DETECTING GLASS FIBERS USING COMPUTER VISION

Alfred S. Malowany Martin D. Levine Mussa R. Kamal†
Ronald Kurz Abdol-Reza Mansouri

Computer Vision and Robotics Laboratory
Department of Electrical Engineering
McGill University
Montréal, Québec, Canada

† Department of Chemical Engineering,McGill University, Montréal, Québec, Canada.

Abstract

We present an application of computer vision in the field of chemical engineering, the processing of images of glass fibers. Fibers are used as reinforcement for several polymer products to enhance their mechanical and thermal performance. In order to evaluate the properties of the polymer products however, a quantitative measure of fiber length and orientation is required.

The algorithm presented here locates the fibers present in a given image and thus enables their quantification in length and orientation. This algorithm operates in two steps. In the first step, feature points in the image are extracted in order to enable the generation of hypotheses as to the possible presence of fibers. In the second, the generated hypotheses are verified and the hypotheses that yield the highest confidence are retained.

1. Introduction

We present an application of computer vision in the field of chemical engineering, the processing of images of glass fibers. Fibers are used as reinforcement for several polymer products to enhance their mechanical properties and thermal performance. However, the process induced orientation distribution of fibers is anisotropic, and this results in direction dependent mechanical properties. Thus, a quantitative measure of fiber length and orientation is required for predicting the mechanical properties of the product.

In a typical image of glass fibers, the fibers have a high intensity, as opposed to a noisy background which has, on average, a low intensity. In addition, the fibers appear as straight line segments of differing lengths. Thus, looking for fibers, we actually focus our attention on detecting high intensity straight line segments on a low intensity and noisy background. A number of techniques have been developed in order to solve this problem. Relaxation labelling is one of them[1]: using this technique, an initial set of orientations is assigned to the image points, and these orientations are iteratively altered until the orientation at each point becomes the one dictated by its neighborhood. The main drawback of relaxation labelling and similar techniques lies in their excessively large computational complexity, which renders them impractical for most applications. Another technique that is in use for detecting line segments is the Hough transform[2]. The Hough transform is a mapping from image space to parameter space (usually taken as distance and orientation) in which colinear points in the image appear as clusters of points in the parameter space. However, the Hough transform does not distinguish between connected and non-connected points. This results not only in interference between points residing on different segments, but also in ambiguous interpretations of the clusters in parameter space.

We take a different approach for solving the fiber detection problem. For this, we adopt the well known hypothesis prediction/verification paradigm [3].This algorithm consists of two major parts: a predictor, which predicts possible fiber locations, and a verifier, which verifies the predictions and discards those which do not satisfy all of the constraints. Bearing in mind that template matching is itself a form of hypothesis prediction/verification (although of an exhaustive nature), it becomes clear that the main requirement that is imposed on the predictor is that it should significantly limit the scope of its predictions. In other words, the image-based features that are used by the predictor in formulating hypotheses should be closely correlated with the hypotheses themselves. In what follows, the features that are used in formulating hypotheses are described, and the cost function which allows the verifier to prune hypotheses is presented.

The complete algorithm is implemented in the C programming language on a VAX 11-750 running the UNIX operating system.

2. Predicting Fiber Locations

Consider a slide of glass fibers. The slide is illuminated by direct lighting and the reflected light pattern is viewed with a camera mounted on a microscope. This yields an image in which the regions occupied by the fibers have a large gray value, as opposed to the background gray value, which occupies, on average, the low

intensity range. Thus, in such images, the histogram is bimodal, and it is possible to differentiate between the fibers and the background by simple intensity thresholding. The threshold is selected using the p-tile method, whereby a certain percentage of the pixels are assigned to the background, while the remaining pixels are assigned high intensities and are identified as forming the regions occupied by the fibers.

In our experiments, the threshold is chosen in such a way as to assign 85 percent of the pixels to the background, and only 15 percent to the fibers. This specification is directly related to the density of the fibers. Since this density is relatively constant, the percentage value specification is invariant with respect to the type of images that is being dealt with. As a result of intensity thresholding, the image is partitioned into a set of regions, all of which (except for the background) have resulted either from the glass fibers or from noise.

In order to partially eliminate noise, regions with very small areas are discarded. In order to detect the fibers, even those inside large clumps, and then measure their length, the geometrical properties of the remaining regions have to be analyzed. These geometrical properties are captured by first thinning the regions. The thinning algorithm used is Tsuruoka's sequential algorithm for 4-connected thinning of binary images[4]. The advantage of thinning lies in the fact that it reduces the amount of data to be processed at later stages, while preserving its integrity as far as the number of fibers, as well as their individual length and orientation are concerned. In addition, thinned regions allow us to formulate hypotheses as to the placement of the fibers. The observation underlying this statement is that thinning imposes a structure on the local arrangement of pixels inside a region. Since our objective is to detect glass fibers, which are straight line segments, knowledge of the endpoints of a glass fiber is sufficient to determine its position, length and orientation. Thus, in order to produce hypotheses as to the placement of the fibers, we will, in the first step, detect those pixels which could be located either at the extremities of fibers, or at their intersection. Knowledge of these endpoints and intersection points will allow us to predict all possible configurations of fiber arrangements, and then retain only the one configuration that best matches the fiber arrangement in the original image.

