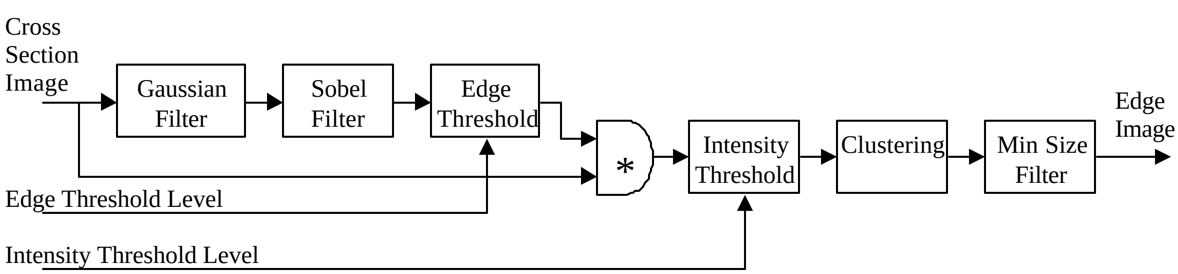

Figure 6. Edge extraction algorithm.

After edges have been isolated, a set of features uniquely describing the cross section is extracted. The features extracted are height, slope of centerline, and 100 width measurements corresponding to width from 1 to 100% of the height. By taking the width measurements equally spaced over the height of the line, the resulting width features are normalized between lines of various height allowing lines with similar sidewall shape and different heights to have similar line width feature arrays. While the algorithm is fairly complicated, users can quickly learn the effect of the two thresholds on edge extraction with experience, and often these parameters require little adjustment across a FEM training set.

3.3 Top down feature extraction

Top down feature extraction has been fully automated in the SSE. This full automation requires the user to provide pixel size, design rule, and design pitch information to accommodate line location and feature extraction. The algorithm follows a series of steps: rotation, line edge location, sub image extraction, and feature extraction for each subimage. A detailed description of the top down feature extraction can be found in [5]. In this process, a set of subimages is extracted from each top down image and features are extracted for each subimage to account for any variations that may occur across the image. Figure 7 presents a top down image with all of the subimages marked for extraction. The subimages are extracted and features are calculated for each subimage. These features consist of 13 image moments, 128 normalized features of an average linescan, and pixels corresponding to a 32x32 resized version of the subimage. Thus, feature extraction results in a 1165-dimension feature vector. For the example in Figure 7, the feature set consists of 12 feature vectors (1 for each subimage) having 1165 elements apiece.

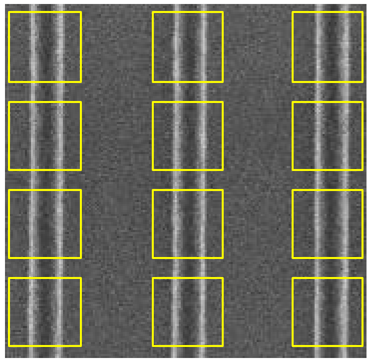

Figure 7. Top-down CD SEM image with sub-images to be extracted.

Although it is possible to compare top-down feature vectors in this 1165-dimensional space, it is computationally demanding and unnecessary since there are redundant features as well as features that are not helpful with respect to the cross-section estimation problem. Motivated by techniques that have been successfully applied to template-based face recognition [6,7], a principal component analysis (PCA) plus linear discriminant analysis (LDA) approach for dimensionality reduction was selected. With the above process in mind, we would like for the reduced-dimensionality feature vectors to preserve as much of the original separation between classes as possible. In the current software, each cross section in the training set is considered an individual cluster where each subimage extracted from the corresponding top down is a member of the cluster. Improvements are currently being tested in which clustering is performed on the cross sections to reduce the number of clusters and prevent separation of cross sections that have the same or similar shape.

