Feature Recognition Using Correlated Information Contained in Multiple Neighborhoods

Andrea Califano
Exploratory Computer Vision Group
IBM Thomas J. Watson Research Center
PO Box 704, Yorktown Heights, NY 10598

Abstract

Parameter transforms play a very important role in the recognition of geometric features in image data. Local operators devised to compute parametric descriptions of geometric entities using a small neighborhood p(x,y) about points of interest have been succesfully employed.

These operators fail to exploit the long distance correlations present in the image (distant points belonging to the same feature). Thus, their accuracy decreases with the order of the parametric properties (e.g., position, direction, curvature, torsion, etc.) and they are very sensitive to noise.

This paper presents a generalized neighborhood concept that allows parameter-extraction operators to use the joint information of different portions of the same feature. This produces up to a few orders of magnitude improvement in accuracy (signal/noise ratio) and a smoother response of the transform.

A general framework, based on a connectionist approach, is presented to deal with the complex response in parameter space generated by such operators.

A layered and concurrent scheme to extract 3D surfaces intersection curves is presented which, exploiting the properties of these operators, is able to reconstruct lines and conic sections in three-space.

1. Local Parameter transforms

The usual approach for the recognition of complex geometric features in images is to introduce a, possibly partial, parametric description of the feature of interest and then devise an operator to extract parameters from sets of image data. For example see [Kimme Ballard and Slansky, 1975; Slansky, 1978].

If multiple features are present in the data, the portion of image used as an input by the operator must be small enough to avoid, on average, the interference of more than one feature. Thus small, possibly overlapping, neighborhoods of fixed size are generally used; these are chosen such that the entire image is covered. Usually, for discrete samplings, the input to the operator is composed of the data points contained in a window of small size.

Several problems limit the usefulness of local operators:

- **Poor resolution and high noise sensitivity:** due to the multiple noise sources, namely measurement and quantization noise, and due to the limited size of the neighborhoods, only zero and first-order properties of the features can be extracted with sufficient accuracy. For a curve in three-space for instance, these would be the location and the orientation of a curve element. Recognition of features with higher-order properties (e.g., curvature for conics sections) is extremely difficult [Bolle, Kjeldsen and Sabbah, 1987].
- **Small number of points:** some features like surface intersection curves, given their one-dimensional nature, extend over $O(N)$ data points (pixels) of a $N \times N$ image. This results in a small number of neighborhoods (computational units) taking part in the transform.
- **Poor handling of complex parametrization:** complex parametric features are handled by leaving some of the parameters undetermined; this introduces a one-to-many mapping from image to parameters space. This works well providing that the dimension of the undetermined portion of parameter space is small [Shapiro 1978a; Shapiro, 1978b]. It can be experimentally shown that, when this dimension exceeds two, the parameter space tends to become overloaded so that no structure appears.

2. Long distance parameter transforms

For the above reasons, a different, non local approach is introduced. Such an approach must preserve some important characteristics of local operators, for instance their intrinsically parallel nature.

Let us consider a bidimensional image where each point (x,y) has an associated value z (intensity, range, etc.). We define $p(x,y)$ a neighborhood centered about the point (x,y), which is small compared to the image size. Let $P = \{p(x_i,y_i) ; i = 1 \dots N\}$ be a partition of the image, that is, a set of neighborhoods containing all the points of the image with possible overlapping. Let us define a new partition MP ("multi-partition") whose elements are all possible (unique) combinations of N neighborhoods $p(x,y)$ in groups of K:

$$MP = \{p(x_i,y_i) \cup p(x_j,y_j) \cup \dots \cup p(x_k,y_k); i = 1..N, j = i+1 \dots N, \dots\} \tag{2.1}$$