In order for a pixel **P** to be an end-point, it should possess the following neighborhood configurations (modulo reflections and rotations):

$$\begin{pmatrix} & & \\ & \mathbf{P} & \times \\ & \times & \times \end{pmatrix} \quad or \quad \begin{pmatrix} & & \\ & \mathbf{P} & \\ & \times & \end{pmatrix}$$

where the $\times$ sign denotes a pixel which belongs to the same region as **P**. Also, in order for **P** to be an intersection-point, it has to have the following neighborhood configurations (modulo reflections and rotations):

$$\begin{pmatrix} & \times & \\ & \mathbf{P} & \times \\ & \times & \end{pmatrix} \quad or \quad \begin{pmatrix} & \times & \\ \times & \mathbf{P} & \times \\ & \times & \end{pmatrix} \quad or \quad \begin{pmatrix} & \times & \\ & \mathbf{P} & \times \\ \times & \times & \end{pmatrix}$$

Let **R** be a thinned region, defined by its constituent image points. Let **S** be the set of end-points and intersection-points which also belong to **R**. Since a line segment can be defined by a pair of points, generating hypotheses as to the presence of fibers is equivalent to mapping the set **S** into the Cartesian product of **S** with itself, i.e. $\mathbf{S} \times \mathbf{S}$. The elements of $\mathbf{S} \times \mathbf{S}$ are of the form (P_i, P_j) where P_i and P_j are end points or intersection points of the thinned region **R**. Note that the hypotheses which correspond to lines of length 0, i.e. to elements in $\mathbf{S} \times \mathbf{S}$ which are of the form (P_i, P_i) are not considered. In addition, since for our purposes the line segment predicted by the pair (P_i, P_j) is the same as that predicted by the pair (P_j, P_i), we need consider, in total, less than half of the generated hypotheses. Once all possible hypotheses have been generated, they are forwarded to the hypothesis verifier which retains only those which could be valid, according to some specific criterion.

3. Verifying Hypotheses

Let **H** be the set of all possible hypotheses. We have $\mathbf{H} = \mathbf{S} \times \mathbf{S}$. As was mentioned previously, we consider only those hypotheses which are neither trivial nor redundant. Evaluating the generated hypotheses and retaining those which could be true is equivalent to mapping **H** into a set of hypotheses $\mathbf{H}^*$, where all hypotheses belonging to $\mathbf{H}^*$ are true according to some specific criterion. This criterion should reflect as much as possible the information in the original image $\mathbf{I}(x, y)$, as far as the placement of fibers is concerned. One criterion that could be used in evaluating the hypothesis $\mathbf{h} \in \mathbf{H}$ would be the following:

$$\mathbf{F}(\mathbf{h}) = \frac{1}{N} \sum_{(x,y) \in L_h} \mathbf{I}(x, y)$$

where L_h is the line segment predicted by hypothesis **h**, and N is the number of points on L_h. Since the background has, on average, a lower intensity than the regions occupied by the glass fibers, the above summation will yield a large value whenever the line segment corresponding to a predicted hypothesis actually overlaps with a glass fiber, and a much lower value if the overlap is only partial. We would thus have:

$$\mathbf{h} \in \mathbf{H}^* \equiv \mathbf{F}(\mathbf{h}) \geq T$$

where T is some threshold imposed on the criterion **F**. This approach could work in some limited cases. Its main drawback is its dependence on variations in the input image, namely, in the gray levels of the regions occupied by the fibers. In order to remove this dependence, the predicted hypotheses are not matched with the original

image, but rather with the thresholded image, where regions occupied by the fibers have a gray value of 1 and the background has a gray value of 0. Thus, our new criterion is:

$$\mathbf{F(h)} = \frac{1}{N} \sum_{(x,y)\in L_h} \mathbf{I_T}(x,y)$$

where $\mathbf{I_T}$ is the thresholded image. The above criterion returns a value between 0 and 1. Again, as before, a hypothesis **h** is retained if and only if its associated criterion function **F(h)** has a value larger than a threshold T. The closer this threshold is to 1, the better the match is between the selected hypotheses and the actual fiber placements. Conversely, if T is chosen close to 0, a large number of false hypotheses will be retained. Thus, the tradeoff on the choice of the threshold T is the traditional tradeoff between accepting a hypothesis given that it is false, and rejecting it given that it is true. Note that we could also have matched the predicted hypotheses with the thinned image. The main drawback with doing this however, lies in the lack of robustness of such a match to slight variations in line orientations (which would be due to noise), bearing in mind that line segments in the thinned image are only one pixel wide, while they are many pixels wide in the thresholded image.