The SSE training database contains the full 1165-dimensional feature vectors for each subimage extracted from a top down. In the software, the feature space reduction using PCA and LDA is called a search structure build. The build step produces a matrix used to reduce the feature vector from 1165 elements to 150 elements and attempts to separate the clusters as much as possible while keeping the components within clusters as close as possible. Although the build step is computationally intensive, it can be considered as part of the training and must only be performed once before beginning to estimate cross sections from top down SEM images.

3.4 Cross section estimation

The cross section estimation portion of the system accepts a top down SEM image as input, locates similar top down images (neighbors) in the training database, and combines the cross sections corresponding to the neighbors to obtain an estimate of sidewall shape for the submitted top down. In addition to the top down image for the query and the corresponding pixel size, critical dimension, and pitch parameters, the query requires input of the training datasets to search, the distance measure to be used, and the number of neighbors to return. All training sets entered into the software are available for search during the query. Distance measurement choices are minimum and average. When comparing two top down images, there are a number of subimages extracted from each. Minimum distance seeks to find the distance between the two closest subimages.

$$D_{\min} = \min(\mathbf{S}_1(i) - \mathbf{S}_2(j)) \quad i = 1 \cdots N_1 \quad j = 1 \cdots N_2$$

where $\mathbf{S}_n(i)$ is the feature vector for subimage i of top down image number n, and N_n equals the number of subimages extracted from top down image number n. Average distance is defined by

$$D_{avg} = \frac{1}{N_1} \sum_{i=1}^{N_1} \left\langle \min(\, \mathbf{S}_1(i) - \mathbf{S}_2(\,j)) \middle| \forall j = 1 \cdots N_2 \right\rangle.$$

Thus, the average distance is the average of the minimum distance between each subimage of the query top down and all the subimages in the database top down. Finally, the number of neighbors determines how many nearest neighbors will be returned from the query and used to estimate sidewall shape. All of these setup parameters are likely to stay constant when the system is being employed in a fab environment, so little if any user interface is required to obtain sidewall estimations once the system has been set up.

With the query setup parameters in place, a query (sidewall estimation) can be performed. The results displayed in the current SSE system are shown in Figure 8. This interface shows the query image (submitted top down), the nearest neighbors, and cross section estimation graphs and parameters. For each nearest neighbor returned, the top down image, the cross section image, and statistics associated with the pair are shown. These statistics include the height of the line, sidewall slope, radius of curvature for top and bottom, filenames within the database, and a score value. The score value provides a relative measure of the neighbors closeness to the query image calculated by $Score_i = D_1 / D_i$, where D_1 is the distance to the nearest neighbor. Therefore, the nearest neighbor has a score of 1 and a neighbor that is twice as far from the query image has a score of 0.5.