An element of the partition contains all the image data points of $P(x_i,y_i)$, $P(x_j,y_j)$ to $P(x_k,y_k)$. The total number n of (unique) elements of the partition MP is given by the combinations of N objects in groups of K:

$$n = \binom{N}{K} = \frac{N!}{(N-K)!K!} \tag{2.2}$$

where an element is considered not unique if it can be obtained from another by a permutation of the indi ces $i,j, \dots, k$. For in-

1 - Thanks to Ruud M. Bolle for his valuable contributions

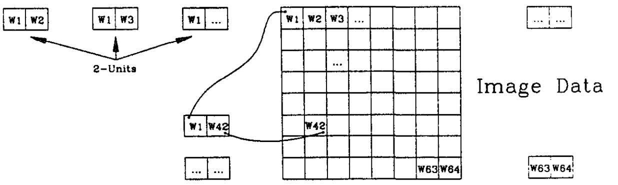

Figure 1

stance the two elements:

$$MP_{ji...k} \equiv MP_{ij...k} \tag{2.3}$$

would be considered equivalent and only one would appear in the partition.

For convenience we will use the word "unit" to indicate a local neighborhood and "K–unit" to indicate one of the generalized neighborhoods obtained by combining the local ones in groups of K (unit and 1–unit are synonyms). As an example, let us consider the case of a discrete image of size $L \times L$. If we choose P to be the set of all possible non overlapping windows of size $M \times M$ [$N = (L / M)^2$ is integer] we have N windows in the partition. If we choose a group size of two, the multi-partition MP is composed of $N(N-1)/2$ possible combinations of two windows as shown in figure 1.

3. Properties of nonlocal operators:

Using K-unit neighborhoods for parameter transforms has many advantages, among which:

Enhanced accuracy

The accuracy of the extracted parameters increases with the number of points taking part in the transform; the accuracy also depends on the relative position of the points on the feature. A K-unit contributes to the extraction of the parameter of a feature if all the composing local neighborhoods $p(x,y)$ contain only points on the feature. Thus, since a K-unit contains the points of K different units, the computation has an expected increase K in the number of active points, compared to a local transform where only the points of a single unit would be considered.

Another factor contributing to the enhanced accuracy is that distant points on the same feature are jointly used for the parameter extraction. In general, due to the complexity of the parameter transform model, it is impossible to quantify this contribution. However, using a rough and simplified model, we can quantitatively estimate the decrease in error when K-units are used, for the special case of a 2D–line of specified lenght parallel to the x-axis.

Suppose we have a set of points $\{(x,y)\}$ evenly spaced with respect to the x-axis over an interval of length L with the value of y distributed within a finite interval $2\Delta y$:

$$y(x) = y_0 \pm \Delta y \tag{3.1}$$

If Δy is small with respect to L, these points "correspond" to a line parallel to the x-axis. Now choose the partition to be the set of windows $\{W_i\}$ of size $M \times M$ centered about the points $\{(x_i, y_0)\ ;\ i=1\ ..\ N=L\,/\,l\,\}(l = x_{i+1} - x_i)$, with $l \le M << L$ (see fig. 2a). With these assumptions, all the points in a window are contained in a rectangle of size $2\Delta y \times M$, If we compute the angle of the line using 1-units, the worst case estimate for the angle is $\theta = \operatorname{atan}(2\ \Delta y\,/\,M)$ as shown in figure 2b. If we use the result of the computation for all possible N windows on the line to generate a distribution in parameter space, we can estimate an upper bound for the error by taking the average of the worst case errors of the single measures:

$$\sigma_1 = \frac{1}{N}\sum_{i=1}^{N} |\operatorname{atan}\Delta\theta| = \frac{1}{N}\sum_{i=1}^{N} \operatorname{atan}\frac{2|\Delta y|}{M} = \operatorname{atan}\frac{2|\Delta y|}{M} \tag{3.2}$$

which is, as one would have expected, the same of the single–measure case.