Once the set $\mathbf{H}^*$ of fiber locations has been generated, additional processing is necessary in order to group together line segments which have been artificially split. This happens whenever many fibers intersect. The result of each intersection is an additional intersection point, which yields additional hypotheses. In order to overcome this, the cosine of the angle between any two segments in $\mathbf{H}^*$ which have a point in common is used as a similarity measure to decide whether or not two segments are actually part of a common larger segment.

A summary of the complete algorithm is shown in figure 1.

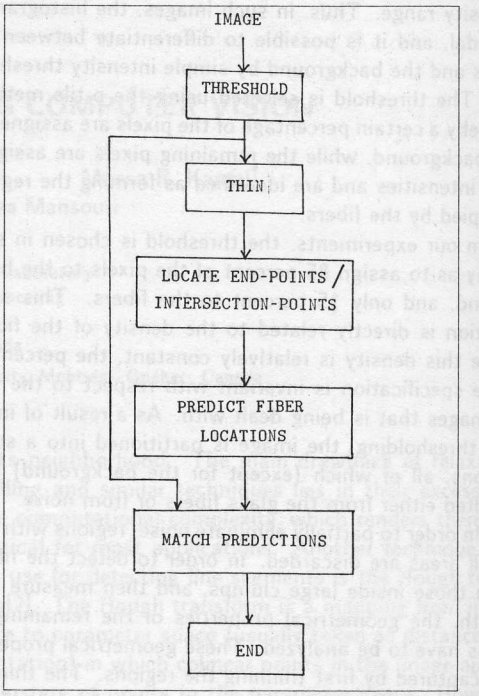

Figure 1 Flowchart of the algorithm

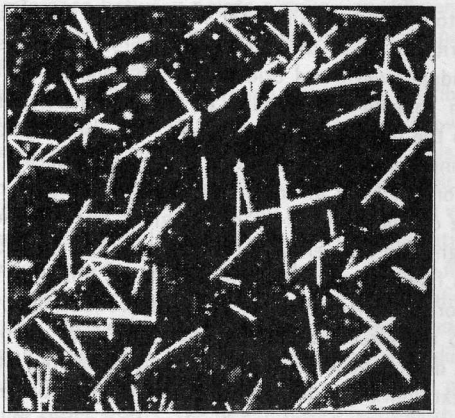

Figure 2 Original image

4. Experimental Results

Figure 2 shows a typical image of a slide of glass fibers. Note that in addition to the clusters formed by the fibers, complications are introduced by the presence of background noise.

Figure 3 shows $\mathbf{I_T}$, the result of thresholding.

Most of the fibers have been retained during the thresholding process owing to their large gray value; however a number of regions which are due to noise have also been detected. These regions, which are small in size, are eliminated in a subsequent step. Figure 4 shows the result of thinning the thresholded image.

As expected, the amount of data to be processed has been drastically reduced, while shape information relating to fiber placement and orientation has been preserved. Figure 5 shows the result of the hypothesis prediction/verification process. The threshold T that is imposed on the quality of the match is chosen equal to 0.9. In other words, in order for a certain hypothesis to be verified as true, there must be at least 90 percent of the line predicted by that same hypothesis overlapping the region occupied by the fibers in the thresholded image.

As can be seen, a large number (80 to 85 percent) of fibers have been properly detected. Fibers which have not been detected are mostly associated with noisy regions, or even regions with an excessively large number of intersecting fibers, where the shape information is distorted during the thinning process. This distortion results in a series of false hypotheses, which are then rejected during the hypothesis verification process, owing to the rather poor match between the predicted line segments and the actual fiber locations. Conversely, a number of line seg-