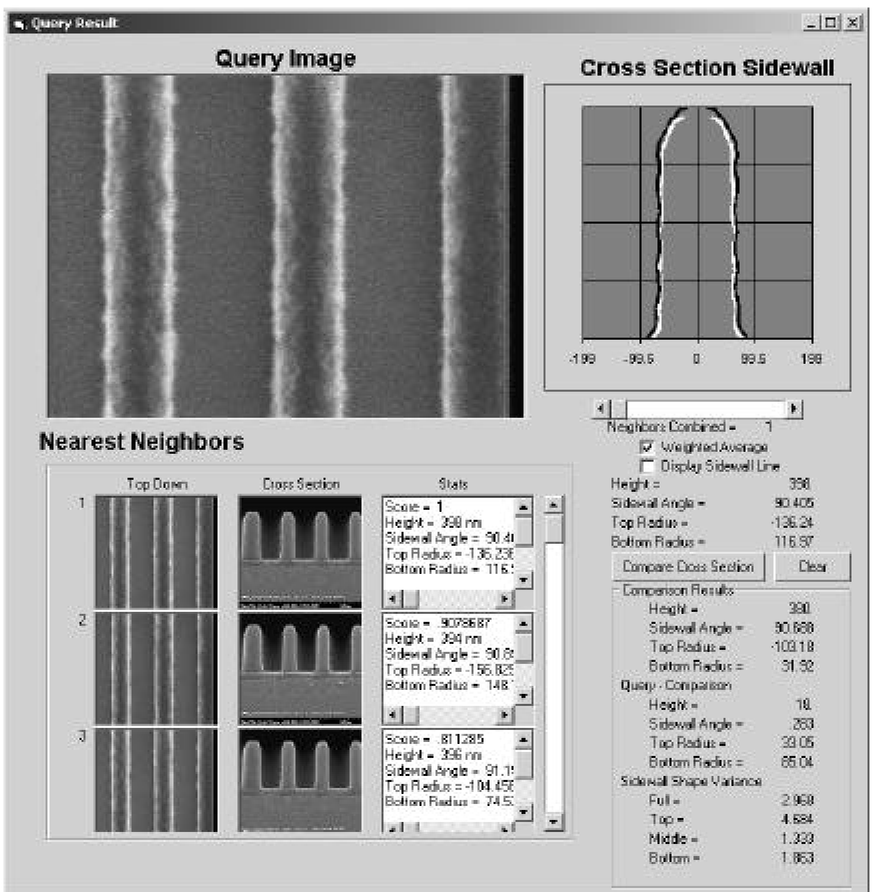

Figure 8. Query results interface.

The sidewall shape estimation is presented in the top right of the interface, and is represented by the black line in this graph. Using the slider immediately below, the operator can choose the number of neighbors to combine. Since the height is normalized by measuring a width value from 1 to 100% of the sidewall height, neighbor sidewalls are combined by averaging the cross section width vectors. Height is calculated independently by averaging the height of the neighbor cross sections. Two averaging choices are possible, straight averaging and weighted averaging. Weighted averaging weights each neighbor to be averaged by the score returned from the query for the associated top down. Using weighted averaging, the average cross section width is calculated as

$\mathbf{W} = \frac{1}{\sum_{i=1}^{N} Score_j} \sum_{j=1}^{N} \mathbf{W}_j * Score_j$, where $\mathbf{W}_j$ is the vector containing the 100 cross section width values for neighbor *j*, and $Score_j$ is the score for neighbor *j*. Weighted averaging is the recommended mode of operation since cross sections with top downs most similar to the submitted top down influence the result more than those farther away. Below the estimated sidewall graph, several other parameters of the estimation are presented. These values are height, sidewall angle, top radius of curvature and bottom radius of curvature and attempt to provide numerical values describing the returned sidewall shape.

The SSE has been developed to test the CBIR approach to sidewall estimation. For testing purposes, the software allows the user to evaluate the results by loading the known cross section associated with the top down query image. The example in Figure 8 shows the true cross section shape as the white line. Additionally, the numerical measures of the shape are included in the box at the bottom right along with the differences between these values for the estimated cross section and the true cross section. This comparison results box also contains four measures of sidewall shape difference. These four measures represent the same calculation over different regions of the sidewall. The four regions are the full sidewall, the top quarter of the sidewall, the middle half of the sidewall, and the bottom quarter of the sidewall. To measure sidewall shape difference, one could simply subtract the width vectors of the true sidewall and the estimated sidewall and calculate the average error, but this would include any width variations, which are not errors in sidewall shape. Therefore, we calculate the standard deviation of the width variations. This calculation provides a measure of sidewall shape difference with the average width variation removed.

4. EXPERIMENTAL RESULTS

In order to test the SSE, a field test was performed with data from wafers manufactured by International SEMATECH (ISMT). The following subsections outline the field test data collection procedure and the results obtained by the SSE system for the test.

4.1 Field test data

ISMT produced 25 wafers with a resist-on-poly pattern containing 180nm lines having a 1 to 1 pitch. For training, the first of these wafers consisted of a FEM. Top down images of the 180nm 1 to 1 features were captured on every die in the FEM using the Applied Materials NanoSEM tool. Next, each die was cross-sectioned, and a cross section SEM image was captured.