Let us now use the generalized neighborhood concept with two windows per group. As shown in figure 2c, considering two windows W_i *and* W_j, all the points are contained in a rectangle of size $2\Delta y \times [\operatorname{abs}(x_i - x_j) + M]$. Thus the maximum error for θ on the single measure is given by:

$$\Delta\theta_{ij} = \operatorname{atan}\frac{2\Delta y}{l|j-i|+M} \tag{3.3}$$

The N windows in groups of two generate $N\ (N-1)\,/\,2$ 2-units that can be used, like in the previous case, to produce a distribution on the parameter space; this time, the upper bound for the error becomes:

$$\sigma_2 = \frac{2}{N(N-1)}\sum_{i=1}^{N}\sum_{j=i+1}^{N} |\Delta\theta_{ij}| = \frac{2}{N(N-1)}\sum_{i=1}^{N}\sum_{j=i+1}^{N} \operatorname{atan}\frac{(2|\Delta y|}{l(j-i)+M} \tag{3.4}$$

Figure 3 shows the behavior of the ratio between the error in the two cases, compared to the number of windows over the line (in the example $\Delta y = M = l$).

Increased number of computational units

In some cases, especially with small images or low–dimensional features, the number of units that can take part in a local

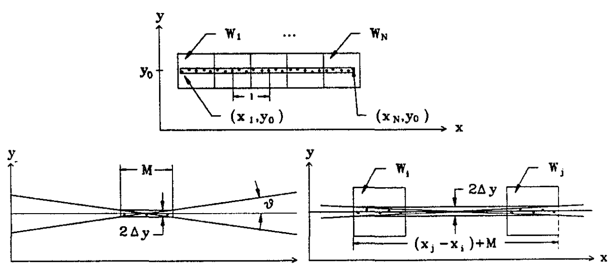

Figure 2(abc)

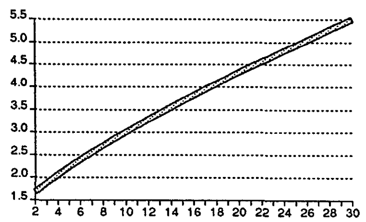

Figure 3: Accuracy increase.

parameter extraction can be very small. The generalized neighborhood concept allows for an increased number of such active units. If a feature has N units on it, the total number of K-units is given by (2.2):

Statistical analysis of local neighborhoods.

If we keep one of the neighborhoods $p(x,y)$ fixed and generate all possible K-units obtained by varying the other $(K-1)$ units we arrive at:

$$n = \frac{1}{N} \cdot \binom{N}{K} = \frac{(N-1)!}{(N-K)!K!} \qquad (3.5)$$

such K-units. This can be used for a statistical analysys for the fixed unit. For instance, then, it can be established that a unit has voted a certain number of times for the parameter vector of some specific features. This information is very important to define a competition process between mutually exclusive features generated by the same data points. This mechanism is described in detail in the next sections.

4. Disadvantages of non-local transforms

The trade-off for the enhanced accuracy is a complex response of the parameter transform, an increase in computation time and a nonlinear response of the transform.

Correlation between different features.

When generalized neighborhoods are used on images where several features are present, we have two possible configurations for a K-unit.

- All the units in the K-unit are located about different points of the same feature.
- The units are shared among different features.

In the first case, the transform produces an estimate for the parameters of the feature. In the second case, the units are still used to produce the parameter vector of a single feature and they generate an unpredictable value.

Figure 4 shows the result obtained by fitting lines to points on two different features. As shown, lines are scattered in all possible directions and positions generating a quasi-random response in parameter space (due to the deterministic nature of the process the distribution is pseudo-random [Lamperti]). Usually, due to the nature of the parameter vectors produced in this second case, these K-units contribute to the general noise background and no accumulation is produced in parameter space . However it is experimentally shown that false confidence peaks can be generated if several features are present in the image.

This is a consequence of the highly structured information present in the image which, when jointly used, produces correlated noise and hence false confidence peaks.

