TR-653
MCS76-23763

April 1978

BONE MARROW SMEAR IMAGE ANALYSIS

Peter Lemkin
Image Processing Unit
Division of Cancer Biology and Diagnosis
National Cancer Institute
National Institutes of Health
Bethesda, MD   20014

Computer Science Dept.
University of Maryland
College Park, MD   20742

COMPUTER SCIENCE
TECHNICAL REPORT SERIES

UNIVERSITY OF MARYLAND
COLLEGE PARK, MARYLAND
20742

PREFACE

This report presents an approach to the analysis of bone marrow smear images. Specifically, it presents a segmentation strategy for fields from normally aspirated Wright's stained bone marrow smear images of the type found in clinical laboratories. Problem domain knowledge is used in selecting the segmentation strategy. The choice and sequence of segmentation algorithms may be expressed as a decision tree. Although the decision tree is potentially open-ended, an analysis of an actual data base shows that most segmentation problems were of particular types. The algorithms implemented were restricted to a subset of the latter types.

The contents of the later Sections include: statement and scope of the problem with discussion of image and morphologic data bases (1 and 2); review of the literature on cell image segmentation (3); proposed segmentation stragegy (4); detailed algorithms for RBC removal (5); large object analysis (6); single cell articulation (7); small object analysis (8); conclusions and recommendations (9). An Appendix is included with a detailed description of the image data base. Other details covered in the author's dissertation include: the semantic data base (B); image processing software (C); specific details on selected algorithms which contribute to the dissertation (D, E, F); and the procedural description language developed for use with models of the type proposed in the dissertation (G).

i

TABLE OF CONTENTS

iii

1. PROBLEM STATEMENT AND SCOPE

1.1 The problem

This report presents a segmentation strategy for fields from normally aspirated Wright's stained bone marrow smear images of the type found in the clinical laboratories. Problem domain knowledge is used in selecting the segmentation strategy. The choice and sequence of segmentation algorithms may be expressed as a decision tree. Although the decision tree is potentially open-ended, an analysis of an actual data base shows that most segmentation problems were of particular types. The algorithms implemented were restricted to a subset of the latter types.

1.1.1 Selection of the problem

Bone marrow smears are sufficiently complex in their composition as to require structure-based segmentation rules. Normal marrow contains over 70 morphologically distinct elements recognizable as different classed by a trained hematologist. Cells present themselves in more complex ways in marrow (which retains some of the properties of tissue) than in peripheral blood. The latter consists primarily of isolated cells and involves far fewer cell types. Because of their tissue-like presentation, normal marrow smears contain many regions of clustered leukocytes that touch one another, which complicates any segmentation process [BesM73].

1

The successful segmentation of complex marrow scenes is a necessary step if tissue analysis is to be automated. This is because the types of analysis done in marrow are a subset of those required for tissue analysis.

1.1.2 Long term goals

The complete analysis of marrow images is beyond the scope of any single research project. The subproblem selected here has been restricted to that of segmenting (by isolating, labeling and measuring) a set of primitive objects found in marrow images.

The resulting segmentation can then be passed to a higher level system incorporating more hematology-specific knowledge in order to continue the segmentation and classification process. For example, once touching cells are separated, measurements can be made on the individual cells leading to their classification.

The primary data in a full marrow analysis system would consist of not just images, but problem domain knowledge as well. Such knowlege would enable the model to adapt to global changes in the slide or patient sample. Disease condition, staining shifts, cell density on the slide, artifact or dirt level could be taken into account in such a model. Such a modelling system might embed knowlege procedurally in such a way that patterns in the data could trigger the evaluation of global hypotheses about the model.

The segmentation methods suggested here were suggested during discussion with biologists of cells in a field of view (called a frame). Compartmentalization of objects is based on the occurrence of membrane boundaries (both visible and inferred).

During analysis, the model of the scene becomes more informed and is used to iterate toward an increasingly knowledgeable model of the frame. Basic assumptions about the number of cells in the frame; which cells are touching each other; existence of platelets; existence of hemoglobin-formed objects; etc. are used in filling in a basic model of the frame. Once this initial model is defined, questions can be asked as to whether the touching cells should and can be split; whether the hemoglobin formed objects are red blood cells (RBCs) or are cells in the developmental sequence of the RBCs.

The segmentation method is extensible in that it permits new knowledge about the bone marrow domain (normal and abnormal) to be added by its incorporation into the segmentation algorithms and non-image domain model. Such extensibility may make the model inconsistent. However, often biologists' concepts of particular morphological models are inconsistent, as evidenced by the variation in definition of particular cell classes.

The initial model is concerned with a single marrow slide field (represented by scans taken at multiple wavelengths). It is concerned only with normal marrow cells.

The restriction to normal marrow cells is not severe, because the majority of abnormal marrow slides consist of abnormal distributions of normal cells, or contain many cells which are morphologically indistinguishable from normal.

The model is extensible to multiple fields (and multiple slides) and to abnormal cells. The extension to multiple fields requires the ability to ignore information about previous fields below a given level of abstraction. The extension to abnormal cells requires additional information about particular instances of such cells or about the changes in cells which lead to abnormal cells. Much of this information is conjectured to be quite non-contradictory to that of the normal cell morphology, so that it may simply be added without destroying the consistency of the model. Other information has to do with numbers of cells and other types of population statistics which can also be added to take the effects of abnormal situations into account.

1.1.3 Restrictions on the problem domain - short term goals

The segmentation algorithms developed here isolate, label and measure some of the objects in the frame, including:

    (a) Red blood cells.
    (b) Complete nucleated cells (including separating
        touching cells).
            (b.1) Nucleus identification.
            (b.2) Cytoplasm identification.
            (b.3) Nucleolus identification.
    (c) Platelets.
    (d) Cells touching the frame (with the option of
        ignoring them).
    (e) Dirt and other artifacts.

(f) Ill-formed regions not belonging to any of (a) to (e).

The initial segmentation algorithm is insufficient to account for all pixels in all fields. Further segmentation algorithm enhancement will be required for use in a production envirnment. The initial segmentation algorithm goals are:

(a) RBCs and normoblasts are accounted for.
(b) Nuclei are correctly accounted for (cells may or may not be segmented).
(c) Cytoplasm is correctly accounted for in being assigned to the correct nucleus.
(d) Platelets are accounted for.
(e) Prominent nucleoli are accounted for.
(f) All pixels in the field are designated as belonging to artifacts, biological components, or background.

The result of frame segmentation is a set of connected component (CC) images for various cell components (including complete cells). Cellular objects and subobjects may be extracted using the CC images and features computed at that stage. The CC images to be created are: nuclei (CCnuc), cytoplasm (CCcyt), nucleoli (CCnucleoli), and platelets (CCplate).

1.1.4 Open-endedness of the problem

The open-endedness of the problem is due to the fact that the number of classes of objects is open-ended (in types of artifacts as well as in abnormal types of cells). In addition, the number of situations in which cells present themselves is open-ended.

The diminishing returns associated with infrequently

seen situations should determine what is implemented and what is not. This is the approach taken here, where those situations most commonly occurring were selected for implementation. A simple analysis of common types of situations was performed on the image data base, as discussed in Section 1.3.2. Failures in the experimental results can be explained most often as belonging to cases not treated in the current implementation.

1.2 Bone marrow domain

The marrow problem domain is introduced here in clinical terms. This discussion is intended to give the flavor of some of the specimen preparation problems and some of the reasons for the wide variation found in samples from the same slide and between slides.

1.2.1 Clinical use of bone marrow smears

Normally, a marrow smear is requested in a clinical context when other patient symptoms indicate that it is needed for confirmation of a diagnosis or for following the progress of treatment. The smear is prepared by taking part of an aspirated bone marrow sample and placing it in the center of a microscope slide. Another slide is placed on the top of this and pulled away in a parallel direction, thus spreading out the sample. The smeared out sample is then stained. Generally many smears are prepared from the same aspirated sample. Some are stained immediately and some are air dried and saved for later staining with special stains, none of which are usually absolutely required for slide diagnosis.

As viewed in white light with a 100X oil immersion objective (about 1000X magnification), the typical marrow smear frame contains 25 to 50 marrow-formed elements (biological objects of the size of a cell or platelet). Most of these are red blood cells, and a much smaller number are nucleated cells. There is a higher concentration of cells toward the center of

the smear with more cells touching each other in that region. The marrow elements can include:

(a) White blood cell clusters or isolated cells and their precursors.
(b) RBCs and their precursors.
(c) Abnormal cells.
(d) Artifacts of various types.
(e) Megakaryocytes and platelets.

1.2.2 Selection of fields for analysis

Generally, the clinical practice is to count and classify 200 nucleated marrow cells taken from different regions encountered on the slide. The first and last regions are on the periphery and the others are near the center of the smear. The hematologist normally scans the slide by moving it in a vertical raster pattern from the left to right side of the slide with the longest dimension of the slide being horizontal. The smear usually appears as an oval region in the center of the slide with the longest dimension of the oval aligned with the longest dimension of the slide.

The knowledge gained by looking at the first region is used as a normalizing factor in looking at the central regions. The last region, also being a peripheral region, acts as a check on the assumptions made in examining the others.

Thus, cells on the periphery of the slide tend to be more isolated and can serve as examples of class-typical cells. Similarly, cells in the center are subject to more clumping and distortion. The gross analysis of the slide is included in the

hematologist's report as comments on his overall impression of the slide.

1.2.3 Magnification and frame size

Bone marrow examination under the optical microscope with Wright's or Giemsa stain is normally performed at about 1000X magnification, obtained with a 100X oil immersion objective lens with numerical apertures of 1.2 to 1.4 giving a resolution of less than 0.5 micron. At this resolution, with cell sizes of 10 to 20 microns, an image frame size of 256x256 pixels (approximately 0.2 micron/scan pixel) yields about a 50 micron square field - adequate for viewing about 30 complete cells.

1.2.4 Staining as a selective contrast transform

In the context of marrow, a major reason for staining is to increase the contrast of the material. Normally the index of refraction of the background and cytoplasm is close to that of the nucleus so that without some form of contrast enhancement no structure can be seen. Little information can be extracted from unstained microscopic images without using methods such as phase or interference microscopy.

The Romanowsky stains (which include Wright's and Giemsa stains) consist of several constituent-specific staining materials which stain the DNA, histone, and RNA complex. The colors resulting from the staining are generally in the red and

blue/violet regions of the spectrum. A first approximation segmentation may sometimes be made based on color separation taking into account the changes in staining intensity and color between slides.

Wright's/Giemsa stain is the one most commonly used in hematological laboratories, and is the stain on which the largest body of standard hematologic knowledge is based. Other stains will not be discussed here.

Color filters (in the visible spectrum) are often used in order to enhance the contrast of stains for specific cell components. This will be discussed in more detail in Section 5.

1.2.5 Abnormal cells and abnormal situations

An abnormal cell is a cell not usually found in a healthy person, or alternatively, an undistorted cell which is not classifiable as normal. An abnormal cell could be some variant of an otherwise normal marrow cell, or a cell from outside the marrow. The handling of abnormal cells in a marrow modelling system can be accomplished by an extension of the model semantics to include their definition.

It is possible that a smear may be abnormal in the absence of morphologically observable abnormal cells. For example, a case of severe posthemorrhagic anemia presents an abnormal situation with an increased number of erythroblasts [LEL]. They are morphologically normal cells, and the relationship between different stages of maturation has been

shifted to the left (toward the precursor cells) in terms of maturation features.

