DEPARTMENT OF ARTIFICIAL INTELLIGENCE
UNIVERSITY OF EDINBURGH

DAI Working Paper No. 137
Date: March, 1983

Title:

Anomalies Occurring in the Image Formation and Analysis Process

Author:

Robert B. Fisher

Abstract:

Many object recognition and scene analysis programs fail to work except under severe constraints, and only then with considerable laboratory adjustment. This paper takes the view that the anomalies which often cause such programs to fail are actually an important part of the semantics of intermediate level vision. A collection of the anomalies are listed and categorized according to their major source: object, scene, illumination, imaging, model or processing properties. It is suggested that the problems should actually be regarded an integral part of the image recognition process, and that any reasonable interpretation program has not only to cope with, but also explain the problems.

Acknowledgements:

The author was appreciatively supported during the course of this work by a postgraduate studentship from the University of Edinburgh. Special thanks also go to R. J. Beattie, W. Caplinger, J. Hallam, J.A.M. Howe, R. J. Popplestone, and B. L. Yin for their discussion and advice on both the content and exposition of this paper.

1. Introduction

Recognition of isolated simple man-made objects, with bold and distinct features, is reasonably straightforward. Examples of successful recognition projects include: teacups ([Bar71]), blocks world configurations ([Fal72]), two-dimensional industrial parts ([Per77]), telephones ([Shi78]), roadways ([Bol79]), three-dimensional industrial components ([Vam79]), airplane profiles ([Bro81b]), and rectangular tables ([Fis82]). Unfortunately, their success has been due to near-perfect input data (i.e. constrained by high resolution, high contrast, high dynamic range and low noise), perfect segmentations, special light sources or configurations, largely unobscured visibility, no coincidental alignments, and limited ranges of object orientations. However, when turning to the problem of recognizing objects in more interesting or complex scenes, recognition often fails because the existing methods depend upon adjustment of scene and processing parameters. Yet, in retrospect, it is usually clear why the interpretation process failed (at least to the implementor).

In practice, the majority of real images have anomalies sufficient to cause the recognition process to fail (i.e. the programs do not degrade gracefully). These anomalies include: incomplete edges, unconnected boundaries, low contrast surfaces (which blend into the background), obscured objects, self-obscured components of objects, coincidental alignments (which confuse the segmentation and recognition processes) and shadows (which create new regions and obscure boundaries). As some of the problems occur even in ideal images (e.g. obscuration), it is not sufficient to blame the equipment or algorithms.

Thus, it seems that the range of images input to a given image understanding program must meet rather severe constraints in order for the program to be successful. On the other hand, if the program had greater knowledge of the actual constraints of physical objects and the image formation process, it would not be so limited by the artificial constraints listed above.

This paper starts by examining the anomalies which occur when analyzing an image. These are classified according to the types of anomaly (e.g. self-obscuration) and the result observable in the recognition process (e.g. invisible structures). Also included is a breakdown of affecting factors: illumination, image formation, object models, scene arrangement, scene objects and processing algorithms. This is followed by a consideration of a schematic model of the recognition process, detailing the information path and where instances of the anomalous data first occur. (We use the term information path to denote the sources of information, the processes which affect it, and the process interconnections. The path consists of a flow from the elements which make up the raw scene, through the raw image, to the final extracted information. The path is illustrated in figure 1 and is discussed later in the paper.) Lastly, there is some discussion about why the anomalies should be viewed as typical occurrences in images, and that reasonable scene analysis programs should be expected to cope with them.

To some extent, the types of anomalous data that appears in the information path is dependent upon the scene analysis goals. For example, systems which calculate the shapes of visible surfaces are probably not concerned with the obscured elements of the scene. Consequently, this loss of information doesn't limit the output of the process. A second example is systems which measure the area of crop regions, and are thus probably not sensitive to spurious boundary segments. In the discussion of anomalies which follows, a reasonably specific type of scene analysis is considered. The overall goal of the processing is to analyze an image (the type of which is discussed below), using a set of models, to give a description of the scene. The description includes locating and orienting the objects in the scene, delineating all possible sub-structures contained within the model, and explaining why the other substructures were not found.