The remaining 24 wafers were produced at various focus/exposure levels spread through the FEM. For each of these focus/exposure levels, a wafer was produced and measurements made on four die locations spanning from the center of the wafer out to the edge. Top downs were taken with the Applied Materials NanoSEM tool followed by cross sectioning and capturing of cross section SEM images at each of the 4 die locations.

4.2 Field test results

The FEM from the first wafer was entered as training data, and the SSE was used to predict sidewall shape for each of the remaining field test data top down images. Figure 9 shows two examples of sidewall estimates (dotted lines) and the true sidewall shape (solid line) corresponding to the submitted top down image. Notice the estimates not only provide a good approximation to the top roundness, but the roughness of the sides and the footer shape are also closely estimated.

The standard deviation of the full sidewall shape differences between the estimates and the true sidewall shapes were calculated to measure performance of the system. These performance measures were made for results using 1, 5, and 10 neighbors and cumulative histograms of the resulting measures over the entire field test data set are shown in Figure 10. Two additional graphs are included in the results figure. The "Best Over FEM" line is a cumulative histogram of the minimum difference between the cross sections corresponding to the search top downs and all of the cross sections in the training data. Therefore, "Best Over FEM" represents the best possible cross section match in the training data. The "Average over FEM" line is a cumulative histogram of the average distance between query sidewalls and all of the cross sections in the training data. These results indicate that SSE performs well over 85 to 90% of the field test data; however, approximately 15% of the results have greater than average difference. Closer analysis of the results show that the 15% of cross sections that perform poorly are either much more narrow or much wider than the other 85% of the field-test data. Furthermore, the field test FEM does not contain cross sections as wide or as narrow as the images in the field test product data. Width only varies by 18nm in the field test FEM, but width varies 35nm in the field test product data.