An abnormal situation on a bone marrow smear may involve the observation of too many or too few objects of particular types in the smear. There are many correlations between such abnormal situations and disease states. Abnormal situations are much more common than abnormal cells.

The number of cells that must be found in a sample of a given size in order for a situation to be called abnormal depends on the type of situation suspected by the observer. It should be noted that an observer has different motives in doing an analysis of the smear depending on whether the analysis is from the point of view of the morphologist or the clinician.

1.3 Image data base

The image data base used for the initial design of segmentation algorithms for marrow consists of thirteen multispectral scanned images. These images were selected to present typical difficult situations involving clustered white blood cells normally found in marrow.

1.3.1 Selection of fields

The frames in the data base are Wright's stained normal bone marrow scans at white, 546 (green), 420 (blue) and 620 (red) nm (+/- 10 nm) wavelengths. They were scanned using a Zeiss Axiomat microscope at 800X magnification (100X oil objective lens NA 1.32, condenser lens NA 1.4, giving a magnification of 0.159 microns/pixel). The segmentation algorithms discussed in the dissertation did not use the red and white scans, but only the green and blue scans. selected to be Each frame was analyzed by itself, taking no information from the analysis of other frames into account. Figure 1.1 illustrates a 256x256 pixel square field taken from a larger marrow field photographed from the TV display during data acquisition.

The image data base, presented in Appendix A, shows the white, green, blue and labeled green image for each frame, as well as discussing peculiarities of each frame.

The frames in the data base were selected taking into
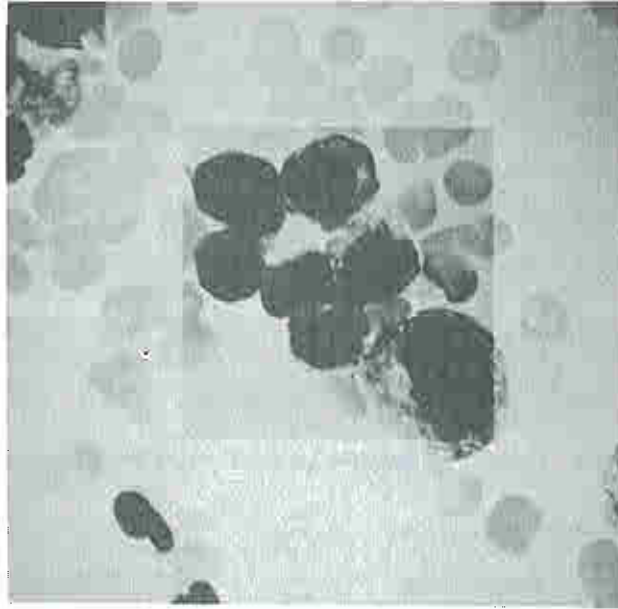
Figure 1.1 Sample of frame taken from larger context of marrow slide during data acquisition. The darkened region is a 256x256 pixel frame at 800X in green light.

account features of the interest from both the cytologic and image processing points of view. They are representative of the types of fields normally found. They do not include all of the types of cells and situations that can occur, but they do reasonably represent many of the major problems often seen. The selection criteria were as follows:

(a) Worst case and moderate case examples of cell types were used so as not to bias the image analysis algorithms in favor of "easy" cases.
(b) Cases were selected involving intact but touching cells of various types.
(c) Some smashed cells and isolated nuclei were included to represent some of the artifacts normally seen.
(d) Fields with varying degrees of rouleaux (RBC overlap) indicating degree of cell population density were selected.
(e) In many of the frames, cells were clustered and of varying gray levels so that simple strategies depending on isolated cells or on thresholding could not be used to perform the segmentation.

1.3.2 Frequency of particular data base configurations

In order to decide which aspects of the bone marrow segmentation process should be of primary importance, an informal analysis of the acquired data base was performed. The results of this analysis were then used to help decide which parts of the decision tree segmentation procedure to implement first.

The analysis is summarized in Table 1.1 below where the nine columns indicate the occurrence of a condition for the frame given by the row number. The nine conditions are:

(a) Distortion due to touching cells. The shape of the cytoplasm and possibly the nucleus is questionable because of the forces of deformation.
(b) Cytoplasm forced between touching cells.
(c) A lightened region occurring between touching cells.
(d) A darkened region occurring between touching cells.
(e) Lack of distinctive region between touching cells.
(f) Overlap occurring between RBCs (rouleaux pattern).
(g) Overlap of RBCs and leukocytes.
(h) Congregation of platelets near the edges of leukocytes.
(i) Wetting artifacts compressing cells near a clear region.

Frame	a	b	c	d	e	f	g	h	i
126-3	yes	no	no	no	yes	yes	no	no	no
126-4	yes	yes	no	no	no	no	no	no	yes
126-5	no	no	no	no	no	yes	no	yes	no
126-6	no	no	yes	no	yes	no	no	yes	no
126-7	yes	yes	yes	no	no	yes	no	yes	yes
126-8	yes	yes	yes	no	no	no	yes	no	no
126-9	yes	no	yes	no	no	yes	no	yes	yes
127-10	yes	no	yes	yes	no	no	no	no	no
127-11	no	no	no	no	no	yes	no	no	no
127-12	no	no	yes	yes	no	yes	no	no	no
127-13	yes	no	yes	no	no	yes	no	yes	no
127-14	yes	yes	yes	no	no	yes	no	no	no
127-15	yes	yes	yes	yes	no	no	no	yes	no

Table 1.1 Partial list of characteristics of frames in the image data base.

Note was also taken of the occurrence of various other conditions which can cause problems in distinguishing between cytoplasm and nuclear objects.

Some objects such as naked nuclei have no cytoplasm (126-5, 126-6, 127-14, 127-15). Cells such as normoblasts have very little observable cytoplasm (126-3, 126-7). This could conceivably be a problem in defining cytoplasm boundaries.

Cells such as eosinophils and basophils are so heavily granulated that the nucleus is obscured (126-5, 126-8). On the other hand cells such as mature neutrophils have a nucleus which is easily distinguishable from its cytoplasm since the latter is almost clear and of a much lower density than in other leukocytes (126-3, 126-4, 126-5, 126-6, 126-8, 126-9, 126-10, 127-11).

Nuclei at certain stages of maturation tend to be displaced in the cell to an edge of the cytoplasm (126-3, 126-4, 126-5, 126-6, 126-7, 126-8, 126-9, 127-11, 127-12, 127-13, 127-14, 127-15). This condition makes labeling of cytoplasm near the nucleus/cytoplasm intersection difficult.

2. APPROACH

The general approach used to segment marrow images is discussed in this Section. This approach has been implemented at a preliminary level; however, the discussion here encompasses more sophisticated levels of the model (cf. Section 2.1). The more complex levels of analysis require the use of explicit morphologic knowledge of marrow images. This type of knowledge is introduced in Section 2.2. Finally, the contributions of this work toward machine analysis of biological cell image analysis are discussed in Section 2.3.

2.1 Overall Method

The paradigm proposed here is to use the syntax and semantics of the problem domain to aid in the segmentation of images for which the model semantics hold. The discussion here will cover a proposed long term approach (which was not implemented) as well as a short term implementation.

2.1.1 Long term approach using the paradigm

Problem domain knowledge in both the visual and biological domains can be used to aid in the analysis of a frame or set of frames from a slide. This can be extended to include knowledge gained from looking at a collection of slides. One question immediately arises. How is such knowledge to be (1) organized and (2) used? The solution suggested here is to use the framework of a procedural

17

description system to embed the knowledge in such a way that several goals can be achieved. Other formalisms, such as decision trees or production systems could also be used to embed the structural knowledge. Segmentation of the image data uses knowledge about how particular biological materials present themselves. Information gained about the frame (or set of frames, or set of slides in the extended cases) should be used in analyzing the current frame. The model of the frame becomes more informed about the contents as the analysis progresses and is able to use this information to help in the segmentation.

Various procedural description systems have been developed and have been reviewed in [LemP75]. A procedural description system developed in part by the author, called PRDL [ShapB77] was used in this research.

Knowledge of the problem domain is coded in such a system as static facts or as theorems (coded as procedures). Static facts may be defined as binary relations or associative triples (using the PRDL notation of ATTRIBUTExOBJECT=VALUE). The following triples might describe some of the relations among objects found in a frame.

```
HAS x NUCLEATED-CELL = NUCLEUS,
HAS x CELL = CYTOPLASM,
INSIDE x CELL = NUCLEUS,
INSIDE x NUCLEUS = NUCLEOLUS,
INSIDE x CYTOPLASM = VACUOLES,
INSIDE x CYTOPLASM = GRANULES.
```

Using additional relations, one can contruct a static

model of the frame which resembles a fixed decision tree. Some of these include

```
HAS x FRAME = CELLS,
HAS x FRAME = ARTIFACT,
IS x CELL = NUCLEATED-CELL,
IS x CELL = RED-BLOOD-CELL,
IS x CELL = NAKED-NUCLEUS.
```

Flexibility is gained in the ability to manipulate such binary relations by using procedures (hence procedural description). These procedures embody theorems concerning relations between relations. A cell is defined by recognizing its component parts in a structured way such that the model's syntactic and semantic conditions are met. A very simple procedural model of a cell might be defined as

```
NUCLEATED-CELL(x) = IS-OBJECT(X) and
                    EXISTS-CYTOPLASM(x) and
                    EXISTS-NUCLEUS(X).
```

The composed procedures embody lower level semantic knowledge about cells at the morphologic image level (taking size, color, shape, texture, etc. into account).

The paradigm first tries to find objects as isolated cells. It then checks to see if some of these are actually touching cells. If this is the case, they are split, the model semantics being checked at each step. The algorithm must iterate, if necessary, in defining a more informed model.

For example, a procedure to test if a region is cytoplasm might invoke different procedures depending on what type of cell it is trying to model. Some cells have abundant

cytoplasm while others have very thin rims of cytoplasm which have a color and texture gradient going out from the nucleus. Different techniques are required in these two cases; the choice would be directed by what has been found so far in evaluating the model. This type of analytical decision is required over and over again in analyzing the complex types of situations found in segmenting marrow.

Also available in a procedural description system is the ability to invoke procedures based on the occurrence of Boolean expressions of patterns instantiated in the course of performing an analysis. This is where the real power of such a system resides. Special cases may be handled by defining pattern invoked procedures for them. These procedures are never evaluated unless the given situation occurs.

The driving procedure in the paradigm is to account for all of the regions of the frame. This procedure would be called ACCOUNTFOR in PRDL. Currently a collection of OS8 batch jobs implements parts of ACCOUNTFOR. ACCOUNTFOR controls the segmentation process and interacts with the morphologic data base, rationalizing what is found in terms of what is possible.

An alternative way of interrogating the system (via PRDL when implemented) would be in using a WHATIS(x) or FINDINSTANCE(y) procedure having previously run the ACCOUNTFOR(frame) procedure. Thus, given an object x or a desired cellular object y, the analyzed scene can be interrogated in a different way. If a trace were kept of the

analysis process then the reasons for the decisions could be extracted. This would permit the modification of the model manually to improve its performance when finding faulty reasoning.

2.1.2 Short term implementation

The short term implementation does not use a procedural description system. This was because of the lack of time to integrate the components of such a system. As a result, the application of marrow domain knowledge to the segmentation of marrow images was performed using the lower level components of such a system, with the integration of various algorithms done manually. The image processing system Buffer Memory MONitor System (BMON2) was used with segmentation procedures coded as either special BMON2 functions or as batch job sequences.