Prior to a discussion of the anomalies, the scene domain should be considered. The images input to a hypothetical scene analysis program are taken from the everyday man-made macroscopic world. This excludes most radar, sonar, holographic, x-ray, satellite, microscopic and telescopic images. The scenes are usually three dimensional, and mostly contain opaque compact objects. These restrictions still allow most images that have been studied by artificial intelligence researchers, with the exceptions of pulmonary, cardiac and industrial inspection x-rays, medical cytographic images and satellite landmass survey. The major types of scenes examined by vision researchers include: blocks world, simple curved everyday objects, industrial components, office and laboratory scenes, and cartoon drawings.

2. Preliminaries

An anomaly is defined to be an occurrence in the formation of an image or its processing which causes erroneous (anomalous) data to pass to succeeding stages of analysis. (More precisely, it is an occurrence in the information path which causes the loss of correct or introduction of incorrect data into the information flow.) Anomalies arise partly because of the natural limitations on obtaining perfect information about the world, and partly because of exceptions to the first order model of the scene, imaging process and data analysis. The anomalies cause three effects to appear in the information path - loss of data, incorrect data, and unexplained data. The sources of the limitations are: incomplete data at the sensors, imperfect symbolic models of the world (including the objects, illumination, and imaging process), selection of an inexact model to explain the data, and the complex process of transforming one model instance into another. The first source is purely physical, the last is purely symbolic, and the other two relate to the barrier between data and its description.

There are three major classes of anomalous data: the "present, but erroneous" (type I), the "needed, but non-existent" (type II), and the "present, correct, but unaccounted for" (type III). An example of the first type is an edge formed by the coincidental alignment of a file cabinet side with the vertical edge in the corner of a room. Examples of the second type are gaps in edge segments. The third type of data would arise from regions formed by shadows, or from texture regions on the surfaces of objects. The type I errors are responsible for seriously slowing down the rate of recognition, by the introduction of spurious hypotheses, and for producing false recognitions. The type II errors cause recognition failures (i.e. insufficient information to trigger the models) and partial recognitions (incompletely instantiated models). The type III errors cause incomplete analyses, i.e. unexplained features in the results and hence image.

Effects of anomalous data include:

- failure to recognize an object
- mis-identification of objects or image structures
- inclusion of components not actually in the object
- omission of subcomponents

The anomalies arise as a result of the various factor processes in the generation of a description of a scene. These form the contributory factors affecting the appearance and location of the anomalous data. Six major categories have been isolated, and each has several specific affecting sub-factors. They can be broadly classified as:

- <u>illumination</u> - its direction, intensity and spectrum
- <u>image formation</u> - sensor location and orientation, sensor response, system noise, scene distance
- <u>scene object properties</u> - surface shape and reflectance, opacity, compactness, structure
- <u>scene construction</u> - object location and orientation
- <u>object modeling</u> - surface shape and reflectance models, boundary types, component relations, flexibility, parametric/model class properties, constructive/descriptive mix, simplifications
- <u>processing algorithms</u> - locality of decision making, sensitivity, exactness of matchings, matching criteria, discrimination criteria

In general, several factors are involved in each anomalous situation, as in obscuration, which is the result of the scene configuration, the user viewpoint, and the relative size of the objects.

3. An Analysis of the Anomalies

This section largely consists of table 1, which contains a breakdown of the anomalies occurring in everyday images. The analysis contains the following information:

Type – the general descriptive title given to the phenomenon which caused the problem, e.g. obscuration. This is typically a class of problems, which occur in a variety of ways and have a variety of observed effects.

Factors – the factors which contributed to the formation of the anomalous condition. These are denoted as primary (P) and secondary (S). The notion of primary is such that, if some relevant criteria is met, then all images will have instances of the anomaly (e.g., any illumination generally means shadows). Secondary factors control the presence or absence of particular instances of the anomalies in the particular images (e.g. the location of objects affects the location of shadows). It is also noted that the primary factor corresponds to the location in the processing schema where the anomalous data is first present (see Section 5 and Figure 1). The factor classes are: illumination (il), image formation (if), scene construction (sc), scene objects (ob), object modeling (md), processing algorithms (pr).