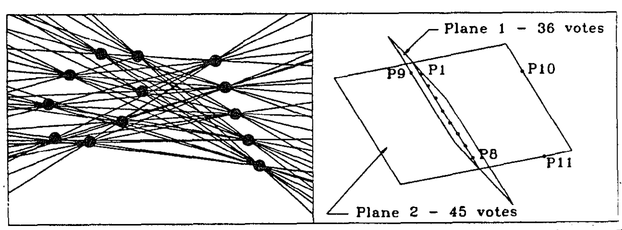

Figure 4,5

In particular, some feature can alter the response in parameter space of other features. For instance, as shown in figure 5, a line in three-space can amplify the response of all the planes containing it. In fact, any single point not on the line can cooperate with different combinations of points on the line producing always the same parameterse (those of the plane containing the line and passing through the point).

If there are N units on the line and one on the point, $\binom{N+1}{K}$ K-units will return the parameters corresponding to the plane as a result . This increases combinatorially with the lenght of the line and can produce erroneous hypotheses.

Combinatorial computational time explosion

The combinatorial increase in computation time, as the number of units per group grows, [given by (2.2)] is counterbalanced by the smoother response of the parameter extraction due to the increased number of K-units taking part to the transform. Computational complexity is inversely related to the size of the units. In fact, if the size is reduced by a factor M, the number of units in the partition increases by the same factor. This is due to the necessity of covering the entire image.

In this case, if N is the previous number of units in the partition, the number of K-units that participate to the transform becomes:

$$n = \binom{M \cdot N}{K} \qquad (4.1)$$

Nonlinear response

If two features of the same kind (two conic section, two planes, etc.) have different "size" (length for curves, area for surfaces, etc.) the local response of the transform is not a linear function of the size. As we have seen above, if there are N units on a feature, the total number $n(f)$ of K-units that produce an estimate for the parameter vector of the feature is given by (2.2). If N is large with respect to K, this value can be approximated by a simple power law:

$$n(f) \cong \frac{N^K}{K!} \qquad (4.2)$$

thus allowing for a simple renormalization rule [Califano and Bolle, 1987]. It is important to notice that due to the pseudo-random nature of the correlation noise, the renormalization should only be performed after having isolated the confidence peaks from the noise background. Since the latter is not subject to the same law the signal to noise ratio would be greatly deteriorated otherwise.

5. Recognition networks

The noise sources and the correlation induced by the interaction between different features in the image produce a large number of hypotheses in the parameter space. Thus a "filtering" mechanism has to be introduced in order to distinguish between noisy, erroneous and true hypotheses.

This can be elegantly accomplished by instanciating the parameter spaces as networks where nodes correspond to hypotheses characterized by the appropriate parameters. The links in the network are connections between nodes and computational units

(K-units) or between nodes themselves. Hypotheses are generated by K-units using these links; they can be partial or complete descriptions of geometric features.

Logically, each node computes the confidence of the network in the existence of the specific hypothesis in the input. This is done by assigning an activation level to the nodes. Parameter transforms are used to create or support nodes in the network correponding to the output parameters. Updating of the activation levels can be performed in a standard connectionist way [Feldman and Ballard, 1982; Sabbah, 1985]. If the acivation level of a unit falls below a noise threshold NT, the unit is deactivated and no longer takes part in the relaxation process.

Quantitatively the activation level of a node at iteration i, denoted by $AL_{node}(i)$, is computed as:

$$AL_{node}(0) = 0 \tag{5.1}$$

$$AL_{node}(i) = AL_{node}(i-1) + BU_{node} + TD_{node} - LI_{node}(i-1) - D \tag{5.2}$$

where

$$BU_{node} = \frac{WV_{node}}{k_{bu}E} \tag{5.3}$$

represents bottom-up reinforcement an TD_{node} top-down reinforcement. WV_{node} is a measure, ("weighted vote") assigned to the unit by the input data (K-units) through the parameter transform operators. That is, a measure of confidence that the corresponding hypothesis exists based only on data measurements (see [Sabbah and Bolle, 1986]). E is a normalizing factor and k_{bu} a rate parameter. The term D is the decay term that suppresses spurious hypotheses.