The image processing system hardware used consisted of an Axiomat microcope; the Real Time Picture Processor (RTPP) vidicon TV-image buffer memory; control desk hardware; and the BMON2 image processing software running on a 32K PDP8e computer (OS/8 FORTRAN II).

The segmentation algorithms were developed in the context of the BMON2 software system and are discussed throughout the remainder of this report. As mentioned in Section 1.1.3, those parts of the decision process most frequently occurring in marrow were implemented in a first approximation of the decision tree.

2.1.3 Why the approach was chosen

In general, the more one knows about a problem one is trying to analyze, the more successful that analysis will be. Thus, embedding problem domain knowledge in the segmentation process was felt to increase the likelihood of success of the segmentation process, especially when scenes were fairly complex and consisted of a wide variety of object types and situations.

Although an actual procedural description was not used, procedural knowledge was incorporated in the work being reported. This knowledge was embedded in fixed data structure algorithms. Such algorithms assumed the semantics of the operands on which they worked (for example, the RBC elimination algorithm assumed the existance of green and blue marrow image scans).

Procedural description methodology has existed for several years (see the review in [LemP75] and in [ShapB77]), but it has never been applied to analyzing bone marrow images. As discussed in Section 3, previous attempts at analyzing cell images were primarily concerned with non-touching isolated leukocytes often prepared using slide spinners (see [LemP75]). Other approaches did not attempt to reconcile a large morphologic and meta-morphologic data base with the low-level segmentation of images.

2.2 The morphologic data base

As was mentioned before, the biologist brings a large body of morphologic knowledge to bear in analyzing a marrow frame. This type of knowledge is discussed here to a limited extent. In addition, the acquisition of the morphologic data base is discused, and some examples are presented.

In general, morphologic knowledge may be thought of as syntactic and semantic rules that a biologist brings to bear on analyzing the marrow image. Such expert knowledge is available both from cytologist experts and source books. The cytologist informants were Drs. L. Lipkin, pathologist, and N. Markovic, hematologist. The source books were selected from the point of view of finding operational definitions for the morphology of marrow cells and situations. The best source used was the widely acclaimed work on blood by M. Bessis, "Living Blood Cells and Their Ultrastructure" [BesM73]. Other source books include [DavI74], [SilR70], [TalG57]. As is often pointed out in the literature, these sources may conflict.

As will be seen in Sections 4 through 8, problem domain semantics are useful in aiding the segmentation process. Some of the semantic characteristics of marrow smears mentioned here are used in the creation of the decision tree discussed in Section 4. Other problem domain semantics have not been used at this stage of development of the algorithm.

The following are some examples of morphologic data

base concepts that were used in the construction of the current
segmentation algorithm.

    (a) RBCs are the darkest objects at 419 nm (blue)
illumination because of their hemoglobin content
and the staining characteristics of Romanowsky
stains.

    (b) Nuclei are the darkest biologic objects visible at
540 nm (green) illumination because of the staining
characteristics of DNA with Romanowsky stains.

    (c) Touching nuclei and cytoplasm may sometimes appear
to have a lightened texture region between them due
to diffraction effects.

    (d) Most nucleated cells in a normally prepared bone
marrow smear do not overlap one another. RBCs which
do (rouleaux) indicate too dense a preparation
with respect to cell population.

    (e) Topological completeness of nucleated cells implies
that the nucleus is completely enclosed by the
cytoplasm.

    (f) The normoblast nucleus, having been extruded from
the normoblast, is not surrounded by cytoplasm.

The following are some additional concepts which were
derived from the same sources but are not used in the
construction of the current segmentation algorithm.

    (a) Polymorphonuclear nuclear lobes may appear as tear
drop shaped objects with continuous curvature at
the sharp points of the lobes.

    (b) Nuclear texture coarseness increases with age in
the granulocytic series.

    (c) Nucleoli are most often found in the more immature
cells.

    (d) The Golgi region is indicated by a light texture
region in the cytoplasm next to the nuclear
concavity in all of the stab type cells except the
poly.

    (e) Stain crystals occuring on the slide are very dense
and have sharp corners, unlike most isolated
biological objects at this magnification.

    (f) Depending on the stage in the maturation process,
the nucleus is located off center of the cytoplasm.

    (g) Cells such as normoblasts have very little
observable cytoplasm.

    (h) Mature eosinophils and basophils are so heavily
granulated that the nucleus is obscured.

    (i) Cells such as mature polys have a nucleus easily
distinguishable from its cytoplasm since the latter

is almost clear.

2.3 Contributions

Several contributions are discussed here which indicate significant progress towards machine analysis of complex histologic images. Such images cannot be pre-processed by techniques such as slide spinning because of local distortions and disruption of cell groups, or the use of special stains not commonly used.

2.3.1 Problem domain knowledge driven segmentation

A morphologic data base was compiled in order to have a basis for using the semantics of the problem domain in the segmentation process. The data base is oriented towards machine analysis of marrow images. Only a part of this data base was used in the present implementation, but the data base will be used in further work on the segmentation and analysis of marrow images.

The red blood cell (RBC) elimination algorithm discussed in Section 5 using the green-blue positive difference transform is useful in eliminating formed objects which contain variable amounts of hemoglobin as in commonly found in marrow smears. Simple binary mask differencing as used by [YouI75] and the whitening transform [BacJ76] are inadequate for complex marrow images because of difficulties in mask generation and RBC clustering problems (in the whitening transform) due to RBC overlap.

The cytoplasm mask generation algorithm discussed in Section 7.2.1 uses information about the formed objects as well as multispectral images in forming a mask suitable for a wide range of cytoplasm textures and densities. The algorithm differs from that of [YouI75] and others in smoothing the cytoplasm mask to handle fragmentation due to overlap of background/cytoplasm and cytoplasm/nucleus density ranges.

The splitting algorithm for touching cell nuclei is discussed in Section 6.
It uses a heuristic
boundary crack finder on the nucleus and checks the texture in the crack region to verify the fact that the nuclei are touching. The algorithm was developed independently of that of [BrenJ77]. The major differences between the two algorithms are in the pre-filtering, in the use of a heuristic function, and in the use of texture region verification.

The cytoplasm isolation algorithm, which prepares the cytoplasm-nuclei pre-propagation image for propagation of the nuclei into the surrounding cytoplasm, is discussed in Section 7.2.2It uses a heuristic function to evaluate possible lightened regions nearly perpendicular to normals between pairs of nuclei to decide if and where to cut cytoplasm regions touching two nuclei. The optimal angle at which to split the region is computed as a still better estimate. This initial approach to cytoplasm separation, although finding the optimal points at which to isolate the cytoplasm, approximates the isolation process using a straight line, whereas tracking the

edge to perform the split would be more desirable. The latter approach was not taken at this time because of difficulties due to the variability of the presentation of the cytoplasm material.

The boundary trace transform (BTT) is useful for strengthening weak boundaries using BTT(T). BTT(T) could also be used for detecting the nuclear boundary in heavily granulated cells such as basophils. It is useful for finding large textured objects such as nucleoli inside cells using BTT(F).

2.3.2 Tools produced

Various tools were produced in the course of doing the dissertation which were useful in both implementing the algorithms and which also are contributions to the state of the art in image processing systems in general.

BMON2 (Buffer memory MONitor system) is the buffer memory image processing system used to implement the segmentation algorithms described in the dissertation as well as other algorithms and data acquisition software. The BMON2 system is described in [LemP77a]. It was developed using PROC10 as a model image processing system after the buffer memory hardware was completed.

The RTPP (Real Time Picture Processor) hardware was designed and constructed for use in processing large numbers of

images using complex image processing functions. The hardware
was used to perform the image acquisition and processing for
the dissertation. The RTPP is discussed in and ([CarmG74],
[LemP74], [LemP76e], [LemP77b]).

The PRDL (PRocedural Description Language) language was
designed (and partly implemented) for eventual use in high
level modelling of cell images. PRDL is documented in
[ShapB77]. PRDL was not used in developing the actual
segmentation model.

The PROC10 (PDP10 image processing system) was
developed in order to perform image processing operations
before the RTPP was finished. PROC10 was not used in the final
work because of the high cost of doing image processing on the
PDP10 as well as the fact that enough of the RTPP (image buffer
memories) was completed to make the RTPP a viable alternative.

3. PRIOR WORK RELATED TO MARROW SMEAR IMAGE SEGMENTATION

A review of prior work related to the segmentation of bone marrow cell images is presented in this section. A technical report by this author, "A Literature Survey of the Technological Basis for Automated Cytology", [LemP75], was written covering selected areas of automatic cell analysis systems up to the end of 1974. The areas covered in the report include bone marrow, peripheral blood (both leukocytes and red cells), and cervical cells. Sample preparation is discussed including how preparation affects the ease of performing measurements on the samples. Parts of the report relevant to cell image segmentation are referenced in this section. Additional material written after 1974 on topics related to marrow cell segmentation is also covered.

In [LemP74], the literature on both the image processing and flow systems was reviewed, with emphasis on the image processing approach. The survey reviewed both statistical and non-statistical approaches with emphasis on the latter. A wide range of picture operations and features were discussed, including shape, texture, color, and density features, and various segmentation algorithms were also treated. Image processing systems and special hardware processors were also reviewed. The application of problem solving methods to scene analysis was discussed, with emphasis on procedural description, although theorem-proving methods and mixed structural-statistical methods were also considered.

30

3.1 Segmentation of cell images

The field of biological image processing tends to group image processing segmentation techniques into two areas as determined by the complexity of the images. This complexity in turn is determined both by the type of material being analyzed and by the method of sample preparation. Samples prepared using so-called slide-spinners or ultrasonic disaggregation tend to scatter clumps of cells so that declustering of touching cells is not normally necessary. Unfortunately, such sample preparation methods often bias the sample by shifting the heavier cells to different regions of the slide as well as causing other artifacts.

Other specimen preparation techniques involve the use of special stains for which color segmentation techniques often work quite well. Unfortunately, in using these stains one cannot take advantage of the wide body of literature on cells described using Romanowsky stains. For many types of cells, these are the only descriptions available.

In recognition of this difference in approach within the field of automated cytology, in the proceedings of the past several Engineering Foundation Conferences on Automated Cytology, sessions were split into three main areas of research: (1) specimen processing and cytochemistry, (2) image processing analysis, (3) flow systems analysis ([MayB74], [MayB76], [MayB77]).

3.1.1 Segmentation

Work has been done in various areas of cell image segmentation including RBCs, WBCs, and bone marrow. Other related areas include segmentation of cervical cells and chromosomes. Some of the work which is particularly relevant to the segmentation of bone marrow will be reviewed in later parts of this section.

Techniques using density analysis, multi-spectral color images, boundary tracing and analysis, and region growing are discussed. Several of the papers discussed here ([AggR77], [AusH77], [BacJ77], [BrenJ77], [CahR77], [HolJ77], [MuiJ77], [PoulR77]) were presented at the Fifth Engineering Foundation Conference on Automated Cytology (January 1976, Pensacola Beach, Fla.). Over 50 papers from this conference appear in the July 1977 issue of the Journal of Histochemistry and Cytochemistry [MayB77].

3.1.1.1 Density analysis algorithms

Density analysis of gray scale image data has been used since the early work on cell image analysis to segment cell images. Under optimal slide preparation conditions this approach is a very effective procedure. As has been pointed out in Section 3.1, such conditions are not always present in routine bone marrow slides. Some of the work using density analysis for segmentation is reviewed here.

Prewitt and Mendelsohn [PrewJ66] discuss using analysis of the density histogram to select thresholds for slicing manually selected images of single peripheral blood cell leukocytes into background, cytoplasm and nuclear regions. They suggest selecting the minima between region component peak density values as the threshold values.