Results – these are the observed anomaly phenomenon which occur during the recognition process. They are the actual manifestations which cause the analysis to fail. It is noted that the presence of one result often means others as well, as in the case of shadows, which obscure some boundaries while creating others.

Category – each Result is denoted as type I – present, but erroneous result, as type II – missing, but correct result, or as type III – present, correct, but unaccounted for results.

The following types of results have been found to occur. As well as being separated by type, they are also distinguished as manifesting their effects at the low level (L) or high level (H) of image analysis.

Type I – present, erroneous

new single boundaries (L) – a single non-object boundary in some image region

new single regions (L) – isolated non-object regions

merged regions (L) – a structurally insignificant region formed from the merger of two or more other regions

merged boundaries (L) – boundary consisting of several scene structures

oversegmentation (L) – multiple new boundaries or regions

distorted structure (L) – components present, but with incorrect metrical relationships (as from geometric camera distortion)

unrelated matchings (H) – inappropriated linked segments or subcomponents

non-existent objects (H) – incorrect analyses

indiscriminable hypotheses (H) – several hypotheses covering the same data

unverifiable hypotheses (H) – insufficient data prevents evaluation

multiple usage (H) – same structure appears in several different objects

global inconsistency (H) – improperly constructed object, but all local relations hold according to model

duplicate recognitions (H) – multiple instances of same object found in the same location

Type II – missing, correct

boundary loss (L) - scene boundary not detected
region loss (L) - scene region not detected
missing objects (H) - present objects not reported
incomplete objects (H) - subcomponents missing from reported objects
structure loss (H) - structure not detected (possibly any of above)

<u>Type III - present, correct, un-accounted for</u>
unexplained boundaries (L) - unexplained boundary interior to object
unexplained regions (L) - unexplained regions interior to object, e.g. painted areas, texture areas
unexplained structure (H) - e.g. painted images, photographs of scenes appearing in scenes, mirror reflections, unmodeled objects

Table 1 - Analysis of the Anomalies

Types	Factors (P=primary, S=secondary) il	if	sc	ob	md	pr	Results (I - present, but erroneous) (II - missing, but correct) (III - present, correct, but unaccounted for)
1. overlighting	S	P		S			boundary loss(II), region loss(II), new single boundaries(I), new single regions(I), merged regions(I)
2. shadows	P		S	S			unexplained regions (III), unexplained boundaries(III), boundary loss(II), region loss(II)
3. highlights	P	S	S	S			new single regions(I), new single boundaries(I) boundary loss(II)
4. reflections	S	S	S	P			unexplained boundaries(III), unexplained regions(III), unexplained structure(III),
5. weak obscuration boundaries	S		P	S			new single regions(I), boundary loss(II), region loss(II),
6. opaque obscuration		P	S	S			structure loss(II)
7. granular obscuration (fog, tree branches)		P	S	S			oversegmentation(I), boundary loss (II), region loss (II)
8. coincidental alignments (visual)		P	S	S			merged boundaries(I), unrelated matchings(I)
9. scale (too close to viewer)		P	S	S			unverifiable hypotheses (I), indiscriminable hypotheses(I), incomplete recognitions(II), unexplained structure(III)
10. scale (too far from viewer)		P	S	S			structure loss (II)
11. object proximity		S	P	S		S	unrelated matchings(I), structure loss(II)
12. object co-location (intertwining)		S	P	S		S	merged boundaries(I), unrelated matchings(I)

Table 1 - Analysis of the Anomalies cont.

Types	Factors (P=primary, S=secondary)						Results (I - present, but erroneous) (II - missing, but correct) (III - present, correct, but unaccounted for)
	il	if	sc	ob	md	pr	
13. model omissions (no models)				S	P		unexplained structure(III), missing objects(II)
14. unmodelable objects (inadequate models)			S	P	S		missing objects(II), unexplained structure(III)
15. coincidental self-obscuration (camera viewpoint or flexible objects)		P	S	S			structure loss(II)
16. structural self-obscuration (concavities)		S	S	P			structure loss(II)
17. inherent self-obscuration (object backsides)		S	S	P			structure loss(II)
18. surface texture				S	P		oversegmentation(I), unexplained regions(III)
19. reflectance irregularities (surface markings)				S	P		unexplained boundaries(III), unexplained regions(III)
20. transparent objects		S		P			region loss(II)
21. defective objects				P	S		incomplete objects(II), unexplained structure(III), distorted structure(I)
22. ambiguous components (similar objects)				P	S		indiscriminable hypotheses(I), multiple usage(I)
23. model non-discrimination (undistinguished objects)				S	P		indiscriminable hypotheses(I)
24. algorithm spatial resolution limits (scale effects)		S	S	S		P	merged regions(I), merged boundaries(I), unverifiable hypotheses(I), structure loss(II), incomplete recognitions(II)
25. algorithm grey-scale resolution limits (threshold effects)	S		S	S		P	merged regions(I), boundary loss(II)

