A Multiple-Scale Measure of Static Tactile Texture

R. E. Ellis

COINS Technical Report 85-18

Laboratory for Perceptual Robotics
Department of Computer and Information Science
University of Massachusetts, Amherst MA 01003

ABSTRACT

A method is proposed for extracting features useful in the classification of tactile texture. It is shown that accumulation of simple statistics on step edges and local extrema, calculated at multiple scales of resolution, can provide linearly separable features. The scale can be varied continuously, thus facilitating adjustment of the scale parameters to optimise the effectiveness of the method for a given tactile array sensor.

Publication as a COINS technical report was sponsored by Prof. Allen Hanson.

§ 1.0 Introduction

Characterisation of tactile texture is an important part of the object recognition process in humans [Lederman, 1982], and can be performed by a robot that is equipped with a tactile array sensor. One approach is to analyse a static imprint and find some simple measures of texture [Ellis, 1984]. To be successful, this approach requires a sensor of high spatial resolution, a dynamic range of at least several bits, and good repeatability. A number of such sensors are now available, e.g., [Luo *et al.*, 1984], [Mott *et al.*, 1984], and [Tanie *et al.*, 1984]; the one chosen for this work was developed in our Robotics Laboratory, and is described in [Begej, 1985].

The problem is to find features that can be used by a simple learning algorithm to classify texture into a number of distinct classes, when examples from these classes are available. The learning rule chosen is a simple one, which attempts to classify by the principle of linear separability; what must be found are features that group measurements (samples) into linearly separable classes – the problem of taking some arbitrary set of features and finding a classification scheme for them is not a topic of concern here.

The tactile domain has some characteristics which can be used to advantage. One is that for most tactile sensors, the response of a tactile element, or taxel, is monotonically related to the strain in a sensing medium, which under simple conditions implies a monotonic relation between normal pressure and taxel intensity. This means that local maxima and minima in tactile imprints have direct, and perhaps important, physical meaning, which is not always the case in visual image analysis.

Another useful characteristic is that step discontinuities in the tactile imprint occur as a result of variations in the surface of an object which is pressed normally against the sensor. Step-finding is, however, complicated by strain propogation in the sensing medium that spreads the step response over the array [Fearing and Hollerbach, 1984]. If a step-finding algorithm is to be used, it must be one which provides both good detection and good localisation.

A final consideration is that for a given tactile array sensor, there is a known, fixed distance between any two taxels. For typical tactile array sensors, which are both *stiff* and *planar*, there is a simple relationship between the inter-taxel distance and the distance of texture elements on the touched object. Thus, it may be possible to parameterise a texture classification procedure so that it is applicable to several sensor array geometries.

The method proposed below uses all three of these characteristics to advantage. The features chosen are statistics on the step changes and local extrema, as found after smoothing the data with a two-dimensional Gaussian. This procedure provides excellent detection and localisation of the step changes, provides a scale over which local extrema are defined, and can be adapted from one sensor to another by a simple change in the scale factor.

The succeeding sections discuss the theory behind this method, present the experimental results that indicate the method's utility, and examine the results, indicating possible future work.

§ 2.0 Motivation

Almost all work on computer perception of texture deals with visual images. The approaches include examination of histograms of local features such as step magnitude and orientation [Zucker *et al.*, 1975], applying spatial constraints to co-occurrence of local features [Davis *et al.*, 1979], or measuring the mean absolute value (the energy) of convolutions with simple masks [Laws, 1979].

A difficulty that arises in applying these methods to the tactile domain is that tactile imprints – even those of highly textured regions – exhibit a lower density of step discontinuities, and poorer step localisation, than do typical images of visual texture. This criticism does not apply to the latter method employing texture energy, which instead fails when there are several classes of coarse texture and several of fine texture (failure is due to fixed mask size).

An approach to tactile texture by dynamic sensing was preliminarily explored in [Bajcsy, 1984]. This involved moving a tactile sensor against various surfaces and noting gross differences in the resulting time-varying signal from a few taxels. Although this is an intriguing approach, the work was severely limited by the low spatial resolution of the sensor, and by the taxels' providing only one bit of information. (Static texture imprints are also presented in this paper, but merely to note that the sensor was capable of detecting them; no analysis of the texture was reported.)

Another texture method specifically addressing tactile issues was proposed by [Ellis, 1984], in which crude estimates of static texture were extracted to indicate that very different classes can, in fact, be distinguished. The method relied upon knowledge of the applied normal force, and classified on the basis of simple statistics of both local maxima, and of the mean intensity difference between a maximum and its neighbours. Subsequent investigation, however, showed that for some classes that were distinct to human touch (such as light contact of knitted fabric, and light contact with a planar surface), the measures were inadequate.

The question that arose was what additions, modifications, and alternatives to the above static methods might work. One possible source of inspiration is the psychophysical literature. For example, the psychological results in [Lederman, 1982] indicate that the intensity of step changes *and* the amount of unstrained skin were important in distinguishing amongst various linear metal gratings. These results suggest that a combination of step and local extrema information might produce a feature vector of sufficient richness for classification in the robotics domain.

Another consideration was that the classification desired was not of texture in the *imprint*, but rather texture of the *object*; it was thus necessary to explicitly take into account the geometry of the sensor (rather than its topology). If the algorithm was to be useful on distinct tactile sensor technologies, a method of accounting for the geometric differences was necessary.

The method selected was to extract both step discontinuities and local maxima at a given scale. The original data were smoothed with a Gaussian of a given standard deviation (hereafter denoted σ) which defined the scale of observation; steps and extrema could both be extracted from this smoothed array. The step-edge operator of [Canny, 1984], employs Gaussian smoothing in its first stage, and is thus compatible with this extraction procedure. Gaussian convolution kernels have the advantages of being symmetric, strictly decreasing about the mean, limiting to the signal at small σ and to the mean at large σ, and can be smoothly varied [Witkin, 1983]. Because the geometry and transduction mechanism of a tactile sensor is fixed, and the sensor is stiff and planar, the relation between the texture of the imprint and the texture of the sensed object are related by means of a scale factor which can be easily and regularly varied from sensor to sensor (increasing the ease with which the algorithm can be applied to various sensor sensitivities and geometries).

§ 3.0 Experimental Results

A series of experiments were conducted with tactile imprints of various textured objects. This section describes the tactile imprints used in the experiments, the learning algorithm, the features and how they were extracted, and the results of the experiments.

§ 3.1 Tactile Imprints for the Experiments

The sensor used in these experiments was a very high density tactile array, capable of sensing 128×128 taxels in an area 36 mm^2, with approximately 7 bits/taxel (see [Begej, 1984] for a technical description). Because of the nature of the calibration procedure, only about 32×32 of this array was available at once; texture was thus examined over an area about 9 mm^2.