The lateral inhibition term LI is generated as a weighed sum of the activation levels of competing units. This term insures that semantically incompatible hypotheses inhibit each other so that stronger ones survive while the others are eventually deactivated by the noise threshold term NT in a "winner-take-all" network [Feldman and Ballard, 1982].

In previous connectionist networks involving parameter transforms [Bolle, Kjeldsen and Sabbah, 1987], hypotheses would mutually inhibit each other when their parameters are "close". This implies the generation of a metric in parameter space so that all hypotheses within a certain radius would be connected with inhibitory links.

This approach has some limitations, namely:

- The parametrization must be chosen to ensure the stability of hypotheses in parameter space [Liapunov, 1947]. This means that a small perturbation of the geometric feature in the image must produce a small perturbation of the parameter vector. This is not always possible. For instance, no matter how little the direction of a line is perturbed, the variation i the position of a point on the line (needed for the complete parametrization) can become arbitrarily large depending on where the perturbation is applied.
- Such parametric interaction does not make use of geometrical or topological knowledge (domain knowledge) about the hypotheses. Thus only hypotheses that have the same representation can interact and no inter- submitted to *ICCV Conf.* 1988.parameter space interaction is possible. For instance, the hypothesis of a sphere should support one of a circle of same radius in the image while the hypothesis of a conic section should be incompatible with one of a line if they have both been supported by the same image units.
- A totally symmetric interaction between hypotheses presents some difficult normalization issues. In fact, since the number of competing hypotheses grows as the dimension of the parameter space increases, one with a high level of activation, corresponding to an effectively existing feature in the image could be killed by a large number of competing ones just above the noise threshold. Also, once domain knowledge becomes an active element of the network structure, diffeent sources of inhibition should be normalized separately.

For the above reasons, a more involved interaction model is introduced based on the following rules:

- Hypotheses compete when:
  - Domain knowledge establishes their incompatibility

or

  - Their activation is supported by some identical image I-units (windows) and domain knowledge does not preclude their incompatibility. In this case the strength of the interaction is proportional to the percentage amount of common supporting I-units.
- Hypotheses support each other when domain knowledge establishes their mutual consistence.
- Hypotheses can only inhibit others with a lower or equal level of activation.

It is very important to notice that this kind of approach is consistent with K-unit parameter transforms. In fact, by definition, the K-unit structure allows different hypotheses to be supported by the same I-unit. This implicit knowledge can be used to correctly set up the inhibition network using the above guidelines.

When a hypothesis meets some existence criteria, namely a certain ratio between activation and lateral inibition, the corresponding feature is considered to be present in the image. In case that only a partial description of the feature is produced, the unit can start another process in a higher-level parameter space to extract the remaining parameters. A layered structure of parameter spaces is generated in this way.

The above mechanism allows for concurrent extraction of different geometric features and their mutual support or inhibition through LI - and TD -links.

This model for the interaction has proven very reliable and robust as shown in the section on curves extraction in three-space.

6. Computational complexity.

It is important to reduce the combinatorial explosion of the computation time with an increasing size of the image to a more appealing linear one. To do that we can introduce the concept of a "Search Radius" R_s such that each of the neighborhoods' center points (x,y), in a K-unit, is within a distance R_s from the center point of the next one (see figure 6).