Table 1 - Analysis of the Anomalies cont.

Types	Factors (P=primary, S=secondary)						Results
	il	if	sc	ob	md	pr	(I - present, but erroneous) (II - missing, but correct) (III - present, correct, but unaccounted for)
26. approximate continuity (undistinguished coincidental alignments)		S	S			P	merged boundaries(I), boundary loss(II)
27. noise	S	S				P	oversegmentation(I), structure loss(II)
28. inconsistent reasoning						P	multiple usage(I), duplicate recognitions(I), global inconsistency(I)
29. scene alignments		S	P	S			merged regions(I), merged boundaries(I), unrelated matchings(I)
30. model simplifications				S	P	S	missing objects(II), unexplained structure(III)
31. out of focus		P	S				boundary loss(II), region loss(II)
32. image distortions		P					distorted structure(I)

5. Where the Anomalies Arise

This section gives a schema for the image formation and analysis process, showing at which points the various anomalies arise. It is shown in figure 1. Each of the square boxes represents a set of parameters, and each of the circular boxes represents one of the factor processes discussed in section 2. Each arc is labeled with the type of information delivered to the next level in the processing.

The major worth of this figure is that the numbers beside the arcs designate the anomalous data associated with anomalies listed in section 3. The location of the number shows where the effects of the anomaly first appear in the information path. This means that the anomaly occurs in at least one instance of the information, given random selection of the other factors (such as illumination) as necessary. It is noted that the anomalous data always appears just after its associated primary factor.

The processing portion of the model is contained in the bubbles labeled segmentation and model matching. The segmentation bubble is intended to encompass the extraction of the various low level information (e.g. intrinsic images) such as boundary locations, intensity regions, depth regions, and surface shapes. The model matching encompasses the use of the low level information to support the selection and verification of models for scene structures, and for placement of instances of those structures in a model of the scene. In particular, the model matching should have as output a list of objects, their location, and the linkages from the model components to the related segmentation data.

4. Related Research

There is previous research having to do with the processing of imperfect edge data, the recognition and interpretation of obscuration, and the interpretation of shadows. In most cases, the work focussed upon the anomaly as a side issue, rather than as a key feature of scene analysis.

Some of the best research on working with incomplete edge data was Falk's **INTERPRET** program ([Fal72]). The goal of this work was recognition of objects and structures in blocks world scenes. In the research, he recognized that edges or corners of objects could be potentially missing (partially or wholely) due to obscuration or low contrast induced segmentation failures. Further, he concluded that some corners could be improperly labeled due to coincidental alignments. He then attempted to sort out the problems by use of knowledge of edges (continuity) and of blocks (connection of edges at corners) and by the use of special completion routines. Edges and junctions were labeled according to the topology of visible edges connecting at the junctions. In the case of bad labelings, the type of label was used to help segment the scene into bodies. The program didn't attempt to explain the causes of the problem, only to account for a few and correct them. Further it was largely limited to corrections in the context of blocks world scenes (except for edge completions). More general scenes will not have the highly constraining vertex labeling semantics, which limits the direct use of his techniques.

When it comes to obscuration and self-obscuration, perhaps **ACRONYM** ([Bro81b]) represents the best work. This program interpreted images based upon constraining the parameters of the camera and scene objects. This occurred through a matching of data to models, as mediated by a sophisticated constraint manipulation system. It used a hierarchical object model based around generalized cones as primatives, which were attached using coordinate frame transforms. The model sizes and attachement parameters could be either fixed constants or variables. Image primitives were chosen to be ellipses and ribbons (2D projections of generalized cones). At any stage of its processing, it attempted to predict what objects should look like in the image, given the current value ranges for the model parameters. This was a graph structure with nodes representing image structures, and arcs representing relationships that must hold between the structures. Once a image structure was matched to a predicted structure, then actual size or orientation measurements from the image were used to back-constrain the original parameter estimates. This was useful for both obtaining consistent object interpretations and deducing global camera position parameters.