The data were grouped into four classes, with six members per class. There were three objects used for each class, and each was pressed lightly (about 20 N. normal force) and heavily (about 50 N. force). The classes and object were:

- **Planar**: PLANE1 – stiff conductive foam; PLANE2 – unetched printed circuit board; PLANE3 – a rubber sheet with pressure applied pneumatically to ensure hydrostatic equalisation across the array (this is a tool used in calibrating the array).
- **Rasp**: Rasp1 – wood rasp with teeth arrayed in parallelograms measuring 3.4 mm by 4.0 mm tooth-to-tooth; RASP2 – a different rasp similarly arrayed; RASP3 – a rasp with teeth arrayed 2.4 mm by 3.2 mm tooth-to-tooth.
- **Screen**: SCREEN1 – aluminum screen with holes in a square grid 6.8 mm centre-to-centre; SCREEN2 – aluminum screen with holes in a hexagonal grid 4.9 mm centre-to-centre; SCREEN3 – a hard, thick plastic block with holes in a hexagonal grid 5.2 mm centre-to-centre.
- **Knit**: KNIT1 – a rib-knit sweater cuff, pure virgin wool; KNIT2 – a stocking-knit sweater, pure acrylic; KNIT3 – a garter-knit scarf, composition unknown but suspected to be acrylic.

See Figures 1 through 24 for examples of the imprints taken by the tactile sensor.

These data presented considerable challenge, due to several inter-class similarities. When the imprints were displayed on a graphics device, the scarf and rasps appear similar under light pressure, but the knits were distorted into near-planar contact by the heavy load – particularly the pure wool sweater. The applied force was not available to the classification scheme; the classification difficulty increases significantly when both light and heavy forces are present in the same class. (Due to hysteresis in the sensor, normal force can be deduced from the tactile imprint to an accuracy of about one part in four; the forces used could not be reliably resolved because of this and because some taxels were saturated, which makes the force inference unreliable in any case.)

§ 3.2 The Learning Algorithm

A very simple, perceptron-like rule for distinguishing linearly separable classes was employed. Although a better rule could no doubt be employed, this rule has the advantages of simplicity and of providing a very good test of the separability of the features: because it is so simple, the features must be very good ones for the rule to converge on a solution and to be reliable.

The rule employed was Kesler's construction of the multi-class case of the perceptron learning rule [Duda and Hart, 1975]. In this scheme, the inner product between the input vector $\mathbf{b}$ and each weight vector $\mathbf{a}_i$ is taken; the input vector is classified by assigning to it the class of the weight vector which produces the largest value.

If $\mathbf{b}$ belongs in class i, but was mistakenly labelled as belonging to class j, then only two weight vectors were modified according to the formulae:

$$\mathbf{a}_i \leftarrow \mathbf{a}_i + c \cdot \mathbf{b}$$

$$\mathbf{a}_j \leftarrow \mathbf{a}_j - c \cdot \mathbf{b}$$

where c is some learning constant. It can be shown that if the data are linearly separable, this algorithm will eventually converge to a correct set of weight vectors

from an arbitrary set of weight vectors for any positive c, if the data are repeatedly presented to the algorithm.

For the purposes of these experiments, all a vectors were initially zero, and c was always unity. Convergence generally occurred in fewer than 100 iterations over the data; if 1000 iterations were performed without perfect classification, convergence was said to have failed.

§ 3.3 The Feature Vectors

Once the basic feature types were selected, algorithms for their extraction could be chosen. Of the alternatives available, the preferred step edge operator was the one proposed in [Canny, 1983]. It has the desirable properties of being nearly optimal in both detecting and localising step changes in intensity, and of providing a way of estimating the noise present in the imprint.

The operator essentially finds local maxima of the gradient in the gradient direction, after smoothing by a Gaussian. The window for the Gaussian should be at least three times the standard deviation desired, to ensure proper results; in this implementation, a factor of four times the standard deviation was used. The standard deviation is in units of taxel spacing, i.e., adjacent taxels are distance 1 apart; thus, for the tactile array sensor employed in these experiments, a smoothing scale of 1 corresponds to a real-space Gaussian with a standard deviation of approximately 0.28 mm. The Canny operator used here performs a global estimation of the noise in the image, based on the RMS value of the third derivatives after smoothing; the two threshold values used in the non-maximum-suppression stage of the operator were 1.2 and 1.8 times the noise estimate. (See the above reference for a complete description of the meaning of these values.)

The statistics calculated for step changes were:

- the number present in the imprint;
- the mean strength of the step changes (given by the magnitude of the gradient in the gradient direction, after smoothing);

- the standard deviation of the magnitude;
- the standard deviation of the orientations.

Local extrema were calculated on the basis of a 3×3 window for the square-grid sensor arrays used. The statistics gathered for minima and maxima separately were:

- the number of minima/maxima in the imprint;
- the mean difference between a minimum/maximum and the average of its surrounding eight neighbours;
- the standard deviation of this value.

The extrema were calculated at the same scale of smoothing as the step changes; the presumption is that by so doing the features that are extracted have some close correspondence, and thus are capturing physical information about the object's texture at the given scale.

The area over which the features were extracted was the full 32×32 image. This choice was necessitated by the small physical size of the calibrated section of the sensor, and because some of the textures (particularly the coarse screen) and the entire image were of comparable size. Had finer textures or a different sensor been chosen, a smaller window could have been selected.

Examples of step edges and local extrema are presented in Figures 25 through 28, at different scales of resolution. All of these are extracted from the screen imprint shown in Figure 18. The step edge images show the approximate orientation, as well as the location, of the step discontinuities found by the operator. Local maxima are represented as large dark squares, and local minima as small ones.

§ 3.4 Classification Experiments

A variety of experiments were performed using the above data. The methodology was to remove one or more feature vectors from the data list, attempt to converge to a perfect classification of the remaining vectors, and then classify the test set. The variables in the experiment were the scales of smoothing at which feature extraction was performed.

The experiments conducted were: removal of one vector from the original data set; removal of both imprints of the same object from the original set; removal of one imprint from each class; and removal of both imprints of one object from each class. The successful experiments will be described first, followed by a description of the experiments in which convergence to perfect classification did not occur.

§ 3.4.1 Experiment 1: 1/N

This experiment involved removing 1 of the original 24 feature vectors, learning to classify the remaining 23, and then testing the classifier with the removed vector. All 24 original vectors were subjected to this test, at varying scales of resolution, and with varying richness of feature vectors. This test is probably the most important, for it demonstrates the performance of an algorithm in the presence of the largest amount of training data.

Two forms of test were conducted: one with the full feature vectors and one with some elements removed. Since the weight vectors are found by adaptation, weight elements which are low reflect relatively unimportant elements of the feature vectors; smaller feature vectors, if they can be used, are computationally more efficient. The elements removed in the short version of the feature vectors were the step orientation statistics and the local extrema mean height differences, which had a much lower weight than the other features.

In order to capture all of the textural information contained in the imprints, it seems necessary to extract features at a minimum of *two* scales of resolution. Where '∞' represents failure to converge, the correct classification of the removed test vectors is as follows:

Smoothing Scale(s)	10 features/scale	7 features/scale
0 (no smoothing)	∞	∞
1	∞	∞
2	∞	∞
3	∞	∞
1 & 2	∞	∞
1 & 3	18 / 24 (75%)	19 / 24 (79%)
1, 2, & 4	22 / 24 (92%)	20 / 24 (83%)

The degradation evident due to a smaller number of features is small but expected; for some applications, the decrease in computation time may make such degradation acceptable. Use of three scales of smoothing was better than use of two scales in this experiment, largely because of misclassifications of rasps *vs.* knits – some members of the classes are sufficiently similar that considerable confusion resulted. The increase in performance in the two-scale case, going from 10 to 7 features, was caused by removal of a spurious correlation in orientation statistics. As will be shown below, removing more members from the original set reduces the chance of confusion, and thus changes the relative effectiveness of two and three scales of resolution.