Since R_s can be thought as the distance over which we expect two neighborhoods to possibly contain coherent information, we can think of R_s as a radius of coherence. Let NR be the total number of I-units (local neighborhoods) contained in a circle of coherence of radius R_s. Depending on the metric we select in order to define the circle of coherence, the number n of possible K-units generated by keeping the first unit fixed and choosing the others within R_s from each other is:

$$\binom{NR}{K} \leq n \leq (NR-1)^K \tag{6.1}$$

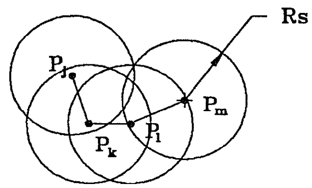

Figure 6

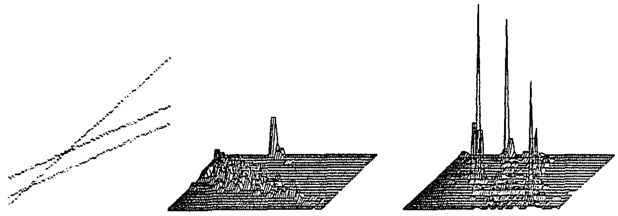

Figure 7(abc)

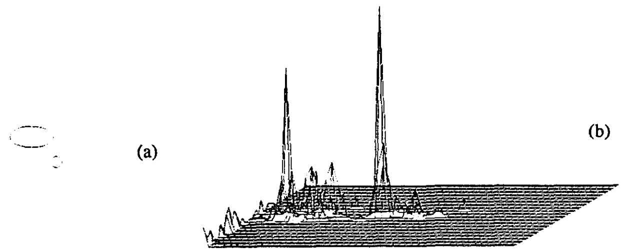

Figure 8(ab)

which does not depend on N, the total number of units on the image. Thus the total number of K-units N_{tot} becomes $N \cdot n$ due to the N possible choices for the first unit. N_{tot} is now a linear function of the image partition size N and thus of the image size itself. However, even if more contained, we still have a combinatorial explosion with respect to K.

It is important to notice that such a definition of the coherence radius allows for K-units to extend on distances much longer than R_s depending on the value of K. In fact, it is not required that all the units have centers within the coherence radius but only that a sequence can be established on the K-unit ordering, such that every 1-unit's center is within an R_s from the next one.

It immediately follows that, if K is larger than two, chains of units of maximum length $K R_s$ can be formed, thus allowing for longer distance correlation extraction.

Statistical properties of the distribution generated by K-units within a coherence radius are under further investigation. However the experimental results are very promising and the reduction of computation time is significant.

7. Experiments

Figures 7bc show the distributions generated in parameter space from a line parameter transform using the 128×128 pixels size synthetic image in figure 7a as an input. Lines are parametrized using their angle (x-axis) and their signed distance from the origin (y-axis). The parameter transform is based on a least square error line fitting algorith using points contained either in 1-units or in 2-units on the image. Here the units are windows of size 8×8. Due to the high level of noise in the image and to the small difference in the lines parameter vectors, the result for 1-units does not show a significant accumulation. In the distribution obtained with the multi-neighborhood approach, on the contrary, the parameter vector values corresponding to the three lines in the image show significant accumulation with respect to the noise background.

Figure 8b shows the distribution on a partial parameter space generated by an ellipse parameter transform. The x-axis corresponds to the ratio between the sqares of the two axis of the ellipse (a ratio of one corresponds to a circle) while the y-axis to the rotation angle of the ellipse with respect to the image x-axis. The parameter transform uses a fitting algorithm proposed by Bookstein [1979] based on the scatter matrix of the image data. The image is shown in figure 8a and it has the same size of the previous case. It contains a circle and an ellipse (ratio between the axis 2). The fitting algorithm extract in a single step the five-dimensional parameter vector for an ellipse. Since it is impossible to display a distribution on a five-dimensional parameter space, its projection along the two parameter axis is shown.

Even when only quantization noise is present, the local approach produces just scattered points in parameter space. Figure 8b shows only the results using 2-units. The same unit size of the previous case is used. As it can be seen, the values coresponding to the two image features (ratio 1, angle 0 for the circle and ratio 4, angl 0 for the ellipse) show a significant accumulation.