Brooks argued that by using geometrical reasoning, based upon estimates of object position and camera parameters, **ACRONYM** could deduce that certain features would not be not visible (prediction of invariant and quasi-invariant features). In particular, it should be able to deduce that certain components are not visible due to obscuration or self-obscuration.

ACRONYM's geometric reasoning and image recognition has not yet been adequately demonstrated for 3D objects. The most successful examples show identification of parts of several airplanes, including subclass identification and estimates of the camera parameters. However, as the airport scene is viewed from a large vertical height, the analysis is largely two-dimensional. Further, as only a few parts are detected, the analysis seems inadequate. Brooks places much of the blame for this on poor segmentation.

Waltz ([Wal75]) considered several anomalies in his celebrated work on interpreting blocks world images. The major work involved assigning a set of candidate labels to each edge junction, and algorithms for finding a compatible, preferably unique labeling for the image from these initial assignments. A prominent feature of his scenes were shadows cast by the blocks. His junction labelings explicitly included shadow constructions. In fact, he pointed out that shadow information was needed to determine some contact and support relations. Further, he included some junction types caused by common missing line and coincidental alignment configurations. This approach is again limited to the highly constrained blocks world and required essentially perfect line/vertex data.

Recent work in machine vision has started to concentrate upon what are the relevant data, symbolic descriptions and relationships (i.e., semantics) at the various levels of image understanding. The work discussed in this paper is a direct result of this philosophical change. In particular, the analysis of anomalies contained herein contributes to an understanding of the information flow in the image formation and analysis process. Further, many of the anomalies are distinct features of the semantics of scenes and images. These include shadows, highlights and obscuration. In the analysis, a description of actual process of scene analysis is considered to be part of its own semantics, in that any analysis process is bound to make errors.

We briefly survey here some of the relevant research into semantic vision. It is classified into low level (extraction of knowledge of scene from image features), intermediate level (deduction of scene objects) and high level (interpretation of scene).

At the low level, there are good specific examples in the work of Horn ([Hor75]) on surface shape, Woodham ([Woo77]) and Brooks ([Bro81a]) on surface orientation, and Beattie ([Bea82]) on the semantics of boundaries during the image forming process. Barrow and Tenenbaum ([Bar78]) considers the low level in general in their discussion on the various intrinsic images present (reflectance, illumination, orientation, distance), and then some of the semantics of boundaries in images.

At the intermediate level is some recent work by Lowe and Binford ([Low81]) on interpreting shadows and boundaries as evidence of three-dimensional structure. The work of Shirai ([Shi75]) on segmenting blocks world scenes is also relevant. He used the semantics of rectangular planar surface objects (i.e. the types of junctions) to suggest where to look for the boundaries separating the planes. Though this is knowledge based interpretation, the knowledge is conceptually close to the task. Binford ([Bin81]) suggested a direct approach to the intermediate level semantics in his work on inferring surfaces. He used three sources of information: depth ordering cues, relative orientations of regions from relative shadow orientations, and region boundary junction interpretations. A second important point he made concerned the preferential ordering of hypotheses, according to scene likelihoods, when interpreting ambiguous phenomena.

This author is uncertain as to what constitutes good semantics at the highest levels of interpretation (i.e., "What is this scene?"). Perhaps Tenenbaum's ([Ten73]) scene based constraints was an early example (e.g. "Door is hinged to wall on one vertical edge and adjacent to floor on bottom").