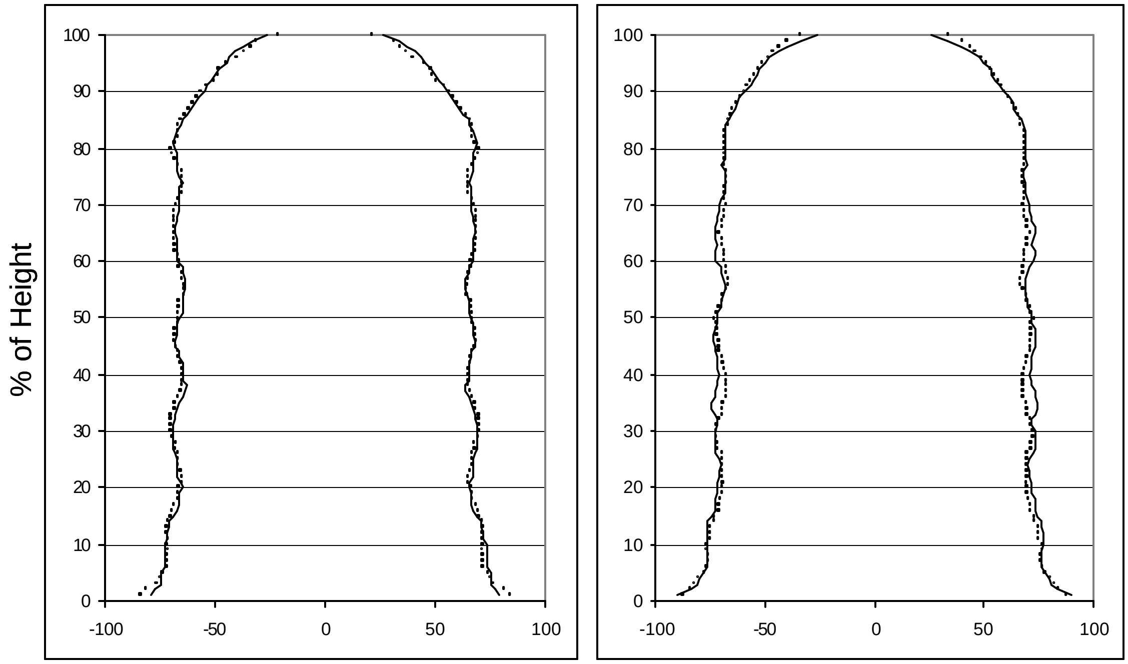

Figure 9. Examples of sidewall estimates (dotted line) versus true sidewall shape(solid line).

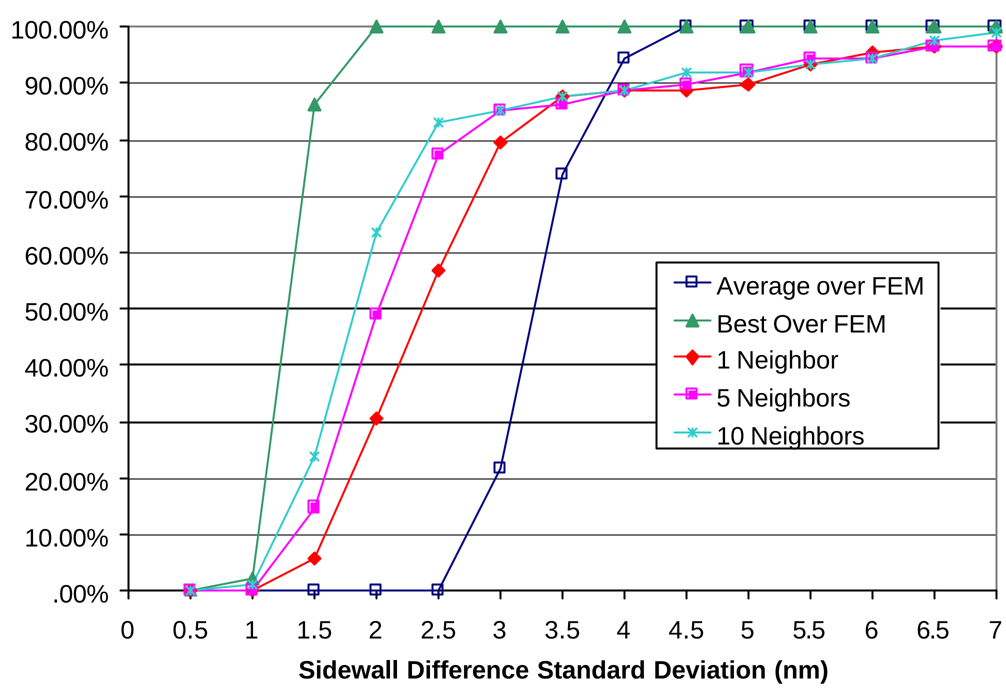

Figure 10. Field test results using field test FEM for training.

Many of the features selected for searching the top down image database have an inherent dependence on line width. Therefore, when a top down corresponding to a wide cross section with vertical sides is submitted and no cross sections with similar width and sidewall are present in the database, cross sections with overcut sidewalls appearing to have similar width in the top down are likely to be returned. In previous testing, we have performed hold one out testing on a given FEM, so the FEM always contained cross sections of similar width. Thus, the field test has proven very valuable in identifying this issue, and methods for dealing with this issue are being considered such as decoupling the width from the top down features to make the system immune to width differences. However, the premise that the training data must represent all possible production regions still holds and a more representative FEM is needed.

Through the course of this project, ISMT has provided several FEM datasets for testing. Therefore, we performed testing to determine the feasibility of using a historical data set containing cross sections with a large range of width variation. A FEM of 180nm 1 to 1 lines captured with a Hitachi CD SEM tool at ISMT having a width variation of 60nm within the set was selected for testing. Two test runs were performed. The first test run used only the Hitachi FEM for training, while the second used both the field test FEM and the Hitachi FEM for training the system. Results of these two tests are shown along with results from the previous test over the field test FEM in Figure 11.