§ 3.4.2 Experiment 2: 2/N

For this experiment, both of the imprints for a given object (with light and heavy imprinting forces) were removed, the classifier constructed, and independent classification of the removed vectors was performed. This was done for all 12 objects; thus, all 24 imprints were eventually tested.

This experiment tested the system's ability to generalise to an entirely new object in each class. As is to be expected, the results are slightly poorer than those of the previous experiment, but not exceptionally so:

Smoothing Scale(s)	10 features/scale	7 features/scale
0 (no smoothing)	∞	∞
1	∞	∞
2	∞	∞
3	∞	∞
1 & 2	∞	∞
1 & 3	17 / 24 (71%)	16 / 24 (67%)
1, 2, & 4	18 / 24 (75%)	18 / 24 (76%)

Given the small number of objects available, and the nature of the test data, this degradation in classification correctness is of an appropriate magnitude. The two-scale methods degraded less, both absolutely and relatively, than did the three-scale methods.

§ 3.4.3 Experiment 3: 1/CLASS

This experiment involved removing one feature vector from each class, learning the classes of the remainder, and independently classifying the test set. Each learning phase was thus on 20 feature vectors, with 4 subsequently tested. All 24 were, of course, eventually removed for testing.

It was becoming apparent, when running the experiments, that single-scale testing was a waste of time. The abbreviated test results are:

Smoothing Scale(s)	10 features/scale	7 features/scale
1 & 3	19 / 24 (79%)	20 / 24 (83%)
1, 2, & 4	20 / 24 (83%)	21 / 24 (88%)

Note again the increased performance in the two-scale test, due to removal of a spurious correlation.

§ 3.4.4 Experiment 4: 2/CLASS

This last experiment entailed removing both imprints of an object, for one object from each class. The learning was thus done on only 16 feature vectors, with 8 independently tested afterward. The results were:

Smoothing Scale(s)	10 features/scale	7 features/scale
1 & 3	17 / 24 (71%)	18 / 24 (75%)
1, 2, & 4	21 / 24 (88%)	22 / 24 (92%)

The change in the two-scale case is consistent with those of the previous scenarios; the surprising result is that the three-scale case is slightly better than it was for other trials. This appears to be artifactual, and is a result of having simultaneously removed a particular pair (a knit and a rasp) whose feature vectors were readily confused.

§ 3.4.5 Misclassifications

Some of the data were consistently misclassified, while others were misclassified in only one or two scenarios. The persistent misclassifications seem due to the cross-class similarity of some imprints; occasional errors tend to indicate sensitivity of the method to small differences in the input data.

There was a strong tendency to misclassify 3 of the 24 imprints: SCREEN3-SOFT (Figure 17), KNIT2-HARD (Figure 22), and RASP3-SOFT (Figure 11). When these were misclassified, it was without exception as a knit, a screen, and a screen respectively; at least two of these three were misclassed in every scenario.

Visual examination of the imprints gives some idea of the reasons for the errors. The knit imprint has been deformed so much that holes – darker, or less pressured areas – seem to form, leading to a screen-like appearance. The screen imprint fails to have the excellent definition of others in its class; this is artifactual, for when a hard block of stiff plastic is pressed against an unsupported glass plate the glass bends slightly, and fails to contact uniformly. The rasp was pressed so lightly that for large smoothing scales, the high-intensity features blend into each

other, giving "virtual holes" and thus appearing more like a screen.

None of the other misclassifications appear systematically. It is most likely that they occur because such a simple learning algorithm is being used; although the classes are by and large linearly separable, the hyper-planes are often very close to the data vectors. This could probably be addressed by using either a minimum-squared-error procedure, or some non-linear classification scheme. For the present purposes, though, it is sufficient to note that reasonable reliability can be achieved with a very simple classification technique.

§ 3.4.6 Convergence Results

One of the more surprising consequences of the experiments was that for these test imprints, rather rich feature vectors are needed in order to distinguish the classes. It is quite apparent that not only must multiple scales of resolution be used with features of this nature, but the features extracted at each scale must provide a great deal of information.

A number of other tests were conducted, solely to establish conditions under which convergence occurred. Convergence to perfect classification, in the 1/N case described above, failed under the following conditions:

- Long form of the vectors, single scale of resolution.
- Short form of the vectors, single scale of resolution.
- Either form of the vectors, if multiple scales of resolution were used but were too close.
- Step edge statistics only, single or multiple scales of resolution.
- Extremum statistics only, single or multiple scales of resolution.

Also tested on these data were the texture energy measures proposed in [Laws, 1979]. There was no convergence for the 3×3 neighbourhood energy measure, and convergence was erratic for the energy measure with a 5×5 neighbourhood – when it did converge to an answer, it gave approximately the same correctness of classification as the short form of the multiple-scale feature vectors.

§ 4.0 Conclusions

Given the small number of classes and samples in the test set, no firm conclusions can yet be drawn; these results should be considered to be preliminary to an extensive test. However, one may tentatively conclude that multiple-scale features can be reliably used to classify the tactile texture of objects, providing that sufficiently rich feature vectors are available.

As had been expected, a single scale of resolution is inadequate for the texture classes examined. Knits, rasps, and screens all exhibit both a fine structure and a gross one; multiple scales of resolution are capable of extracting these structures where a single one cannot. For the classes tested, three scales of resolution were best in one experiment, and two scales best in all others. This is probably because the scales of the texture were, quite simply, best captured at the scales of $\sigma = 1$ and $\sigma = 3$; that is, the physical texture was largely at these scales, and the scales of the triplet test did not capture this textural information.

One unexpected result was that mean step edge strength and mean extremum height from the surround are relatively unimportant, whereas the standard deviation of these values *are* important. It had been anticipated that the metallic rasps and screens could be distinguished from planar and knitted surfaces by having very strong extrema; what actually was a useful discriminant was the consistency of the edges and maxima, rather than their relative strengths.

Several improvements to the algorithm are possible. One definite weakness is in the learning scheme – the rule employed ceases to improve once correct discrimination occurs; it does not attempt to optimise the discriminator. Other approaches, such as a multi-class minimum squared error scheme, could possibly improve performance.

Another improvement might be to conduct a multi-dimensional search of scale space to find optimum scales of resolution for the data at hand. This could be very costly, but if it was desired to implement this system as a working texture discriminator for a known number of classes, an improvement in running performance

could well be worth the extra training cost. (The scales studied here were selected by hand, relying on the author's familiarity with the results of applying step-edge operators to tactile imprints.)

Finally, the approach is essentially one of static imprint analysis. The data provided were at different imprinting forces, but this was not used to improve the discrimination. Various dynamic tactile texture techniques, including touching an object with varying force to establish the texture as the object deforms, are topics for future investigation.