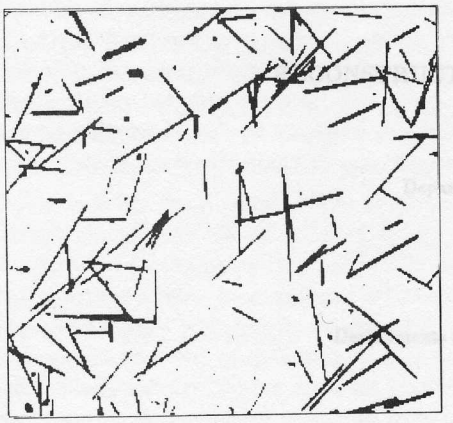

Figure 3 Thresholded image

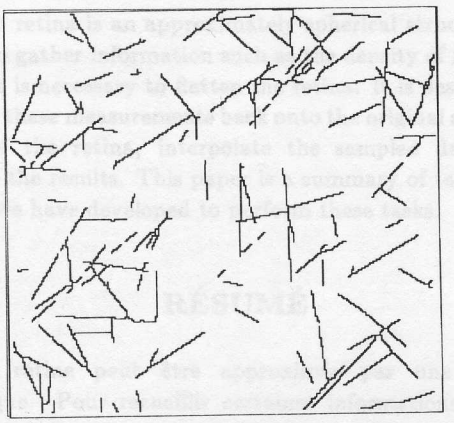

Figure 4 Thinned image

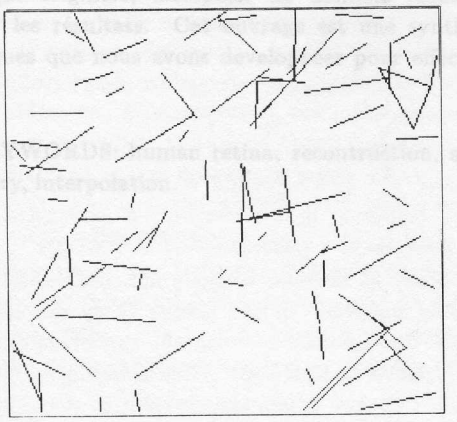

Figure 5 Detected fibers

ments which are not associated with fibers but rather with noisy clusters are also detected.

Figure 6 shows another image of glass fibers.

In figure 7, the result of the processing is illustrated. Again, the same proportion of glass fibers has been detected.

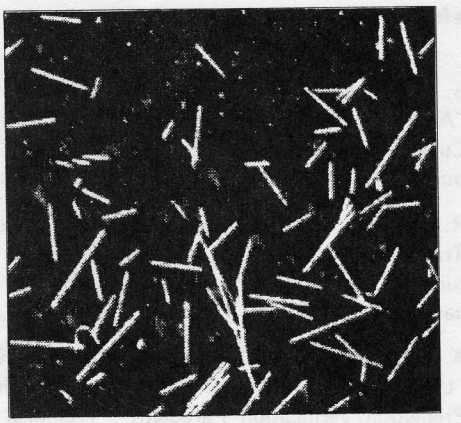

Figure 6 Original image

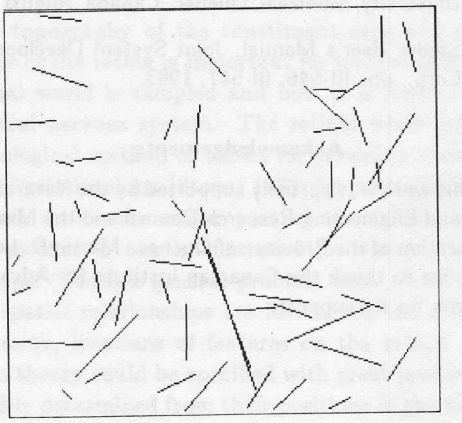

Figure 7 Detected fibers

The procedure currently distinguishes between 75 and 95 percent of fibers present in a given image, depending mainly on its complexity. The use of computer vision for glass fiber detection has yielded a dramatic increase in the analysis speed of the cross-sectioned experimental samples.

5. Conclusion

In this paper, an efficient algorithm for detecting glass fibers in polymer products was presented. The algorithm consisted of two parts: a predictor, which predicted possible fiber locations, and a verifier, which matched the predictions against information derived from the original image. The results of applying this algorithm to images of glass fibers were shown to be successful.

Although this algorithm currently distinguishes between 75 and 95 percent of the fibers present in a given image, possible improvements can be made by refining the segmentation procedure used in this algorithm; this will lead to a better discrimination between fibers and background, and will hence reduce the loss of shape information.

6. References

1. S. W. Zucker, R. A. Hummel and A. Rosenfeld, "An Application of Relaxation Labelling to Line and Curve Enhancement," *IEEE Transactions on Computers*, Vol. 26, No. 4, April 1977, pp. 394-403.

2. R. O. Duda and P. E. Hart, "Use of the Hough Transformation to Detect Lines and Curves in Pictures," *Communications of the ACM*, Vol. 15, January 1972, pp. 11-15.

3. A.-R. Mansouri, A. S. Malowany and M. D. Levine, "Line Detection in Digital Pictures: A Hypothesis Prediction/Verification Paradigm," Technical Report 85-17R, Computer Vision and Robotics Laboratory, Department of Electrical Engineering, McGill University, Montreal, Quebec, Canada, August 1985.

4. Spider User's Manual, Joint System Development Corp., pp. III-546, III-547, 1983.

Acknowledgements

This work was partially supported by the Natural Sciences and Engineering Research Council and the Ministry of Education of the Province of Québec. Martin D. Levine would like to thank the Canadian Institute for Advanced Research for its support.