Taylor et al. [TayJ75] discuss automatic nucleus finding routines using density histogram analysis to find suboptimal thresholds for cervical cells. They point out that the problem with analyzing the histograms for these cells is that the histograms are extremely variable and complex and the choice of threshold is not always obvious. They find the most prominent minimum of the histogram using an adaptive algorithm as follows:

[1] The histogram is first smoothed to remove most of the local minima due to noise.

[2] The figure of merit for a local minimum is chosen to be the lesser of the areas of the two peaks on either side of the dip above the value of the histogram at the dip.

[3] If there is more than one dip within a search range, then the smoothing process is repeated until there is either only one dip left within the search range or the ratio of the largest figure of merit to the next largest is greater than a preselected value.

[4] The threshold value selected is the value of the minimum having the largest figure of merit.

They then propose analyzing three types of density histograms: the normal image density histogram, its logarithm, and density transition (adjacent pixels) histograms in the UP, DIAG (diagonal), and DN (down) directions. The last three

histograms (UP, DIAG, DN) are further analyzed to shift the threshold selected either up or down with varying success. Two functions f and g are then computed:

$$f(k) = \log(DIAG(k)+1)-5*(UP(k)-DN(k))/(UP(k)+DN(k)+1),$$

$$g(k) = \log(DIAG(k)+1)+5*(UP(k)-DN(k))/(UP(k)+DN(k)+1),$$

The second term in f and g tends to shift the dip at the nuclear threshold to a greater or lesser value. The minimum finding algorithm specified previously is then used to find the minimum threshold.

For all three types of histograms, the search range limits for the algorithm are restricted to make the nucleus to cell area ratio fall into empirically found limits as a function of cell size. If no suitable minimum is found in this region, then the search limits are expanded.

All these algorithms showed a bias towards areas smaller than the visually determined values, which they point out is correctable. Part of the problem is that it is difficult to say where the cervical cell nucleus ends, especially in a digitized image, when the image has low contrast.

They further note that the major causes of thresholding problems can usually be traced to unusual input data which confuse the algorithm or cannot be analyzed by thresholding at all. They give several examples, such as cytoplasm having regions darker than parts of the nucleus which results in overlapping histograms for the nucleus and cytoplasm regions; or the presence of nucleoli or other small dark objects

(stained with Papanicolaou stain) inside the nucleus which can be mistaken for the nucleus (this can be compounded by the relatively light nucleus and the virtual absence of identifiable cytoplasm).

3.1.1.2 Color analysis algorithms

Color has often been used in cell image analysis both for classification ([BacJ69], [IngM68], [YouI72], [BrenJ74], [CheG74] and others) and for segmentation ([YouI75] and others). The motivation for this is that stains are taken up differently by different components of different cells. In fact, the common Romanowsky stains are actually multicomponent stains where different colors are taken up by different parts and types of cells.

Color information, then, may be used for contrast enhancement. It plays an important role in both segmentation and identification schemes. Specifically, with Romanowsky stains:

Cell cytoplasm and especially cell nuclei are enhanced at 550-540 nm wavelength (green) of light.

Red blood cells are enhanced at the 419 nm wavelength (the blue "Soret" band) because of the hemoglobin.

Young and Paskowitz [YouI75] discuss this and give a table (reproduced below) for the three wavelengths (570, 530 and 420 nm (yellow $Y(x,y)$, green $G(x,y)$, and blue $B(x,y)$)). Some of the problems associated with using this technique are that meaningful thresholds must have been established that are

reproducible and distinct. With very granular cytoplasm textures and overlapping RBCs in marrow, this latter requirement is difficult to achieve.

Illumination	Nucleus	Cytoplasm	RBC	Background
570 nm	dark	gray	gray	light
530 nm	dark	dark	dark	light
420 nm	light	light	dark	light

They then derive some simple Boolean equations for cytoplasm points $C(x,y)$, nuclear points $N(x,y)$, RBC points $E(x,y)$, and background points $B(x,y)$.

$$N = Y$$
$$\text{or } YB'$$

$$C = GY'B'$$

$$E = B$$

$$W = N + C$$
$$= GB' + Y$$
$$\text{or } GB'$$

Note that N, C and E should be disjoint since they represent different sets of points in the image. i.e.

$$NC = Y[GY'B'] = 0$$

$$EC = B[GYB'] = 0$$

$$NE = YB'[B] = 0$$

Although in principle this technique would seem to be adequate, it will be seen in Section 5 that simple thresholds cannot be used reliably in marrow smears because of the variation in frame intensity and the wider range of cell types which occur in marrow than in the peripheral blood.

Tycko et al. [TycD76] report on the automatic

classification of peripheral blood leukocytes using a new three-component cytochemical stain with relatively narrow absorption bands at 460 (fluorescein), 540 (pyronine y) and 640 (methyl green) nm. The classification strategy makes use of the properties of the stained leukocytes as follows:

[1] Detect the presence of a leukocyte by looking for a dark object under red light.

[2] Examine the object under blue light. If it is dark, it is either a basophil or an eosinophil; goto [3], otherwise, goto [4].

[3] Examine the object under green light. If it is dark, it is a basophil, otherwise it is an eosinophil.

[4] Examine the object under green light. If it has small area and low contrast, it is a neutrophil. If it has small area and regions of high contrast, it is a small lymphocyte. If it has large area and low contrast, it is a large lymphocyte, monocyte or other.

[5] In the last case of [4], use morphological features (nuclear shape and texture from the red images, cytoplasm shape and texture from the green images).

This strategy should separate the leukocytes into five classes. The remainder of the paper concentrates on steps [1] to [4] using features derived from the three colored image histograms.

Another technique, developed by Bacus [BacJ76], called the "whitening transform", maps two colored images into a "colored" image C and a "density" image D. This is done by computing the principal eigenvectors and eigenvalues for each image to maximize separation. Bacus makes two assumptions: first, that the principal eigenvector lies along a 45 degree line in the original bivariate space, and second, that the variances in either direction are relatively constant. He use yellow (YW) and blue (BL) images for the two colors. The

transform may be expressed as:

$$C = K1(BL-YW)+K2,$$

$$D = K3(YW+BL)+K4.$$

Bacus reports good results in applying the transform to segment leukocyte nucleus, leukocyte cytoplasm, red cells and background from spinner prepared peripheral blood smears. However, although the WT works fairly well for spinner prepared perihperal blood smears (with cells being separated), similar results would not be expected in marrow aspirate smears because of clumping of RBCs, overlap of RBCs and WBCs, and the occurrence of other morphologic stages of cells in the RBC series.

Mui and Bacus [MuiJ77] classified peripheral blood neutrophils using multispectral images as either being band or segmented neutrophils. The data base consisted of 378 neutrophils including unambiguous bands and segs as well as neutrophils with overlapping nuclear fragments. Each cell was scanned to 50 square pixels corresponding to 16.5 square microns in blue (Wratten #44) and yellow (Wratten #15). The RBCs and background were eliminated by using the whitening transform (WT) with analysis of the WT clusters and using spatial information from the image taking the number of clusters into account. The remaining image cytoplasm boundaries were then shrunk by two pixels and then expanded by two pixels to eliminate RBC pixels which were misclassified as cytoplasm. This paper is discussed in more detail later in the subsection

on region growing.

Aus [AusH77] discusses the use of multispectral color transforms in Pappenheim stained marrow smears. Using the fact that Pappenheim stains the granules in a more pronounced way than Wright's stain, scans are taken at 620, 580 and 420 nm. Aus computes bi-variate density histograms of the (620,320) nm scans and tri-variate histograms of the three colors. Using a 3-D coordinate rotation transformation and projection onto a 2D surface presents new clusters not visible in the uni- or bi-variate cases. Specifically, granulocyte cytoplasm was separated using such a transformation. Aus suggests that using additional spectral scans and transformations for specific cell types, additional components of other cells can be isolated. The cost of this approach in doing the additional scans and transformations might be prohibitive in general. A further problem is in the difficultly of using the Pappenheim staining procedure, which is not as simple as the Wright's procedure and is not used as routinely.

Aggarwal and Bacus [AggR77] report using color filters to match the spectral characteristics of Papanicolaou stained cervical cytology smears. Three filters at 569, 528 and 485 nm were used. The images contain white blood cells as well as epithelial cells (which are those to be detected).

First the nuclei are extracted using the 569 nm filter (maximum nuclear contrast). Then nuclear features are computed in an attempt to eliminate the WBCs. The cytoplasm is then

extracted using clusters in the 2D density plot of the (528,569) nm distribution. The clusters are found using a ceiling-lowering algorithm which grows clusters from the peak points of the histogram based on 8-neighbor connectivity. If a pixel is connected to more than one cluster, it is assigned to the cluster with the nearest center of mass. If a pixel cannot be assigned to an existing cluster, a new cluster is started. The algorithm stops when all pixels have been assigned to clusters. Several clusters for cytoplasm may result possibly causing problems. Therefore, the cytoplasm clusters were chosen manually.

Finally, the nuclei remaining are coalesced with the surrounding cytoplasm using a radial line algorithm. A number of radial lines are constructed from the center of each nucleus. The transition point between the background and cytoplasm is marked along each radial line. The marked pixels corresponding to the cytoplasm/background boundary are then joined by straight lines. Considering the complexity and variation of the scenes, the success rate appeared to be fairly good. Problems in nucleus and cytoplasm detection are accounted for by the wide variation found in the nucleus and cytoplasm densities in this material.

3.1.1.3 Boundary analysis algorithms

Boundary analysis of cell nucleus and cytoplasm boundaries has been used previously for features in peripheral blood cell classification [BrenJ74] as well as in the analysis

of chromosomes ([LedR66], [KliA71]).

Gallus et al. [GalG74] propose an algorithm for analyzing the curvature of a chain coded chromosome boundary in order to find the vertices of the chromosome arms and the centromeric constriction regions. The algorithm uses chain codes smoothed over both 3 and 6 pixel windows (the latter removing more detail but preserving the main structures). The peaks are then analyzed as to being concavities or convexities and a decision tree recognition procedure is applied to the data to force the fitting of a chromosome model (consisting of various subtypes) to the data. They analyzed 1820 unbanded and 920 banded chromosomes. The procedure had a total error rate of 4.9% on the unbanded and 8.2% on the banded chromosomes.

Mendelsohn et al. [MenM74] propose a topologic approach similar to that of [GalG74]. Their algorithm computes the chain code of the boundary and then looks for extrema in the chain code difference. The centromere is verified by showing that the region contains the shortest path across the chromosome. Furthermore, because the two chromatids must have equal DNA, the densitometric measures of the two chromatids are used to verify the centromere region selection.

Brenner et al. [BrenJ74] isolated RBCs that touch leukocytes in peripheral blood smears using a geometric technique. The RBCs formed concavities with the leukocyte edges. Their algorithm detected points of maximal concavity along the boundary and replaced pairs of these points with

straight line segments. They reported greater than 90% success on those cells which required boundary correction.

Brenner et al. [BrenJ77] discuss a boundary tracing cell segmentation algorithm used for the analysis of routine bone marrow smears from acute lymphoblastic leukemia (ALL) patients. The object of their study is to subclassify patients with "classical" varieties of ALL on the basis of lymphoblast cell morphology. Such a subclassification might be useful for prognosis and therapy.