§ 5.0 Acknowledgements

The experiments were conducted with the tactile sensor designed and constructed by Stefan Begej, who deserves much thanks for his assistance. Allen Hanson and Les Kitchen made valuable comments on this paper.

This research was supported in part by the Office of Naval Research under contract N00014-84-K-0564.

§ 6.0 References

Bajcsy, R. 1984: "What can we learn from one finger experiments?", in *Robotics Research* (M. Brady & R. P. Paul, eds.), MIT Press: Cambridge, pp. 509-527.

Begej, S. 1984: "An optical tactile array sensor", *Proceedings of the SPIE Conference on Intelligent Robots and Computer Vision*, pp. 271-280.

Begej, S. 1985: "A tactile sensing system and optical tactile-sensor array for robotic applications", Technical Report 85-06, Department of Computer and Information Science, University of Massachusetts.

Canny, J.F. 1983: *Finding Edges and Lines in Images*. Technical Report No. 720, Artificial Intelligence Laboratory, Massachusetts Institute of Technology.

Davis, L.S, Johns, S.A., and Aggarwal, J.K. 1979: "Texture analysis using generalized co-occurrence matrices", *IEEE Transactions on Pattern Recognition and Machine Intelligence* 1(3)251-259.

Duda, R.O. and Hart, P.E. 1973: *Pattern Classification and Scene Analysis*, John Wiley & Sons: NY, pp. 176-177.

Ellis, R. 1984: "Extraction of tactile features by passive and active sensing", *Proceedings of the SPIE Conference on Intelligent Robots and Computer Vision*, pp. 289-295.

Fearing, R.S. and Hollerbach, J. 1984: "Basic solid mechanics for tactile sensing", *Proceedings of the IEEE Conference on Robotics*, pp. 266-275.

Laws, K.I. 1979: "Texture energy measures", *Proceedings of the DARPA Image Understanding Workshop*, pp. 47-51.

Lederman, S.J. 1982: "The perception of texture by touch", in *Tactual Perception*, W. Schiff & E. Foulke eds., Cambridge University Press: Cambridge U.K., pp. 130-167.

Luo, R-C., Wang, F-L., and Fernandez, C. 1984: "Imaging tactile sensor with magnetostrictive transduction", *Proceedings of the SPIE Conference on Intelligent Robots and Computer Vision*, in press.

Mott, D.H., Lee, M.H., and Nicholis, H.R. 1984: "An experimental very high resolution tactile sensor array", *Proceedings of the Robot Vision and Sensory Controls Conference*, IFS (Publications): London, pp. 241-250.

Tanie, K., Komoriya, K., Kaneko, M., and Tachi, S. 1984: "A high resolution tactile sensor", *Proceedings of the Robot Vision and Sensory Controls Conference*, IFS (Publications): London, pp. 251-260.

Witkin, A.P. 1983: "Scale-space filtering", *Proceedings IJCAI-7*, pp. 1019-1022.

Zucker, S., Rosenfeld, A., and Davis, L. 1975: "Picture segmentation by texture discrimination", *IEEE Transactions on Computers*, vol. C-24, pp. 1228-1233.

§ 7.0 Figures

Figure 1: PLANE1-SOFT – Light planar contact (stiff foam)

Figure 2: PLANE1-HARD – Heavy planar contact (stiff foam)

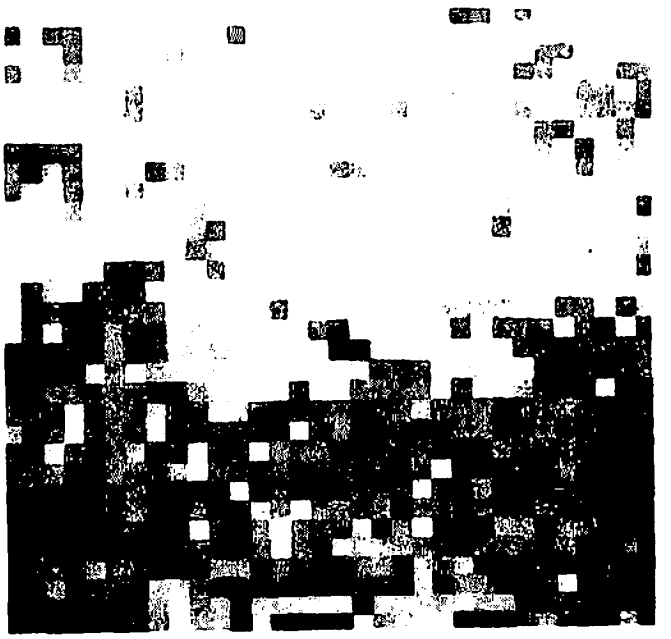

Figure 3: PLANE2-SOFT – Light planar contact (thick rubber)

Figure 4: PLANE2-HARD – Heavy planar contact (thick rubber)

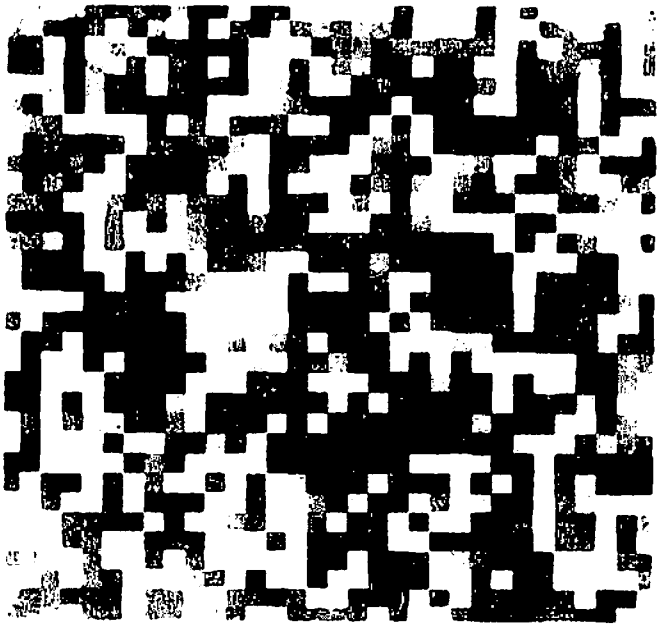

Figure 5: PLANE3-SOFT – Light planar contact (pneumatic pad)

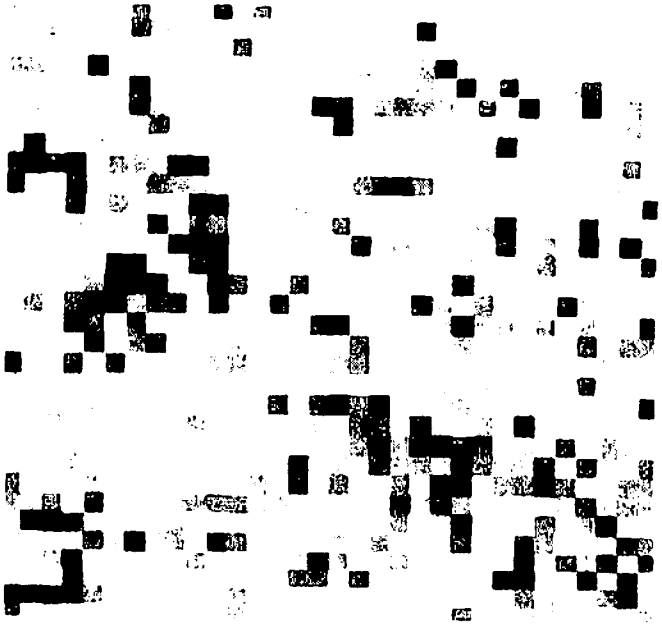

Figure 6: PLANE3-HARD – Heavy planar contact (pneumatic pad)