8. Curves recognition in three-space

We give here a brief outline of a layered and concurrent scheme for extraction of lines and conic sections in three-space with respect to the use of generalized neighborhoods. This system is part of a general effort of our group [Bolle, Kjeldsen and Sabbah, 1987; Bolle, Califano and Kjeldsen, 1988] for the recognition of objects generated by patches of planes and quadrics of revolution in range data images.

The recognition hierarchy for curve parameter extraction is shown in Figure 9. 3D edge detection techniques are used to generate maps of the image discontinuities.

Since the dimensionality of the parameter vectors of interest can range from four for 3D lines to eight for 3D ellipses or hyperbolas we divide our recognition process in two stages.

First using parameter transform based on a scatter matrix fitting algorithms [Duda and Hart, 1973] we concurrently search for (1) lines and (2) planes that contain intersection curves. Whenever a significant plane is found, we search for conic sections contained in the plane, again using a fitting algorithm based on the scatter matrix [Bookstein, 1979]. The best experimental compromise between accuracy and computability is to use 2-units for line extraction and 3-units for planes and conics extrac-

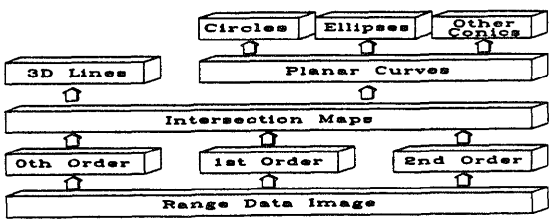

Figure 9

tion. The whole ensamble of parameter spaces is set up as a network with the structure described in section 5.

Figure 10a shows a noisy artificially generated 64×64 depth map (laser-range-finder generated images have also been used and produce similar results [Bolle, Califano and Kjeldsen, 1988]). It contains a cylinder with the axis aligned with the y-axis, with a hemisphere of the same radius at one end and a cone at the other. In this case, units have a size of 2×2 for lines extraction and of 4×4 for planes and conic sections extraction.

Figure 10b shows the situation after two iterations of the relaxation process. At this time 8 out of the 13 initially active line hypotheses and 9 out of the more than 80 plane hypotheses are still active.

Some of the planes, having satisfied our existance criteria, have initiated the search for conics. In fact, three conic section hypotheses associated with two planes have just been created.

After six iterations (see figure 10c) only the correct hypotheses for lines and conics have survived and their lateral inhibition has reached a level of zero so that they will continue to exist. The four lines are the linear limbs of the object, one circle is the limb of thesphere while the other is the intersection between the cylinder and the cone.

The secondorder discontinuity curve, between the cylinder and the sphere, is not found because the low-level edge operators do not detect such discontinuities. A promising operator for second order discontinuities is under investigation.

Figures 10b and 10c display the projection of the active curves at the two stages of recognition. Activation levels are indicated by the gray level (darker = more active). Planes are not displayed since we only use them as an intermediate step for conics extraction.

More than the number of existing planes (2) have been found in the image after our relaxation process. This is due to the interaction between lines and planes as described in section 4. This contribution can be eliminated by changing the lateral inhibition

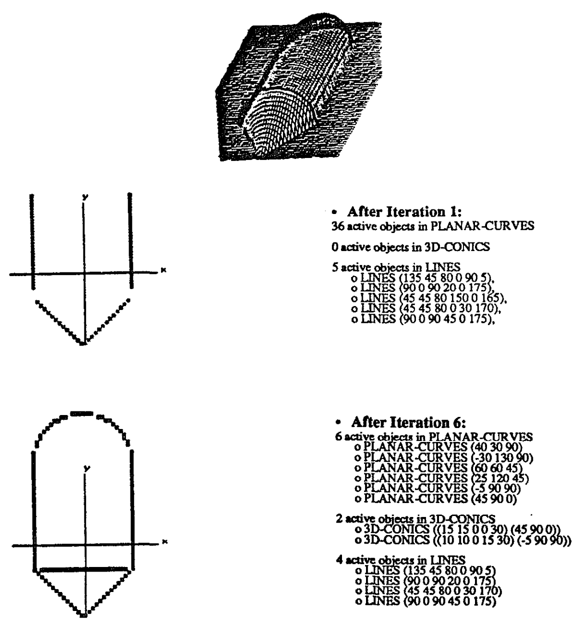

Figure 10(abc)

model for plane hypotheses to include links with the line parameter space. A new model is under investigation.