The images are 20 microns square (0.2 micron/pixel) with 64 gray levels scanned at 600 nm (orange), 490 (blue-green), and 420 (RBC blue reference image). The preparations were Wright's stained aspirated bone marrow smears. The lymphoblast to be measured was put in the center of the field with possibly touching cells adjacent to it. Brenner notes

> "In contrast to relatively 'clear' peripheral blood smear images, those from routine bone marrow preparations are often confusing in terms of segmentation. Red cells and other white cells are frequently touching the cell of interest and there are often defects in the cell boundary itself which must be recognized and dealt with."

He goes on,

> "The most difficult situation in a bone marrow smear is that which occurs when a desired cell is surrounded by a cluster of touching [leukocyte] cells."

In dealing with touching leukocytes, Brenner assumes that the cell of interest is "roughly centered" in the frame in order to start off the segmentation algorithm. In any automatic system

looking at several complete cells in frame at a time, such an assumption is not always valid.

Brenner's segmentation algorithm (denoted [S] here for the purpose of discussion) is best understood if it is explained in three iterations (called here [S.1], [S.2], [S.3] for clarity) with each iteration handling additional cases.

Algorithm [S.1]
Basic isodensity contour following algorithm
----

[1] Compute the histogram of the image and find the cytoplasm threshold Tcyt.

[2] Search outward from the center along a radial line until a significantly large area below the cytoplasm threshold Tcyt is found.

[3] At this "edge point", trace the boundary points nearest in gray value to level Tcyt in the clockwise direction.

[3.1] The (n+1)st point is found by searching clockwise the 8 nearest neighbors of n starting at (but ignoring) the backward direction and selecting the first point above threshold. This point becomes the new n.

[3.2] If point (n+1) is the sames as point (n), then a small loop has occured. The procedure then is to set point n below threshold Tcyt and to start again at point (n-1). When iterated, this step eliminates small loops and thus reduces the boundary noise.

In general, the cytoplasm threshold may not be easily found if several different cells types are in the field. For Brenner's problem, the majority of the cells in the image were lymphoblasts so that the cytoplasm was homogeneous.

Unfortunately, the algorithm also tracks vacuoles (which make small indentations) and debris on the cell surface

(which makes small protuberances). Protuberances are also characteristic of "smashed" cells and are quite common in poorly prepared smears.

The following boundary repair algorithm [R] eliminates these two types of surface artifacts. The main requirement for this repair algorithm to work is that the cell is relatively round or elliptical.

Algorithm [R]
Boundary repair for vacuoles and small projections
----
[1] The boundary chain code is first smoothed and then its curvature is analyzed.

[2] Vacuoles "appear as a region of sharp negative curvature (concavity) between two sharp positive [convex] regions – the breaks in the gently curved cell boundary at the edges of the vacuole. These breaks can be recognized and joined to repair the boundary. A similar procedure detects and repairs the boundary at the projection which is detected on the curvature plot as an inverted 'vacuole' pattern."

In the case of touching WBCs, the boundary traces the outside of the region, missing the region where the cells touch. As WBCs can not be removed by spectral subtraction another technique is used. Brenner analyzes the boundary curvature to find notches, which are then used in pairs to recognize touching cell regions.

Algorithm [S.1] is modified so that after the boundary is initially tracked, the notches in the boundary are identified by analysis of the chain code. The touching cell notches are different from the cell-edge and vacuole-projection curvature differences.

Algorithm [S.2]
Touching cell splitting by notch matching
----

[1] The cell edge notches are eliminated by proximity to the frame edge. The vacuoles (projections) are eliminated by recognizing by short runs of (CCV,CCVX,CCV) or (CVX,CCV,CVX) patterns.

[2] The remaining notches are paired if they point at each other with some tolerance "Pointing directions are established for each notch as the bisector of the angle between the tangents of the boundary at the indentations. These tangents are found by approaching the indentation from the two possible directions."

[3] Then the boundary is retraced with notch pairs being joined by straight lines.

Finally, the case occurs where opposing notches are not detected because a number of cells touch in such a way as to form an untraced central region. A modification of [S.2] is to back up the boundary trace to a notch (not a vacuole or projection) and then trace the central region to generate additional notches to be used for touching cell separation.

Algorithm [S.3]
Touching cell splitting by notch expansion
----

[1] First a search is made of the boundary for a significantly large region with no notches.

[2] Then, starting at the notch, a search is made in the pointing direction (of the notch) for a significantly large region below threshold Tcyt but still within the original boundary.

[3] The region is then traced until the notch point is found.

[4] The new boundary is then analyzed to see if the notches added by the central region line up with the original notch that was expanded. If one does, then boundary repair is effected.

[5] The process is repeated on the new boundary which incorporates the new notches. It terminates when all notch pairs have been removed.

If no enclosed light areas exist for an unpaired notch, a more stringent test is performed to see if it might be a misidentified vacuole. It still remains a notch after the test, then it is ignored (as a variation of the shape of the object). Brenner also notes,

"Complications arise when separate legitimate features occur in close proximity. By a judicious choice of the sequence used in feature identification it was possible to circumvent most, if not all, of these problems. For example, the 'peninsula' of cell material between two neraby vacuoles may look like a projection. If projections are repaired first, then on the next boundary trace the two adjacent vacuoles will have been reduced to one large one. The large vacuole is thereby detected and repaired on the next pass."

In general, however, bone marrow analysis including a wide range of cells cannot be so conveniently adjusted.

Brenner notes that other complications occur which sometimes can be flagged and handled as special cases [no details are presented] by the program while others are handled by the operator. About 10% of the cases processed required manual intervention; about 50% of those were for redefining Tcyt, while the others were geometric problems. Finally, Brenner notes:

"The general boundary trace and repair methodology should be useful for isolating any (possibly touching) approximately round or elliptical objects against a uniform background. This supposes that the texture of the objects [i.e. cytoplasm] is low enough to permitt reasonably smooth boundaries to be found in the initial isodensity contour trace."

In the general marrow problem, some of these assumptions are not true: cells are not always round or elliptical; the

background is not always uniform; the texture of the cell boundaries varies considerably between different cell types; and Tcyt is not uniform for different cell types.

Bowie and Young [BowJ77] discuss a technique for finding the lobes of band and seg neutrophil nuclei shapes. The technique analyzes a smoothed boundary chain code to find points of strong negative curvature on the boundary corresponding to junctions of lobes with the body of the nucleus. A matrix of interpoint distances is computed after the peaks are found. Peaks with the smallest interpoint distance are singled out. In addition, this distance must be less than some threshold and the angles of the curves must point to each other in some sense. After the set of crossover points is found, the objects are cut off the body of the nucleus. This process is iterated until no new objects exist which are to be cut off. If no new objects to be cut exist on a cut body, then that body is defined to be a lobe. Area and (P*P/A) are computed for each of the lobes as a local shape measure.

The lobe segmentation algorithm was applied to a previously reported data base of neutrophils with 64% agreement between averaged observer counts and machine counts. They point out that the disagreement correlated with disagreement among the human observers. The problem was most often due to faulty selection of the fixed nuclear threshold used in the original boundary trace procedure. Such a fixed threshold cannot take local variations of the boundary threshold into account. Such variations might be due to optical shading as well as to the

general increase of threshold near a clump of cells.

Lee [LeeE77] discusses some shape measures for leukocyte nuclear boundaries which might be useful in describing band nuclear shape. He defines an "indented" feature Ui based on the outside tangential angle T2 and inside tangential angle T1 of a band form nucleus:

$$Ui = 1 - Ri*MAX(T1,T2)$$
or
$$Ui' = 1 - Ri'*MIN(T1,T2)$$
or
$$Ui'' = 1 - Ri''*T1*T2$$

where Ri, Ri', and Ri'' are normalization constants set to 1/180, 1/180, 1/(180**2) degrees respectively in order to normalize the values to [0:1].

Lee then defines two other functions of Ui for slightly indented (Usi) and deeply indented (Udi) shapes.

$$Usi = Ui^{1/2},$$

$$Udi = Ui^{2}.$$

3.1.1.4 Region growing algorithms

Preston, Ingram and Norgren ([IngM69], [PresK71], [PresK72]) used the properties of the Golay transforms [GolM69] for extraction of cell features for leukocyte analysis. The technique expands and shrinks regions according to properties of their geometric shapes.

As was mentioned previously, Mui and Bacus [MuiJ77], used shrinking followed by expansion (by two pixels in each

case) in order to eliminate RBC pixels which were misclassified as cytoplasm. They also use a nuclear fragmentation algorithm as the basis for determining the neck region of a band or segmented neutrophil by increasing the threshold with successive segmentations to the point at which the nuclear lobes split. They compute seven features to be used in a combined logical/statistical classifier with the first feature (X1) being whether or not the number of unconnected nucleus fragments is greater than 1.

The cell data base was split approximately in half into training and test groups. Cells in the test group were divided into two subgroups, consisting of those with and without overlapping nuclei. The classifier operated as follows:

```
If (X1 > 1)   {i.e. number of nucleus fragments > 1}
    Then Class <-- "segmented neutrophil"
    Else Class <-- StatisticalClassifier[X2:X7];
```

The classifier worked fairly well on group 1 but not so well on group 2. However, group 2 contains the more difficult cases. They also noted that in a previous study (Bacus's dissertation) some of the segs that were misclassified as bands because of faulty scene segmentation were now correctly classified (using the new shrinking nuclear fragmentation features related to the analysis of the neck region).

Cahn et al. [CahR77] proposed an algorithm for determining the optimal threshold for segmenting the touching cytoplasm of cervical smear cells. The algorithm starts by setting a rough threshold at which to contour trace all objects

above threshold. The perimeters and boundaries are saved. A threshold is rejected if (a) any object is greater than some threshold, or (b) the number of objects is greater than another threshold, or (c) no regions are found. If (a) or (b) is true increment the threshold and try again. If (c) is true, stop, as the frame is void, otherwise, eliminate objects having less than some small area. Finally, the largest remaining object is chosen and its perimeter saved.

The remaining object is tested to see whether it is accepted. If it is not, then the threshold is incremented by 1 and the algorithm is tried again. It is accepted (a) if the current perimeter value is within 20% of the value of the previous threhold value, or (b) if the number of objects (before sizing) is greater than that at the previous threshold (using the previous value), or (c) if the threshold has been increased past some upper limit (in which case the segmentation fails).

After the optimal threshold has been selected for the largest cell, it must be segmented into nucleus and cytoplasm. This is done using an isodensity clustering algorithm which although assuming two clusters (nucleus and cytoplasm) can be expanded to three when the cytoplasm is folded on itself.

The correctness of the segmentation procedure was checked both visually and by comparison with recognition rates obtained by manually segmenting the images into 2-class and 13-class recognition problems. They point out that some of the

problems with the algorithm result from the failure of the contour tracing algorithm to isolate individual cells.

In another paper from the same study, Poulson et al. [PoulR76] discusses segmenting cervical smears using both manual and automatic methods (discussed in [CahR77] above). The manual segmentation consisted of selecting and saving manually set thresholds for the image cytoplasm and nucleus which were later compared with those produced by the algorithm. They point out that for some cells which were unsegmentable by threshold alone, parts of the cytoplasm are denser than parts of the nucleus.

3.2 Statistics syntax and semantics

The use of semantics in cell image analysis was proposed in 1966 by Lipkin, Watt and Kirsch [LipL66]. This paper proposed a sophisticated structural analysis of microscope images based on embedding morphologic knowledge in an analysis system which coded rules in the form of a BNF grammar. The system was to have linguistic input and interaction with a cytologist talking to the system in English.

Others have been more successful in less ambitious applications of semantics to cell image processing. In particular, chromosome boundaries were analyzed using BNF grammars to code the rules of formation of chromosome boundaries ([LedR66], [FuK74]). Shapiro [ShapB78] has used similar techniques using circle transform based features and a BNF grammar to parse RNA micrograph boundaries.