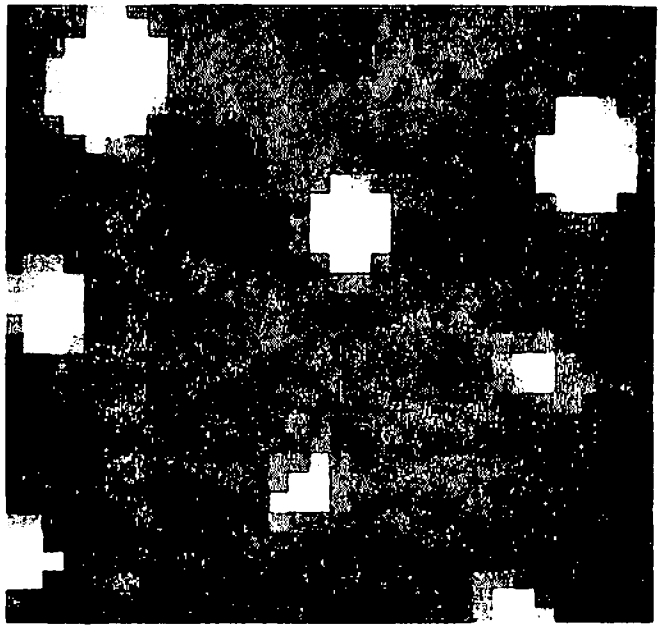

Figure 7: RASP1-SOFT – Light wood rasp contact (coarse teeth)

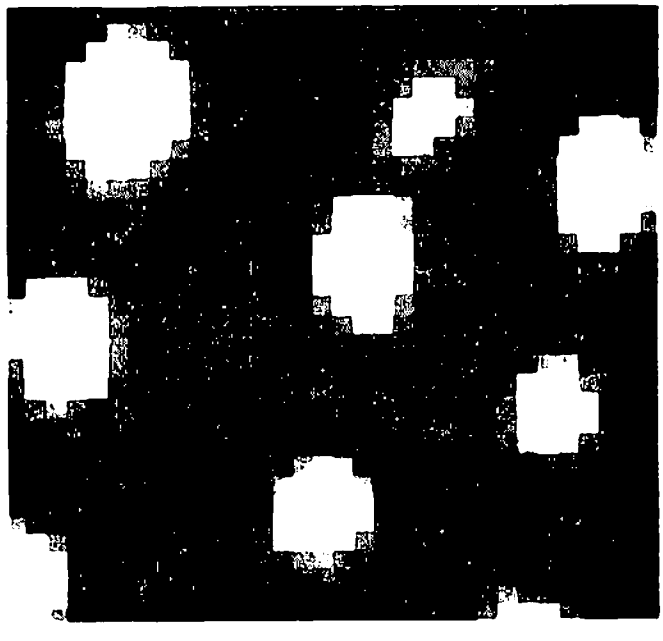

Figure 8: RASP1-HARD – Heavy wood rasp contact (coarse teeth)

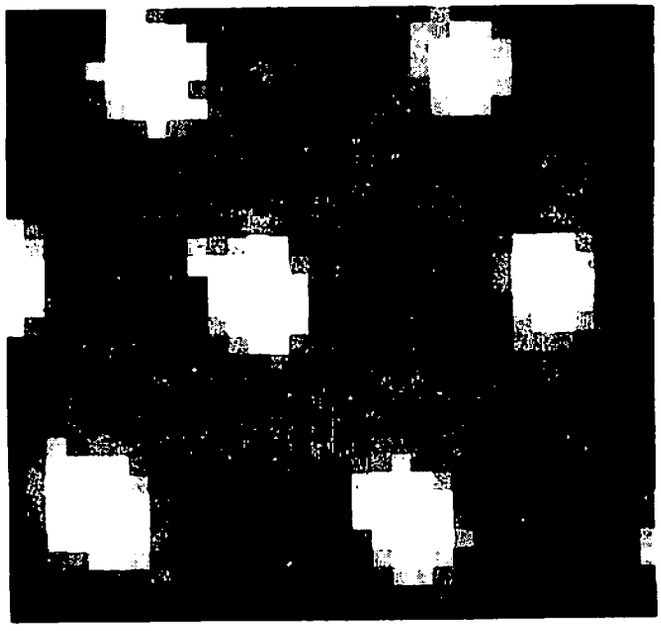

Figure 9: RASP2-SOFT – Light wood rasp contact (coarse teeth)

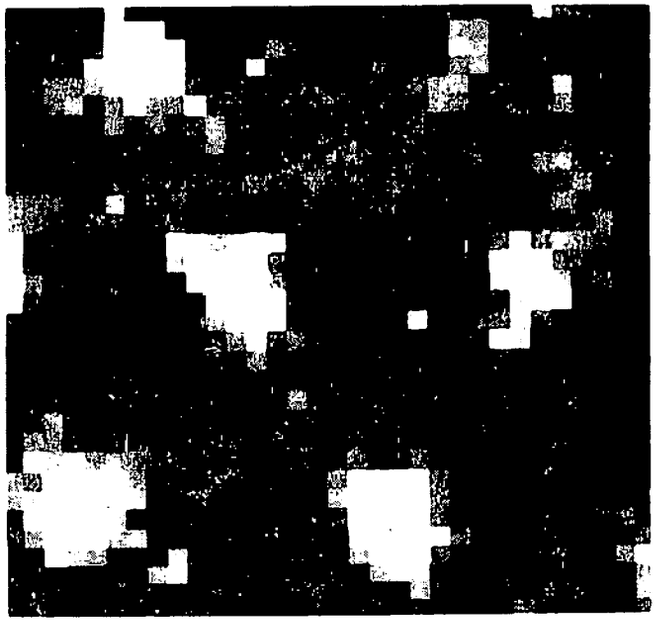

Figure 10: RASP2-HARD – Heavy wood rasp contact (coarse teeth)

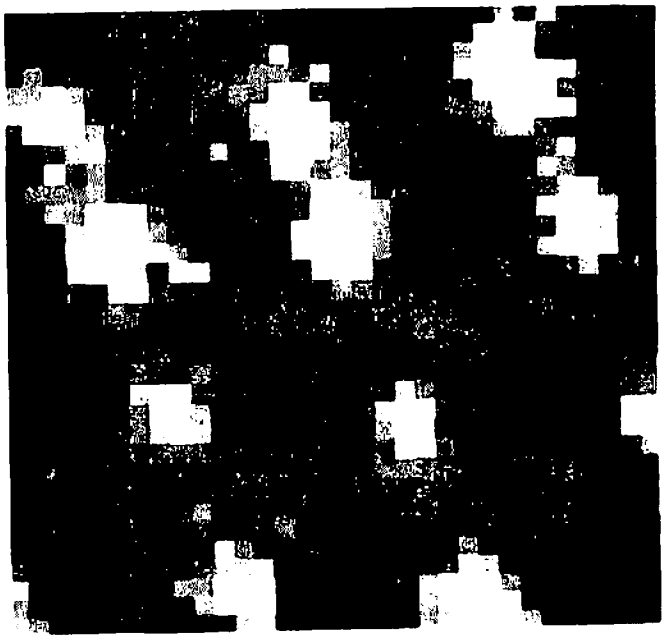

Figure 11: RASP3-SOFT – Light wood rasp contact (medium teeth)