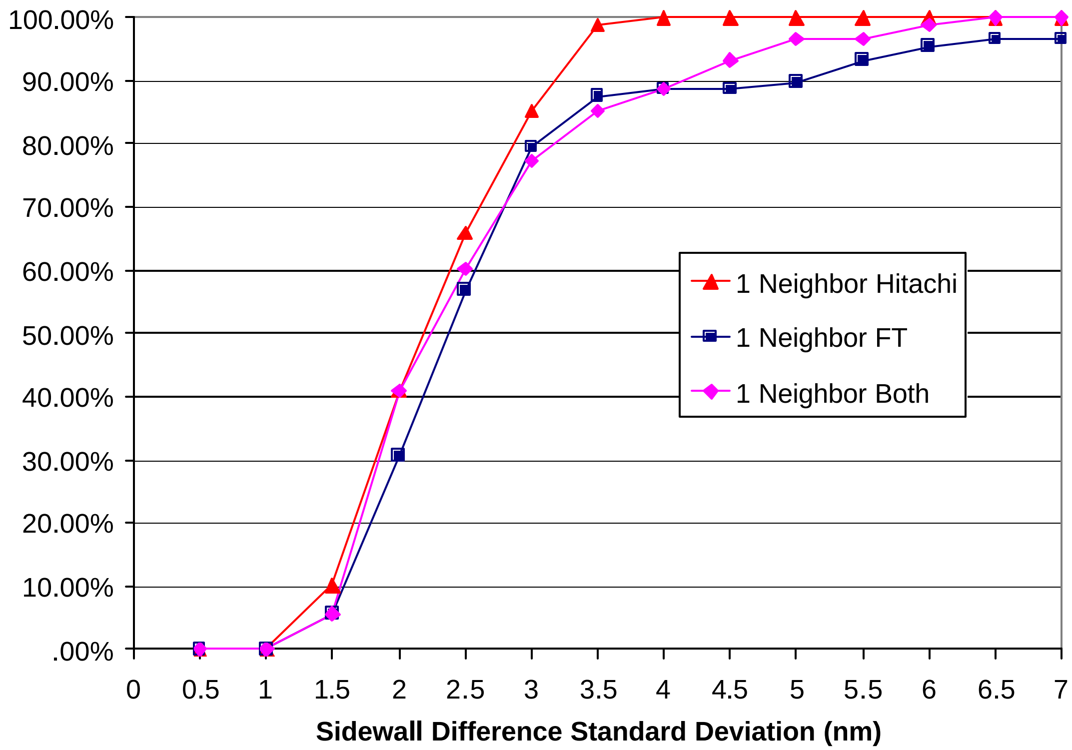

Figure 11. Field test results with various training sets.

Several interesting conclusions can be drawn from these results.

1. Training with both the Hitachi and Field Test FEMs improved the results slightly such that all of the estimates have a sidewall difference standard deviation from the true sidewall less than 6.5nm.
2. Training with only the Hitachi FEM provided the best results such that all estimates are within 4nm.
3. Why would the Hitachi FEM outperform the two together? This is likely due to differences in the images acquired by the Hitachi and Applied materials CD SEM tools. When both are included in the training, the image similarity between the field test product and FEM data overrides the Hitachi FEM. However, when training with only the Hitachi set, all of the training data has equal tool difference between itself and the training data; therefore, other features influenced by the sidewall shape begin to influence the results.

Finally, Figure 12 shows the average performance (standard deviation of sidewall difference) over all the field test product images using the Hitachi FEM for training as a function of combined neighbors from 1 to 10 neighbors. This graph indicates the best performance is achieved by combining the three nearest neighbors and has a value of 1.79nm for the average standard deviation in sidewall difference.