3.2.1 Hybrid statistical-structural methods

The distinction between purely statistical and purely linguistic analysis is discussed in [KanL72] and reviewed in Section 2.6 of [LemP75]. Combined statistical-structural methods are also discussed.

Several papers have been written suggesting the use of hybrid syntactic/statistical methods in constructing classifiers for biological cell images. Bartels et al. [BartP76a] discuss some of the differences between statistical

and syntactic classifiers. Statistical classifiers are highly efficient compared to syntactic classifiers. A statistical classifier maps the unknown sample into a continuous pattern space which is then divided into a fixed set of classes. Syntactic or descriptive classifiers work on an 'open' set of pattern classes. They ascertain which features are present in a given sample and recognize the sample on the basis of the set of features. They map into a discrete pattern space which is easily extendible to allow for additional pattern classes. This easy extensibility is not necessarily true for statistical classifiers. A new pattern will be assigned by such a classifier to one of the given classes based on maximum likelihood or minimum distance. In the syntactic classifier, when a syntactic description is derived which does not fit any existing class description, this condition may be flagged and additional classe descriptions defined if necessary, or old descriptions modified to include the new pattern.

Bartels et al. suggest that a hybrid system employing both methods is probably the optimal solution in analyzing complex histologic tissues. Such a system would employ a descriptive classifier to establish the identity of a tissue and and then use a statistical classifier to establish normality against a statistical data base.

3.2.2 Procedural description

As is evidenced by most of the literature in automated cytology (see the [LemP75] and [PresK76] surveys), the primary

concern in automating cell image processing has been in feature extraction and the use of various types of discriminant analysis for cell classification. However, since 1974 there has been increased interest in mixed procedural-statistical approaches as evidenced by [BacJ77], [TycD76], [BrenJ77], and [ShapB74], [LipL76].

3.2.2.1 Low level semantic image analysis systems

Holmquist et al. [HolJ77] point out that except for a few systems such as TICAS [WiedG70] and SCANNIT [BrenJ76] constructed for research purposes and oriented toward cell classification by feature discriminant analysis, no specific hardware (including on-line microscopy) and software are readily avaiable. They mention that the PEEP-DECIDE-GRAPH system (presented at the 1974 Fourth Automated Cytology Conference) is a step in the direction of a complete software system. It had the ability to to interactively process images and various other types of data structures, but had the problem of running on a large PDP10 computer with no microscope connected on-line.

They present the system design for an integrated microscope/scanner hardware-software system for data acquisition, image analysis and data reduction. Some procedural capability is available for defining procedures as sequences of commands to execute, as well as save, results in variables. The system presented in the paper is an improved version of the SCANCANS system developed by the same group [BengE76]. In

concept and design, SCANCANS and the later system are similar to BMON2 ([LemP76e], [LemP77a]).

3.2.2.2 High level semantic modelling systems

High level semantic modelling differs from low level modelling in that more than just immediately accessible model knowledge may be brought to bear on the analysis of a problem. Various strategies to perform an analysis may be manipulated in accordance with the model. The model may include visual and non-visual information as well as deduction rules connecting pieces of model knowledge and how they are to be applied in an analysis situation. A high level model is an appropriate place to incorporate statistical trends occurring in one or more frames and allow them to modify structural algorithms used for segmentation on a lower level.

An earlier investigation of procedural description of blood cells [LemP72] simulated the image processing feature extraction functions in constructing preliminary structural models of white cells.

For the last several years our group has been designing and constructing a procedural description language (PRDL [ShapB77] running on a PDP10 computer) to bridge the gap mentioned by Holmquist in hardware/software procedural modeling of biological images. The idea is to distribute the computations so that the time consuming image processing/data acquisition operations would be done on hardware suited (and

dedicated) to that purpose. The control and analysis procedures would be distributed to the procedural description system. Thus advantage could be taken of the resources available on a large system (such as the large word size memory and file system) without saturating its resources doing image processing. The hardware image acquisition/image processor is called the Real Time Image Processor (RTPP) and is described in detail in [LemP76b], [LemP77b],it was first mentioned in [LemP74], [CarmG74].

Before the RTPP was at the current state of completion, an image processing simulation system called PROC10 [Lemp76a] was built on the PDP10, written in SAIL. As PRDL can interact with any type of processor for which it has a processor handler, the actual image processing could take place on either the PDP10, PROC10, EMON2/RTPP, or the final version of the RTPP (which could include a general picture processor, GPP).

3.2.3 Decision trees

Decision tree models have been used to incorporate cell image knowledge in both the segmentation and classification of cell images. As discussed in Section 3.1.1.2 Tycko et al. [TycD76] used a decision tree in computing an initial 5-class classification using special stains to aid the decision process. As discussed in Section 3.1.1.3, Gallus et al. [GalG74] applied a decision tree model to analyzing chromosome boundaries in order to label vertices and centromeric constrictions.

Sychra et al. [SycH76] used a decision tree analysis to describe the shape of nuclei and cytoplasm boundaries as morphologic features for cervical cell analysis. Nuclei are categorized as being

    a. Round/oval
    b. Bean-shaped
    c. Elongated
    d. Irregular/none of the above.

while cytoplasm shapes are categorized as being

    a. Round/oval
    b. Convex/concave
    c. Multiconcave/star
    d. Elongated, spindle- or tadpole-shaped
    e. Irregular/none of the above.

They divided the shape analysis process into two levels of descriptors. The first level descriptors are based on Fourier and chain code analysis. The second level descriptors work at a higher level to provide analysis of concavity, elongation, nuclear eccentricity and slimness. Recognition of features such as (a-d) and (a-e) is also accomplished at the second level. Each of the second level shape features is set to one of five degrees of certainty (no, possibly yes, probably yes, yes, definitely yes). The second level features for the nucleus include:

    1. Roundish or oval shape
    2. Number of large concavities
    3. Slim (length/width ratio)
    4. Smoothed contour
    5. Concave
    6. Elongated or lobate
    7. (Spiral) hook
    8. Bean shape
    9. Unclassified (result of decision tree failure).

Bacus and Weens [BacJ77] constructed a decision tree

classification scheme for red blood cells from peripheral blood. The cells were first prepared with a slide spinner to disperse clumps so that single objects were usually single cells. The unstained cells were scanned at 418 nm (the Soret hemoglobin absorption peak) for maximum contrast. Then ten features were computed from the isolated objects, includiing area, perimeter squared/area, spicularity, elongation, hemoglobin concentration, and central pallor analysis. (Red cells have a light inner region which was segmented after the cell as a whole was found. Chain codes for the entire cell and the central region were computed and used with other measures to compute the ten features.)

Using the ten features, a decision tree was designed to classify the cells into 14 classes using 12 thresholds on the features. The decision tree classified cells by assigning morphologic labels corresponding to particular features to groups, subgroups and sub-subgroups of cells as the depth of the tree increased. The tree consisted of a major subclassifier followed by four minor subclassifiers.

Bacus reports on the results of applying this technique to normal slides. The threshold parameters were determined partly by parameters in the literature and partly by experiment. Bacus then presents successful results of applying this procedure to get cell counts for both normal and abnormal red cell specimens.

4. SEGMENTATION STRATEGY

The overall segmentation hierarchy is described in a decision tree formalism. A running example, using frame 126-6, will be used for all of the algorithms. Other frames may also be included in order to illustrate particular points for some parts of the algorithm.

The morphologic data base discussed in Section 2.2 offers some hints as to how one might go about segmenting marrow. Several conjectures as to segmentation strategies are inferred from the marrow morphology.

(a) Marrow is generally visualized as a set of dark formed objects on a lighter possibly non-uniform background.
(b) Formed objects are most often defined by membrane boundaries which may or may not be easily discernible but are known to exist.
(c) The darkest objects are most easily detected. Removing them simplifies the segmentation of what remains by reducing the combinatorics of the scene.
(d) Context should be used in analyzing objects wherever possible, especially in those cases with poor boundary (membrane) definition (e.g. expand a well defined nucleus region into a poorly defined cytoplasm region; find sharp cusps indicating touching nuclei regions so as to indicate where they should be split).

Taking such a top-down context-dependent route leads to a possible segmentation decision procedure. The most easily detected objects are recognized and removed first. Then information gained from detecting objects in various classes is used in analyzing the more difficult objects of the nth class at step n of the segmentation process.

This section gives an overview of the segmentation strategy. The details of the experimental implementation of the tree are given in Sections 5, 6, 7 and 8, which cover RBC removal and extraction, large object segmentation, single cell articulation, and small object segmentation, respectively.

4.1 Segmentation decision tree

The segmenting of an image may be thought of as the successive recognition and elimination of objects from the image (thus accounting for objects in the frame). The algorithm attempts to find those objects which (a) meet the simplest semantic criteria, and (b) are inner membrane objects (regions) capable of being expanded outward to "claim" additional pixels belonging to them.

4.1.1 General approach

Figure 4.1 shows the general segmentation decision tree. This tree illustrates the general types of operations necessary to segment a marrow image.

The frame is presented to the algorithm as a set of multispectral images of the same marrow field. The first objects in the scene which meet the top-down morphologic criteria are the hemoglobin-containing red blood cells. These are removed first. Then large dense objects are identified which are probably individual or touching leukocyte nuclei. The large objects are tested to see whether they are actually
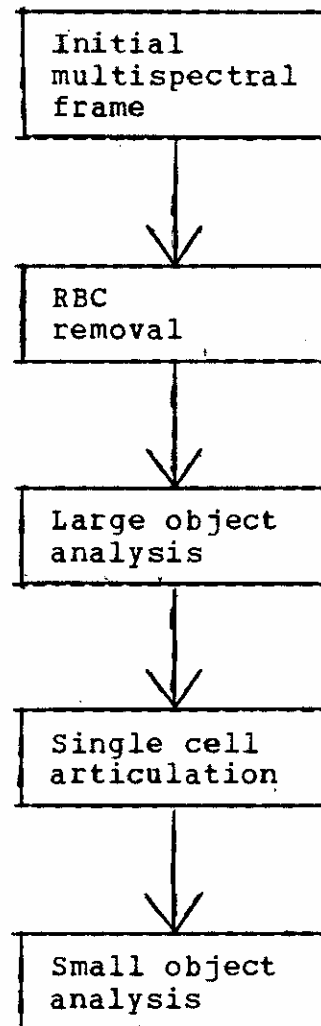
The page number 61 is at the top.

```
                    ┌─────────────────┐
                    │ Initial         │
                    │ multispectral   │
                    │ frame           │
                    └─────────────────┘
                             │
                             ▼
                    ┌─────────────────┐
                    │ RBC             │
                    │ removal         │
                    └─────────────────┘
                             │
                             ▼
                    ┌─────────────────┐
                    │ Large object    │
                    │ analysis        │
                    └─────────────────┘
                             │
                             ▼
                    ┌─────────────────┐
                    │ Single cell     │
                    │ articulation    │
                    └─────────────────┘
                             │
                             ▼
                    ┌─────────────────┐
                    │ Small object    │
                    │ analysis        │
                    └─────────────────┘
```

Figure 4.1 Block diagram of the general segmentation decision tree.

groups of touching nuclei. If they are, they are broken into smaller nuclei. The small cell analysis goes on to assign cytoplasm to the nuclei. Similarly the small object analysis segments platelets outside of cells, and nucleoli inside of nuclei. These formed objects are identified and removed from the image.