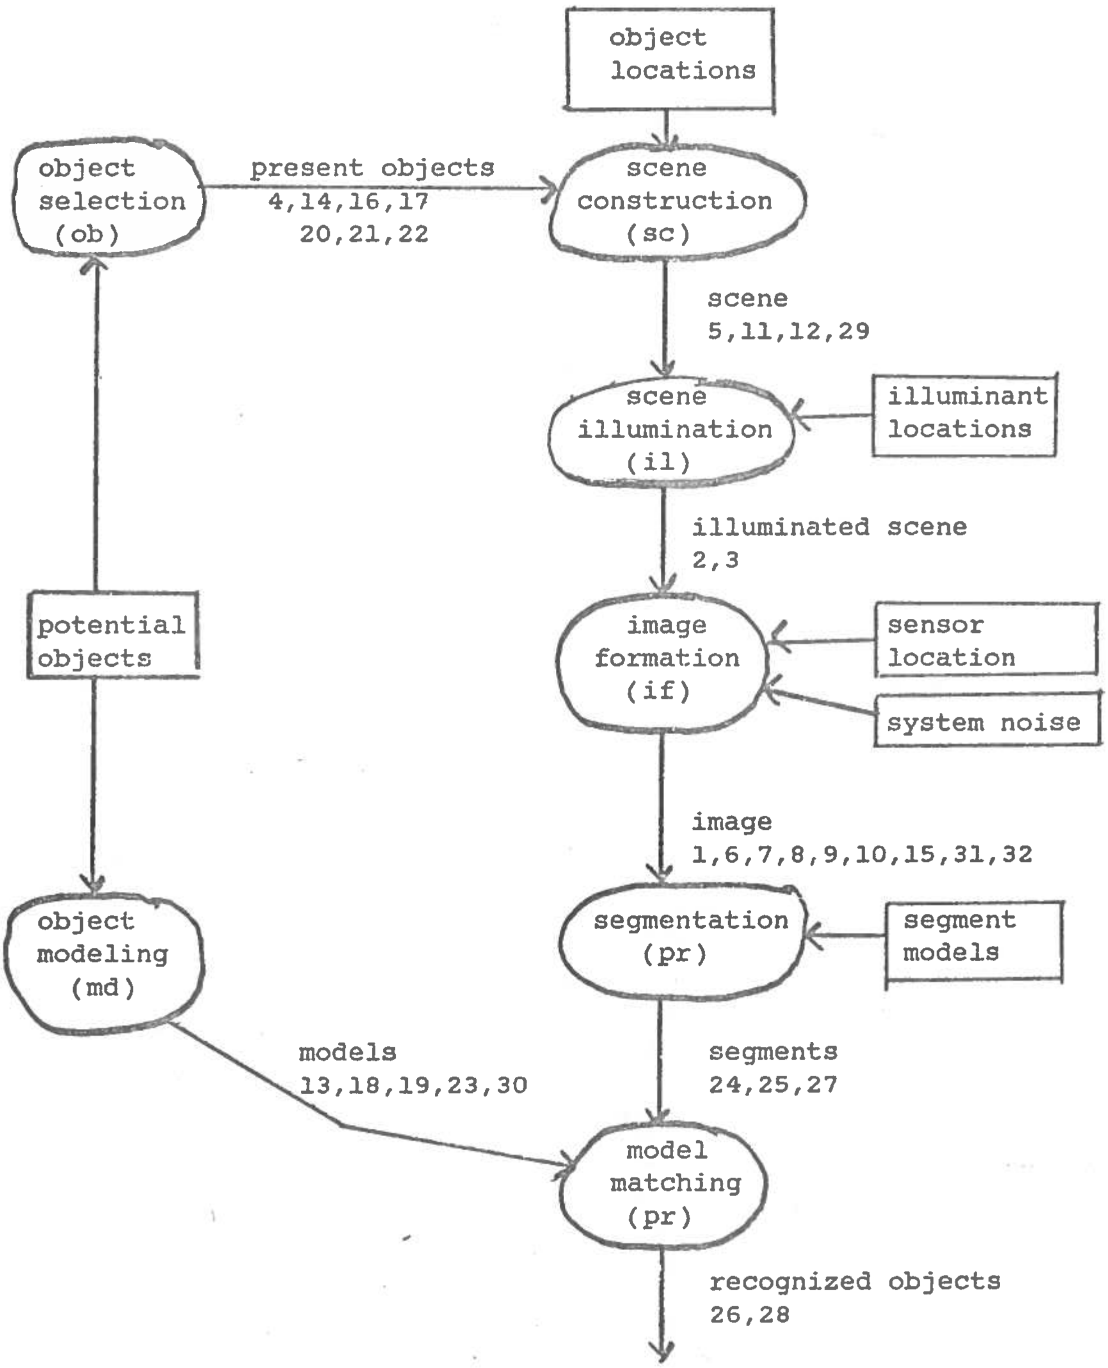

<u>Figure 1 - image processing schema and anomaly appearance</u>
(numbers beside arcs refer to anomalies - see section 3)