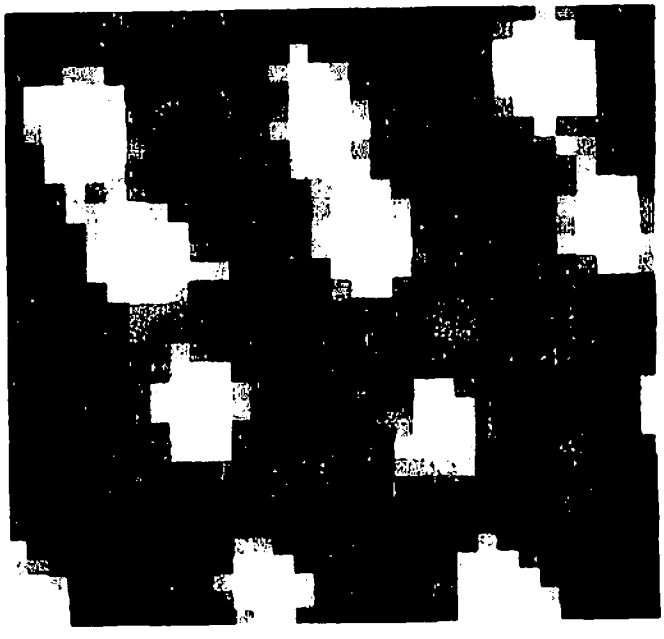

Figure 12: RASP3-HARD – Heavy wood rasp contact (medium teeth)

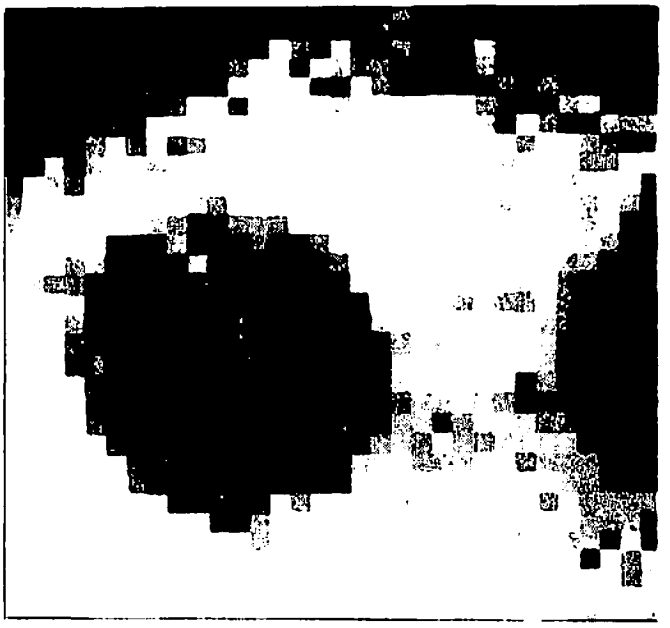

Figure 13: SCREEN1-SOFT – Light screen contact (square grid)

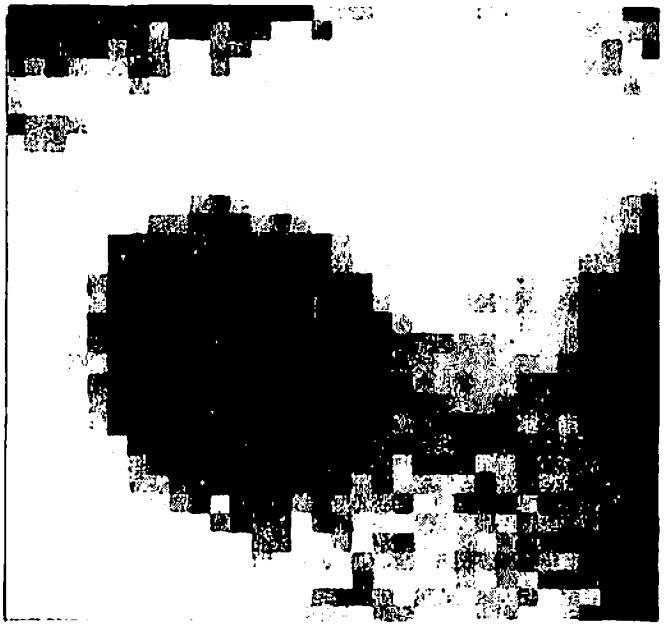

Figure 14: SCREEN1-HARD – Heavy screen contact (square grid)

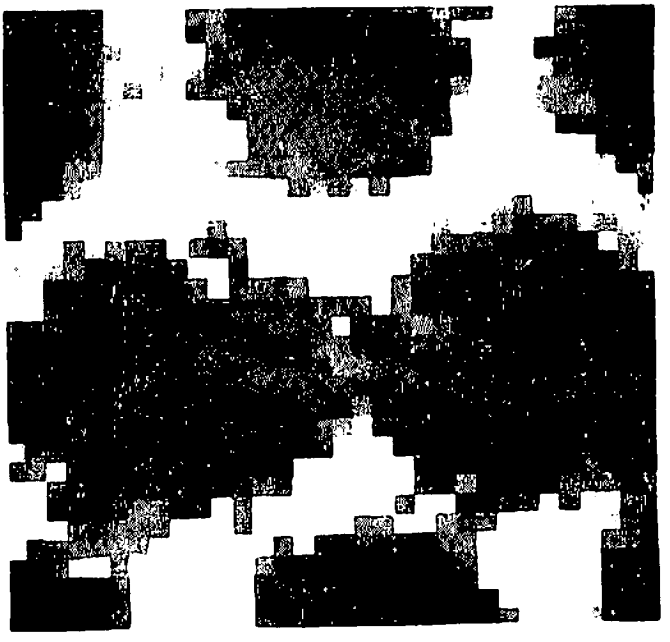

Figure 15: SCREEN2-SOFT – Light screen contact (loose hex grid)