Additional analysis may be performed, if indicated, causing the decision tree to be reentered (thus converting the

' tree to a cyclic graph).

As indicated in Section 1.1.3, the output of the segmentation process is a set of connected component images CCi representing various types of formed objects. In each of these images, all pixels of a given object are labeled with the same value $k > 1$.

Object$(k,i) = \{(x,y) \mid g(x,y)$ of CCi=k$\}$.

The connected component images constructed by the analysis are as follows:

CCnuc=$\{$CCx $\mid$ x is a nucleus$\}$

CCcyt=$\{$CCx $\mid$ x is the cytoplasm of a nucleus y in CCnuc$\}$

CCnucleoli=$\{$CCx $\mid$ x is a nucleolus of a nucleus y
            in CCnuc$\}$

CCcell=$\{$CCxy $\mid$ (x in CCnuc) and (y in CCcyt) and
        (CCx belongs to CCy)$\}$

CCplate=$\{$CCx $\mid$ x is a platelet$\}$

CChg=$\{$CCx $\mid$ (x is RBC) or ((x in CCcell) and
        (x contains hemoglobin))$\}$

CCframe=$\{$CCx $\mid$ (x touching frame) and
        (x < min size for nucleus)$\}$

CCbkgrd=$\{$CCx $\mid$ (x not in CCcell) and
        (x not in CCplate) and (x not in CChg)$\}$

CCartifact=$\{$CCx $\mid$ (x is artifact) and
        (If x in CCi Then remove x from CCi)$\}$.

The last rule is different from the others in that it has the side effect of removing components from other CC sets if after they were mislabeled as non-artifacts, they are discovered to be artifacts and are relabeled as artifacts.

The decision procedure for a set CCi of objects of type i may be thought of as executing a fixed sequence of actions. Assume that the scene has been pre-processed up to the point just before CCi is to be segmented. At this stage, the following conditions should hold:

(a) Only objects in CCi will be segmented by segmentation procedure Sni.
(b) The consistency of objects in CCi may be checked using the global marrow model at step n of the segmentation process. This implies that the model has embedded global knowledge which would permitt such consistency checking.
(c) The morphologic model may suggest mechanisms for backing up the segmentation in order to re-evaluate inconsistent objects in CCi (converting the tree to a cyclic graph).

The basic CCi segmentation procedure can viewed as a three-step process as shown in the block diagram in Figure 4.2. No backtracking is performed in the present implementation.

4.1.2 Implementation

Figures 4.3 and 4.4 show the details of representative pieces of the decision tree which were implemented. These pieces were chosen based on an analysis of the data base to determine the most commonly occurring events (cf. Section 1.3.2).

Images are generally denoted by a two character prefix followed by a lower case semantic label (e.g. CCnuc, where 'nuc' denotes the image having to do with the nucleus). The prefixes used in the following sections are given in Table 4.1.

```
┌─────────────────┐
│ Segment frame   │
│ at step n, Fn,  │
│ to detect       │
│ objects of      │
│ type i ==> CCi  │
└─────────────────┘
         │
         V
┌─────────────────┐
│ Consistency     │        FAIL:
│ check of CCi    │──────> backup segmentation
│ with model      │        using information
│ of Fn           │        from failure.
└─────────────────┘
         │ SUCCEED:
         V
┌─────────────────┐
│ Remove CCi      │
│ from Fn==>Fn+1  │
└─────────────────┘
```
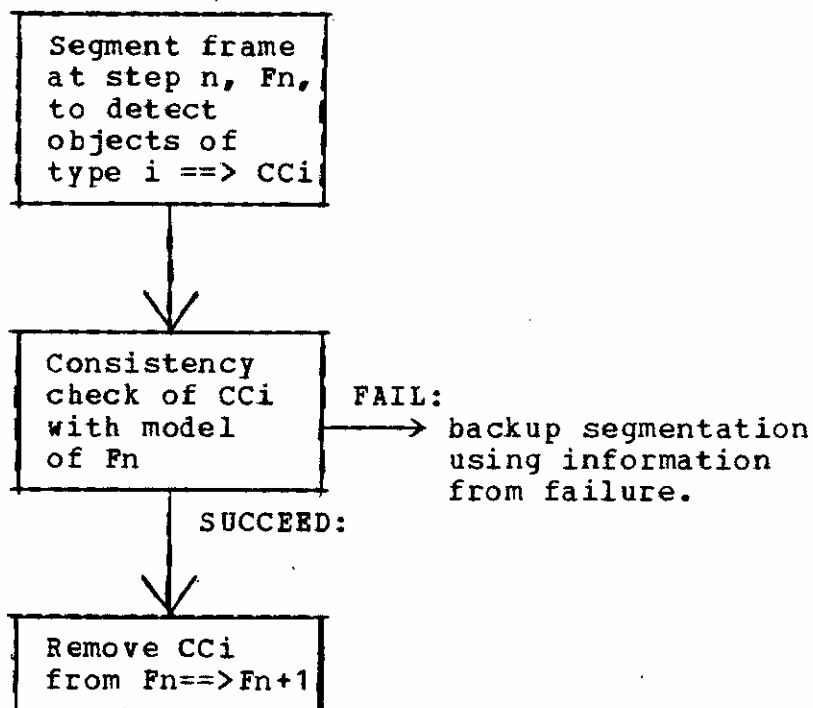
Figure 4.2 Block representation of the three-step primitive segmentation procedure for step n in analyzing a frame F producing image CCi.

Image prefix	Function
GN	green (546 nm)
BL	blue (420 nm)
GB	green-blue positive difference
CC	connected component image
MS	mask image
CN	pre-isolation image
CP	pre-propagation image

Table 4.1 Image prefixes used in the following sections.

The initial frame consists of the green (GN) and blue (BL) unnormalized images. The RBCs are extracted by thresholding the unnormalized blue image, creating a MSrbc. This mask will contain nucleated erythrocytic cell nuclei as

well which have to be articulated later.

The green-blue positive difference image (GB) is computed after normalizing the GB and BL images. Using GB, the large nuclei are extracted in the form of a connected component image CCnuc0. Analyzing these nuclei results in splitting touching nuclei to compute a connected component image CCnuc. This is used to remove the nuclei from image GB creating image GB'.

Nuclei or parts of other dark objects less than 10 microns in area that touch the edge of the image are extracted, creating a connected component image CCedge, and are removed from the GB' image to obtain image GB''.

Similarly, the platelets are extracted creating a connected component image CCplatelet, and are removed from the GB'' image, creating the GB''' image.

Image GB''' and the complement of the object mask MSformed are used to obtain a potential cytoplasm mask image. With the aid of the original green image (to greater differentiate between cytoplasm and background, the cytoplasm image mask MScyto is computed. This is then used along with the CCnuc image to construct the pre-isolation connected component image CNcell. The MSformed mask is the union of the masks defined by CCnuc, CCedge, and CCplatelet. Thus the region corresponding to the cytoplasm and the background, excluding the RBCs, is (GB&(not MSformed)).

Cytoplasm regions are associated with particular nuclei using the isolation algorithm to create the pre-propagation image CPcell. This is used in propagating the nuclei to their associated cytoplasm, creating the cell connected component image CCcell. At this point, connected components for individual cells exist.

Using the MSrbc mask computed previously, the algorithm checks to see if any nuclei have hemoglobin, which may indicate that they are nucleated cells from the erythrocytic cell line.

Using the CCnuc image, the algorithm checks to see whether nucleoli exist in any of the cells, thus creating the connected component image CCncleoli.

```
┌─────────────────────────────────────────┐
│ Initial multispectral frame GB, BL       │
└─────────────────────────────────────────┘
                    │
                    ▼
Continue      ┌─────────────────────────────────────────┐
nucleated  ◄──┤ RBC extraction and removal               ├──► MSrbc
erythrocyte   └─────────────────────────────────────────┘
analysis            │ GB
(Figure 4.4)        ▼
              ┌─────────────────────────────────────────┐
              │ Large object analysis-split nuclei       ├──► CCnuc0
              └─────────────────────────────────────────┘
                    │ GBsplit
                    ▼
              ┌─────────────────────────────────────────┐
              │ Single cell articulation                 ├──► CCnuc
              └─────────────────────────────────────────┘
                    │ GBsplit'
                    ▼
              ┌─────────────────────────────────────────┐
              │ Edge object removal                      ├──► CCedge
              └─────────────────────────────────────────┘
                    │ GBsplit''
                    ▼
              ┌─────────────────────────────────────────┐
              │ Platelet removal                         ├──► CCplatelet
              └─────────────────────────────────────────┘    MSformed
                    │ GBsplit'''
                    ▼
              ┌─────────────────────────────────────────┐
              │ Cytoplasm mask generation                ├──► MScyto
              └─────────────────────────────────────────┘
                    │ CNcell
                    ▼
              ┌─────────────────────────────────────────┐
              │ Cytoplasm isolation between nuclei       │
              └─────────────────────────────────────────┘
                    │ CPcell
                    ▼
              ┌─────────────────────────────────────────┐
              │ Cytoplasm propagation from nuclei        ├──► CCcell
              └─────────────────────────────────────────┘
```

Figure 4.3 Block diagram decision tree of the implemented segmentation strategy.

Continued
(from Figure 4.3)

```
        ┌─────────────────┐
        │ Wait for images │
        │ CCcell, GBsplit │
        │ MSrbc, BL       │
        └─────────────────┘

        ┌─────────────────┐
        │ Nucleated       │──────→ MSrbc'
        │ erythrocyte     │       (RBCs only)
        │ check           │
        └─────────────────┘
```

CCnuc marked as nucleated RBCs
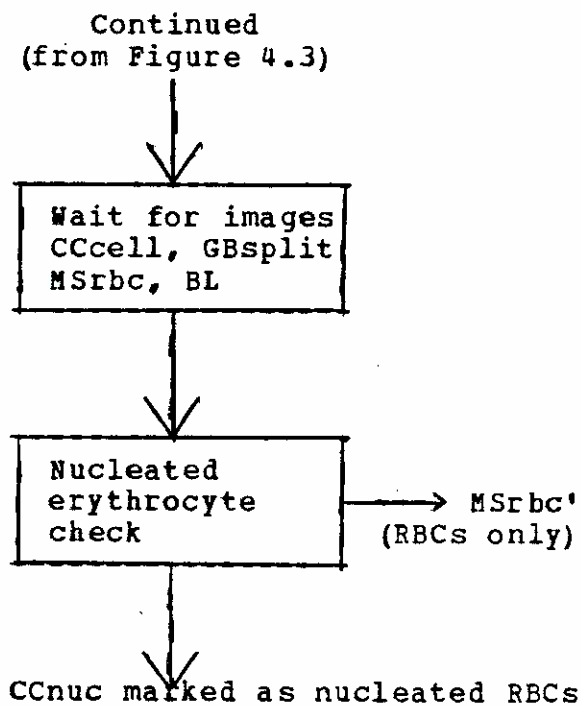
Figure 4.4 Continuation (from Figure 4.3) of implemented decision tree showing nucleated erythrocyte procedure.

4.2 Notation used in algorithms and examples

Before discussing the details of the various stages of the segmentation algorithm, some notational definitions are given. These include the coordinate system and the gray scale conventions. The algorithms are presented in an ALGOL-like language which makes certain assumptions about the composition of picture functions.

4.2.1 Logical coordinate system and grayscale conventions

The image pixel coordinate system used is that of the RTPP. This is called the Logical Coordinate System or LCS The LCS has (0,0) as the upper left hand corner and (1023,1023) as the lower right hand corner. The maximum size of an indivual image used here is 256x256 pixels. This corresponds to a single image buffer memory in the RTPP.