6. Not Anomalies, but Normal Events

This paper started with a definition of anomaly: a situation or event which interfered with the analysis of a image. However, after examination of the factors affecting the anomalies (section 2) and the anomalies themselves (section 3), it seems apparent that most of the anomalies are largely dependent upon factors external to the recognition process and are thus unavoidable. Further, given the nature of most of the processing-based anomalies, such as the effect of system noise on segmentation, many of these are unavoidable as well. Hence, it seems necessary to re-think the notion that these occurrences are anomalies; rather they are just typical features of images which, unfortunately, cause problems for recognition programs.

The question that arises, then, is to what degree do these so-called anomalies occur? Without a detailed study, a quantitative answer is not possible. But, an overview of the list in table 1, shows that many are present in the class of images defined in section 1. These include: shadows, highlights, weak reflectance boundaries, obscuration, scale, object proximity, model omissions, self-obscuration, reflectance and texture irregularities and scene alignments.

So, it seems to be a necessary conclusion that every image is going to contain a selection of these features, and any reasonable image analysis program will have to cope with them. (Most programs of the past have not explicitly coped with most of the anomalies.) Moreover, because of their omnipresence, any program which effectively describes a scene will also have to describe these situations, as they are aspects of the scene. This is in addition to the necessity of coping with them directly, for the purpose of actually reasoning out the scene, as well as for explaining their presence.

In table 1, some of the relationships between the types of anomalies and the results were described. In section 5, the location of the results of the anomalies in the image analysis process were examined. In essence, these two sets of information are the beginnings of some meta-level knowledge of the recognition process. This knowledge embodies the view that a subprocess is not a perfect filter on the information, but actually has characterizable behavior. Hence, when the presence of certain patterns in the output of the subprocess is detected, meta-level knowledge of the subprocess can be used to more effectively reason about what must have been the input to the subprocess. In particular, the presence of specific types of anomalous data at various stages in the recognition process (e.g., the Results in Table 1) suggests a set of possible anomalies (Table 1). These in turn may be discriminable via alternative reasoning methods. The net effect is to deduce aspects of the scene that were not effectively extracted by the base level process. This form of active reasoning (like Freuder ([Fre77])) has the potential for better image analysis. The overall reasoning schema is summarized in figure 2 below.

In this schema, it is suggested that the meta-process can use the output of the process, as well as its knowledge of the process, to generate a more complete and correct description of the scene. This would occur both by augmenting the description of the scene, as output by the base process, and by correcting or adding to the information input into that process.