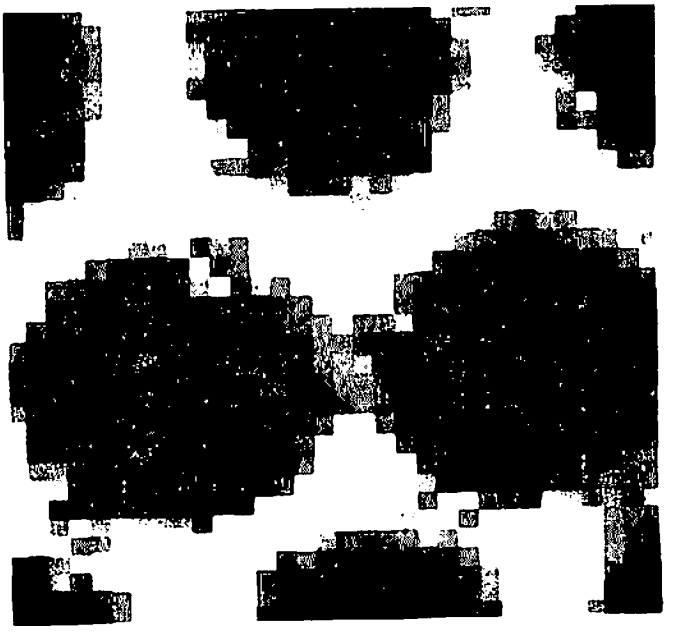

Figure 16: SCREEN2-HARD – Heavy screen contact (loose hex grid)

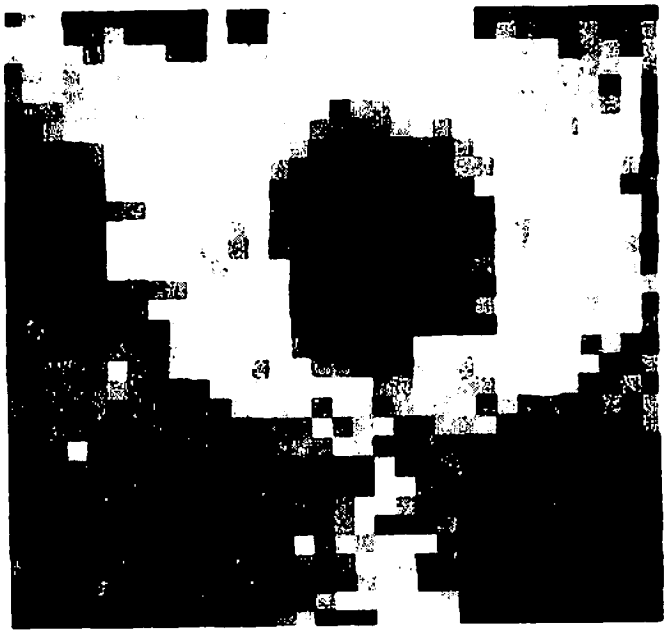

Figure 17: SCREEN3-SOFT – Light screen contact (sparse hex grid)

Figure 18: SCREEN3-HARD – Heavy screen contact (sparse hex grid)

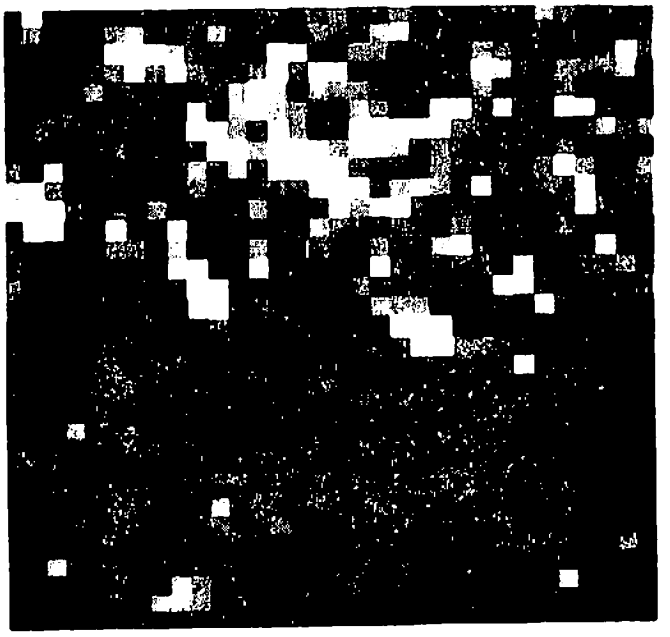

Figure 19: KNIT1-SOFT – Light knit contact (rib-knit wool)

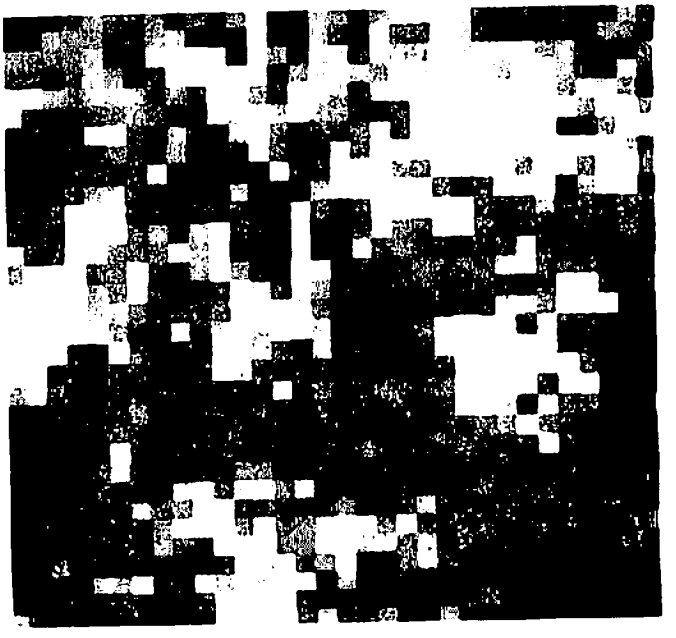

Figure 20: KNIT1-HARD – Heavy knit contact (rib-knit wool)

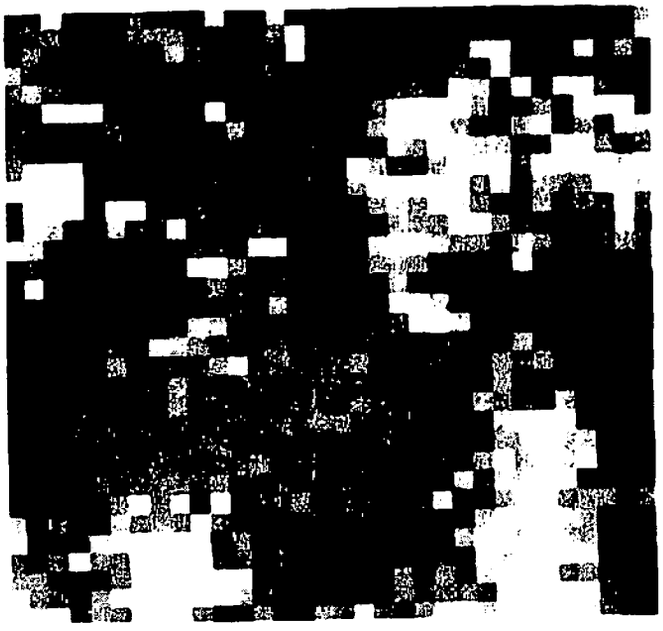

Figure 21: KNIT2-SOFT – Light knit contact (stocking-knit acrylic)

Figure 22: KNIT2-HARD – Heavy knit contact (stocking-knit acrylic)