9. Conclusions

A new approach to parameter extraction in images has been proposed where using correlated evidence from distant part of the image allows for complex parametric feature recogniion. Up to eight-dimensional features have been experimentally reconstructed from synthetic and laser-range-finder range data images.

References

[Bolle, Kjeldsen and Califano, 1988] Ruud M. Bolle, R. Kjeldsen and A. Califano, "Evidence Integration of Curves and Surfaces for 3D Object Recognition," IBM TR RC 13592, 1988.

[Bolle, Kjeldsen and Sabbah, 1987] R.M. Bolle, R. Kjeldsen and D. Sabbah, "Primitive shape extraction from range data," in *Proc. IEEE Workshop on Comp. Vision*, Nov.-Dec. 1987, Miami Beach, FL, pp. 324–326; also IBM Tech. Rep. RC 12392, AI System Group, IBM T.J. Watson Res. Center, July 1987.

[Bookstein, 1979] F.L. Bookstein, "Fitting conic sections to scattered data," *Computer Graphics and Image Processing*, Vol. 9, No. 1, Jan. 1979, pp. 56–71.

[Califano and Bolle, 1987] Andrea Califano and Ruud M. Bolle, "Localized noise propagation effects in parameter transforms," to appear in *Proc. SPIE Conf. Intell. Robots and Comp. Vision*, Nov. 1987, Cambridge MA.

[Duda and Hart, 1973] R.O. Duda and P.E. Hart, "Pattern Classification and Scene Analysis," New York, John Wiley & Sons, 1973

[Feldman and Ballard, 1982]J. A. Feldman and D.H. Ballard, "Connectionist models and their properties," *Cognitive Science*, Vol. 6, pp. 205–254.

[Kimme, Ballard and Slansky, 1975] C.Kimme, D. Ballard and J. Slansky, Finding circles by an array of accumulators," *Comm of the ACM*, Vol 18, No. 2, Feb 1975, pp. 120-122

[Lamperti] J. Lamperti, "Stochastic Processes", *Academic Press*, New York

[Liapunov, 1947] A. M. Liapounov, "Probleme generale de la stabilite' du mouvement," *Ann. of Math. Studies*, N. 17, Princeton Univ. Press., 1947, Princeton, NJ.

[Sabbah, 1985] D. Sabbah, "Computing with connections in visual recognition of Origami objects," *Cognitive Science*, Vol. 9, No. 1, Jan-March 198, pp. 25–50.

[Sabbah and Bolle, 1986] D. Sabbah and Ruud M. Bolle, "Extraction of surface parameters from depth maps viewing planes and quadrics of revolution," in *Proc. SPIE Conf. Intell. Robots and Comp. Vision*, Oct. 1986, Cambridge MA, pp. 222–232.

[Shapiro, 1978a] S.D. Shapiro, "Properties of the transform for the detection of curves in noisy pictures," *Comp. Graphics and Image Processig*, Vol. 9, 1978, pp. 219–236.

[Shapiro, 1978b] S.D. Shapiro, "Feature space transforms for curve detection," *Pattern Recognition*, Vol. 10, 1978, pp. 129–143.

[Slansky 1978] J. Slansky, "On the Hough transform for curve detection, *IEEE Trans. on Comp.*, Vol. 27, No. 10, Oct. 1978, pp.923-926.