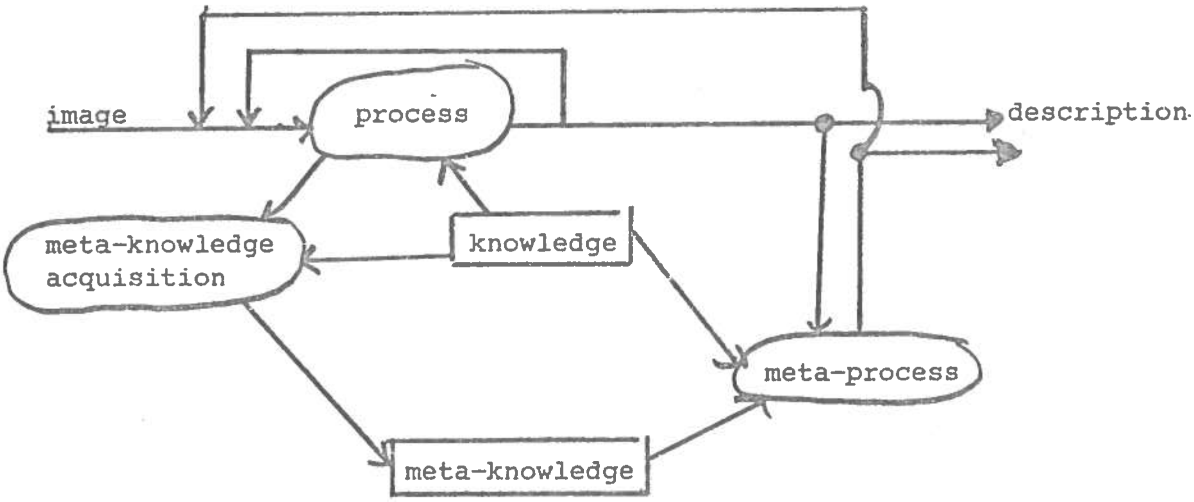

Figure 2 - schematic view of a meta-level reasoning structure

7. References

[Bar71] Barrow, H.G., Popplestone, R.J., "**Relational Descriptions in Picture Processing**", in Meltzer and Michie (eds), Machine Intelligence 6, pp377-396, 1971.

[Bar78] Barrow, H.G., Tenenbaum, J.M., "**Recovering Intrinsic Scene Characteristics From Images**", in Hanson & Riseman (eds), Computer Vision Systems, pp3-26, 1978.

[Bea82] Beattie, R.J., "**Edge Detection For Semantically Based Early Visual Processing**", Proc. 1982 European Conf. on Artif. Intel., forthcoming.

[Bin81] Binford, T.O., "**Inferring Surfaces from Images**", Artificial Intelligence, Vol. 17, pp205-244, 1981.

[Bol79] Bolles, R.C., Quam, L.H., Fischler, M.A., Wolf, H.C., "**Automatic Determination of Image-to-Database Correspondences**", Proc. 6th IJCAI, pp73-78, 1979.

[Bro81a] Brooks, M.J., "**Surface Normals From Closed Paths**", Proc. 6th IJCAI, pp98-101, 1979.

[Bro81b] Brooks, R.A., "**Symbolic Reasoning Among 3-D Models and 2-D Images**", Artificial Intelligence, Vol 17, pp285-348, 1981.

[Fal72] Falk, G., "**Interpretation Of Imperfect Line Data As A Three-Dimensional Scene**", Artificial Intelligence, Vol 3, pp101-144, 1972.

[Fis82] Fisher, R.B., "**A Structured Pattern Matching Approach to Intermediate Level Vision**", Edinburgh Dept. of AI Research Paper #177, 1982.

[Fre77] Freuder, E.C., "**A Computer System For Visual Recognition Using Active Knowledge**", Proc. 5th IJCAI, pp671-677, 1977.

[Hor75] Horn, B.K.P., "**Obtaining Shape from Shading Information**", in Winston (ed), The Psychology of Computer Vision, McGraw-Hill, New York, 1975.

[Low81] Lowe, D.G., Binford, T.O., "**The Interpretation of Three-Dimensional Structure from Image Curves**", Proc. 7th IJCAI, pp613-618, 1981.

[Per77] Perkins, W.A., "**Model-Based Vision System for Scenes Containing Multiple Parts**", Proc. 5th IJCAI, pp678-684, 1977.

[Shi75] Shirai, Y., "**Analyzing Intensity Arrays Using Knowledge About Scenes**", in Winston (ed), The Psychology of Computer Vision, ch 3, pp93-113, 1975.

[Shi78] Shirai, Y., "**Recognition of Real World Objects Using Edge Cue**", in Hanson and Riseman (eds), Computer Vision Systems, pp353-362, 1978.

[Ten73] Tenenbaum, J., "**On Locating Objects By Their Distinguishing Features In Multi-sensory Images**", SRI Technical Note #84, Sept 1973.

[Vam79] Vamos,T., Bathor,M., Mero,L., "**A Knowledge-Based Interactive Robot-Vision System**", Proc. 6th IJCAI, pp920-922, 1979.

[Wal75] Waltz, D., "**Understanding Line Drawings of Scenes with Shadows**", in Winston, The Psychology of Computer Vision, pp19-91, 1975.

[Woo77] Woodham, R., "**A Cooperative Algorithm for Determining Surface Orientation from a Single View**", Proc. 5th IJCAI, 1977.