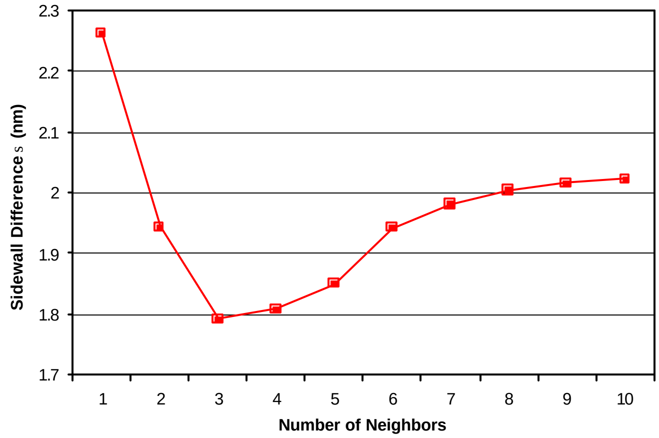

Figure 12. Performance versus number of neighbors.

4. CONCLUSIONS

A software tool has been developed that predicts the sidewall shape of lines on semiconductor wafers from top down SEM images by finding similar top down SEM images in a training database and combining the corresponding cross sections contained in the training database. This tool requires no new hardware and thus can be easily incorporated in the fab using top down SEM images from a variety of CD-SEM tools. Some effort is required to enter training data off line, but once the tool has been set up, the estimates can be provided automatically during production.

Field-testing has shown a dependence of the results on width. This points to the necessity of the training data to contain a good representation of all possible operation points. Additionally, we are investigating the possibility of decoupling width information from the top down feature set to improve performance. A second enhancement under investigation is the clustering of cross sections within the training set to keep similar cross sections near one another during the build instead of assuming every cross section is unique. Even without these enhancements, the SSE has shown excellent results in the field test. An aggregate performance of ±3.5nm was observed using the field test training data, and the performance improved to ±1.79nm with the use of a previous training dataset containing more variability.

ACKNOWLEDGEMENTS

Funding for this effort was provided by ISMT. We would like to thank the member companies of ISMT for their support and guidance. We would like to specifically thank Michael Bishop, Marylyn Bennett, and Hal Bogardus of ISMT.

REFERENCES

[1] K.W. Tobin, T.P. Karnowski, and R.K. Ferrell, "Image Retrieval in the Industrial Environment", IS&T/SPIE's 11th International Symposium on Electronic Imaging: Science and Technology, San Jose Convention Center, January 1999.

[2] T.P. Karnowski, K.W. Tobin, R.K. Ferrell, and F. Lakhani, "Content Based Image Retrieval for Semiconductor Manufacturing", IS&T/SPIE's 12th International Symposium on Electronic Imaging: Science and Technology, San Jose Convention Center, January 2000.

[3] V.N. Gudivada and V.V. Raghavan, "Content-Based Image Retrieval Systems", *IEEE Computer Magazine*, 0018-9162, September 1995, p. 18.

[4] M. De Marsicoi, L. Cinque, and S. Levialdi, "Indexing Documents by their Content: A Survey of Current Techniques", *Imaging and Vision Computing*, Vol. 15, 1997, p. 119.

[5] J.R. Price, P.R. Bingham, K.W. Tobin, ,and T.P. Karnowski, "Estimating Cross-Section Semiconductor Structure by Comparing Top-Down SEM Images," *Machine Vision Applications in Industrial Inspection XI*, Proc. SPIE, Vol. 5011, March 2003.

[6] J.R. Price and T.F. Gee, "Towards robust face recognition from video," in *30th applied Imagery and Pattern Recognition Workshop (AIPR)*, pp. 94-100, October 2001.

[7] P.N. Belhumeur, J.P. Hespanha, and D.J. Kriegman, "Eigenfaces vs. Fisherfaces: Recognition using class specific linear projection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **19**, pp. 711-720, July 1997.