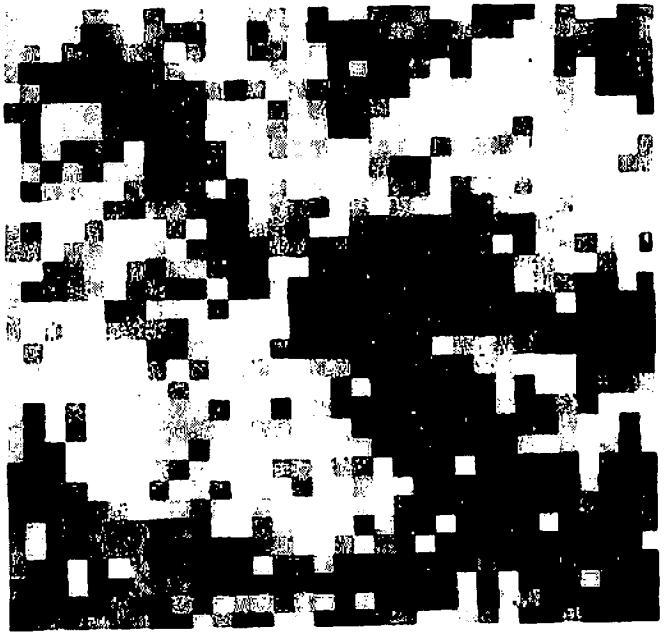

Figure 23: KNIT3-SOFT – Light knit contact (garter-knit acrylic)

Figure 24: KNIT3-HARD – Heavy knit contact (garter-knit acrylic)

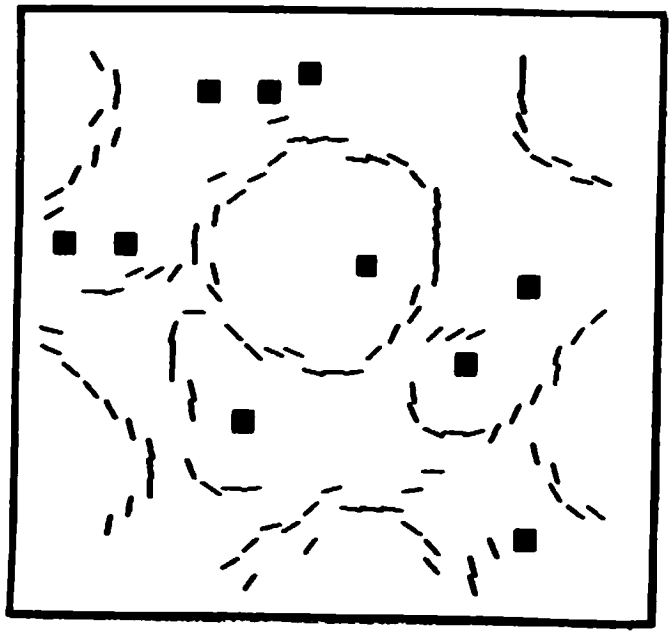

Figure 25: Edges & local extrema found from Figure 18, $\sigma = 1$

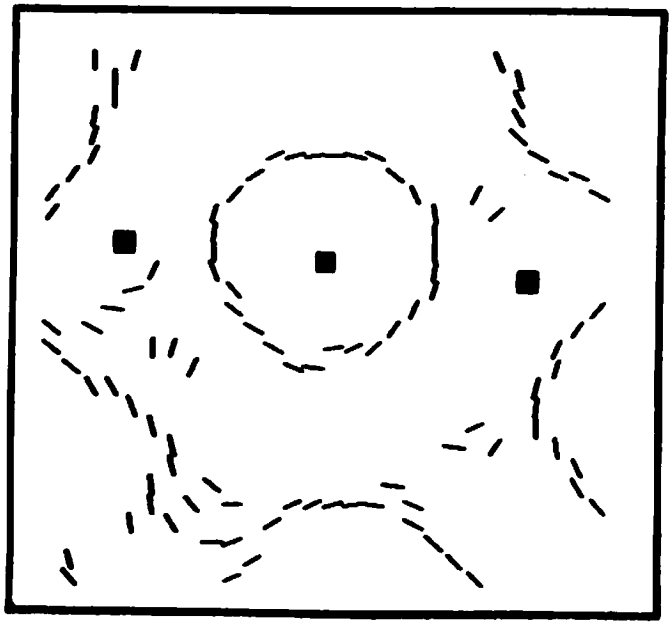

Figure 26: Edges & local extrema found from Figure 18, $\sigma = 2$

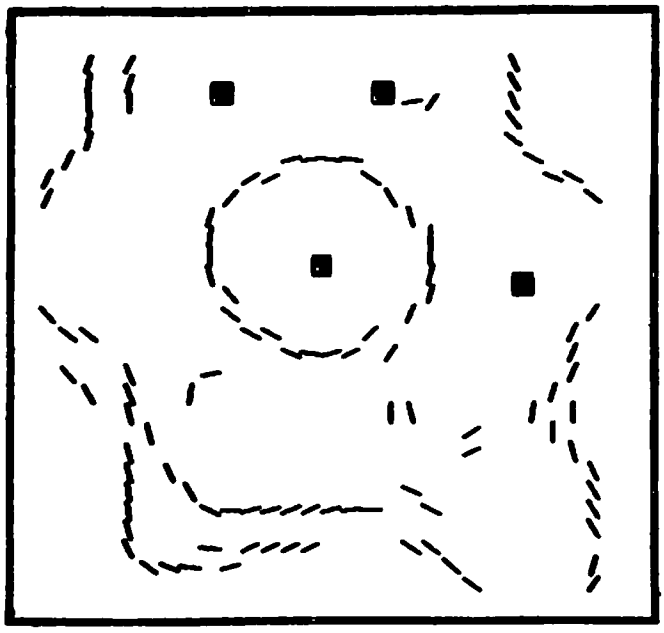

Figure 27: Edges & local extrema found from Figure 18, $\sigma = 3$

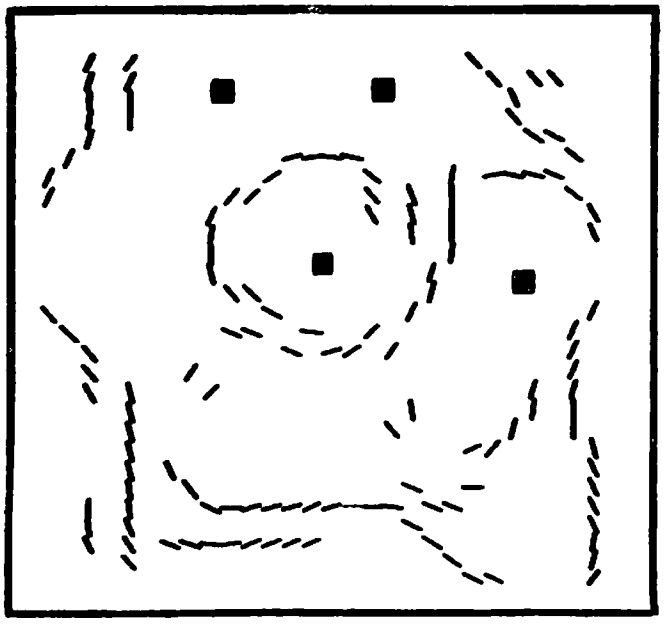

Figure 28: Edges & local extrema found from Figure 18, $\sigma = 4$