An image is defined here to have eight bits (0 to 255) of gray-scale information associated with each pixel position (corresponding to the image buffer memory datum). The value of a black pixel is defined to be 255 while that of a white pixel is 0.

4.2.2 Use of ALGOL-like notation in algorithms

An ALGOL-like notation is used in the algorithms given in the following sections and appendices. It is felt that such notation facilitates the communication of the algorithms

without encumbering the description with unnecessary details.

The comments, which may extend beyond one line, are set off between quote (") pairs. Assignment is made using the "_" symbol, or the "<==" if a matrix (or image) is being assigned a matrix or single value. Assignment embedding is allowed, with the embedding operation having higher precedence of evaluation than the usual arithmetic operators. For example, the following are embedded statements:

```
i_j+(k_3*6);

If (a_NEXT!ELEMENT {alpha}) < q
        Then ...
        Else ...;
```

Both the "Begin/End" notation with lists of expressions separated with ";"'s and expressions separated by ","'s are used to denote sequences of statements to be evaluated. The latter notation is used when the statements may be evaluated in parallel or when evaluation order is not important.

Sets are denoted by "{" and "}" pairs, lists by "(" and ")" pairs, and subscripts by "[" and "]" pairs.

Sometimes an English-like description of the operation to be performed is used instead of detailed code. For example,

```
Ij<==Segment Ii at threshold t;
```

In such cases, either the algorithm described has been discussed elsewhere in the dissertation or the details are obvious.

4.2.3 Composition of picture operations

Picture operations include both point and neighborhood operations. The notation used does not explicitly mention the order of application of the domain of the function to each pixel or neighborhood of the argument image. In most cases this is evident and is not explicitly stated. An example of an implicit computation is:

    Ij<==AVG8(Ii);

    Ij<==SLICE(Ii,t1,t2);

Where the first function computes the 8-neighbor average of Ii and the second function computes the threshold slice function (= g if (t1<g<t2), else =0). The implied (x,y) tessellation is over [0:255x0:255] and is explicitly stated as:

    For y_0 Step 1 Until 255 Do
        For x_0 Step 1 Until 255 Do (operation at (x,y));

The composition of these operations uses the assumptions of algebraic composition (using a similar notation). For example:

    Ij<==AVG8(SLICE(Ii,t1,t2));

computes the slice function first and then the 8-neighbor average.

5. RBC REMOVAL AND ANALYSIS

The marrow segmentation strategy, as discussed in Section 4, concentrates on eliminating the most easily segmented formed objects first. This section discusses the removal of the RBCs from the frame since they constitute a large number of the formed elements and have a simple means of being eliminated. Then discrimination between RBCs and nucleated erythrocytes is discussed.

5.1 Eliminating RBCs

An algorithm is presented for eliminating formed elements with hemoglobin content (mature reticuloytes) from bone marrow Wright's stained smears imaged through an optical microscope. The algorithm subtracts a normalized blue image from a normalized green image using a grayscale positive difference function resulting in a grayscale image.

Methods exist for eliminating red blood cells (RBCs) from peripheral blood images ([YouI75], [BacJ75] and others). These methods use the fact that hemoglobin has a maximum light absorption at 419 nm wavelength. Thus finding the RBCs necessitates finding the discriminant threshold of a blue scan corresponding to the RBC pixels. Using the valley between two peaks as the optimal threshold for separating the components corresponding to the two peaks was suggested in [PrewJ66].

Unfortunately, noise due to optics, digitization, and

72

illumination and sensor shading non-uniformities tends to cause problems with this type of threshold selection. In addition, in complex images such as clusters of cells as seen in bone marrow, other effects tend to make the use of a single threshold difficult. This will be discussed in more detail later. Thus applying thresholding of the blue scan to produce a binary image results in a noisy RBC mask with shading error as well as salt and pepper noise. In addition, image content affects selection of the threshold, making optimal threshold selection difficult.

Six frames (numbered 3,4,5,6,7,8) were selected illustrating complex clusters of leukocytes and reticulocytes in various stages and configurations. The green (GN) and blue (BL) scans are shown in Figure 5.1.

5.1.1 Problems with thresholding the blue image

As just mentioned, the blue image is not always easily segmentable into RBCs and other objects by thresholding, because the RBC threshold for complex scenes is so variable. Figure 5.2 illustrates several thresholds for three different blue images at approximately the RBC threshold Trbc (110) and at two values above it (120, 130). These values were obtained by looking at the blue image histograms shown in Figure 5.3. The blue histograms show additional peaks for some normoblasts in Figure 5.3b while the normoblast peak in Figure 5.3j is not apparent.
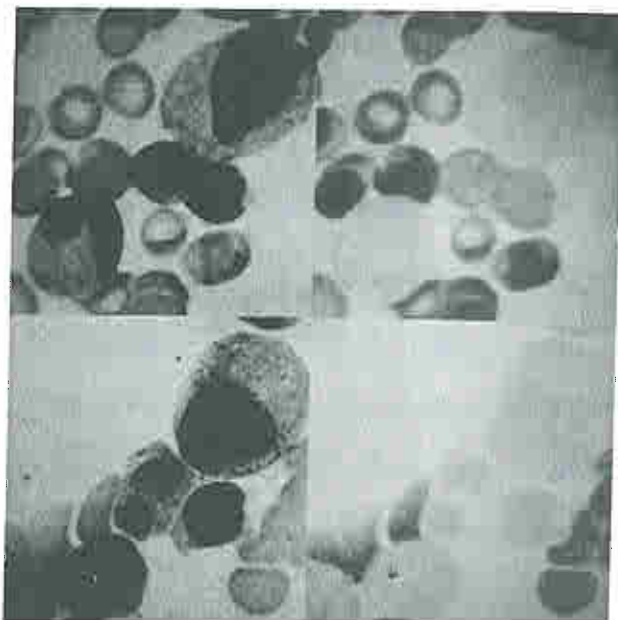
a-d
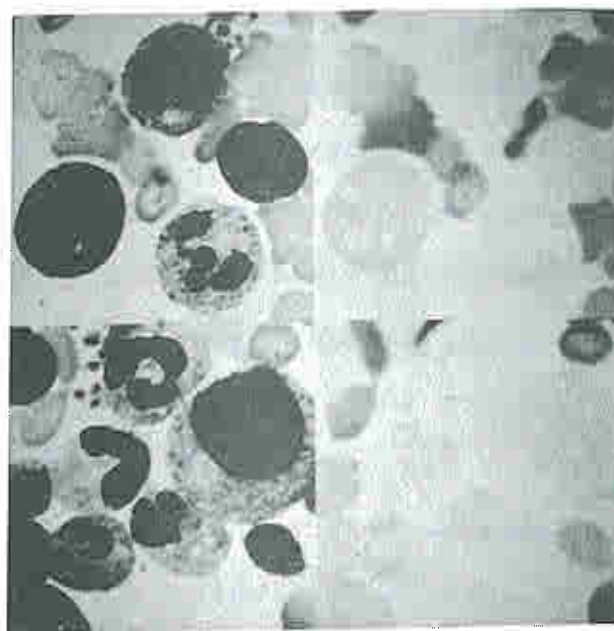
e-h
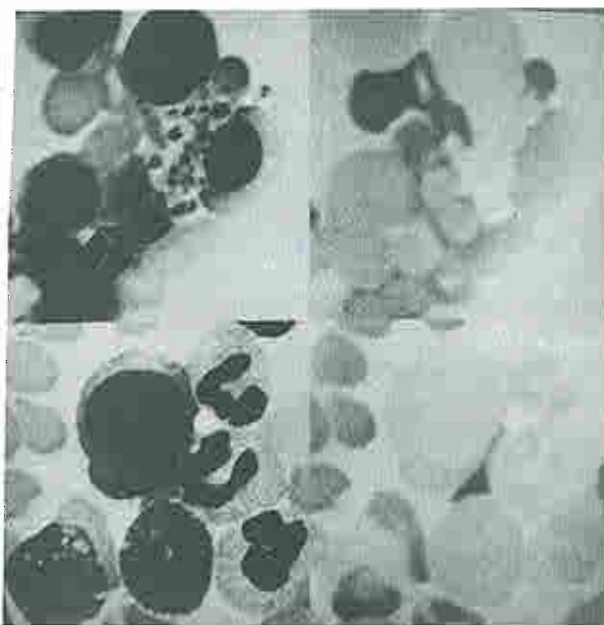
i-l

Figure 5.1 Six original green and blue scanned images. a) 3-green, b) 3-blue, c) 4-green, d) 4-blue, e) 5-green, f) 5-blue, g) 6-green, h) 6-blue, i) 7-green, j) 7-blue, k) 8-green, l) 8-blue.

The two normoblast nuclei (which contain some hemoglobin) in Figure 5.2 a-d are not as dense as the RBCs, yet they will effect a shift in the RBC threshold. If the threshold is raised slightly to remove them, the RBCs start to fragment. In Figure 5.2 e-h, those RBCs overlying one another or squeezed in between other cells as well as the smashed RBC in the center left of the image are denser than isolated RBCs in other parts of the images. In Figure 5.2 i-l, there is also the case of a RBC overlying a white cell, the RBC's density being in between that of the isolated RBCs and those RBCs squeezed or overlying one another.

Thus, there is no unique threshold for eliminating all of the RBC content from the image due to local variation in RBC density depending on how it presents itself in conjunction with other objects in the image.

Another technique, developed by Bacus [BacJ76], called the "whitening transform", maps two colored images into a "colored" image C and a "density" image D. This is done by computing the principal eigenvectors and eigenvalues for each image to maximize separation. Bacus makes two assumptions: that the principal eigenvector lies along a 45 degree line in the original bivariate space, and that the variances in either direction are relatively constant. For the blue and green images used in our experiments, this transform is given by

a-d

e-h
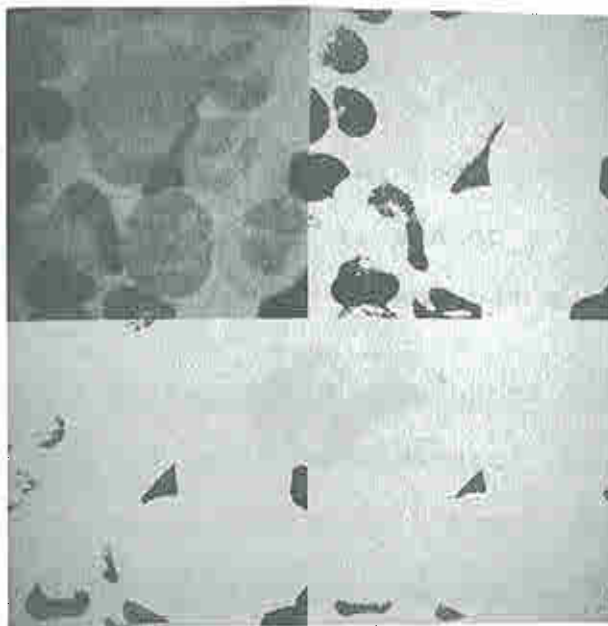
i-l

Figure 5.2 Three blue images sliced at about the RBC threshold, Trbc, and above (110, 120, and 130). a) 3-BL, b) 3-BL sliced 110, c) 3-BL sliced 120, d) 3-BL sliced 130, e) 5-BL, f) 5-BL sliced 110, g) 5-BL sliced 120, h) 5-BL sliced 130, i) 8-BL, j) 8-BL sliced 110, k) 8-BL sliced 120, l) 8-BL sliced 130.
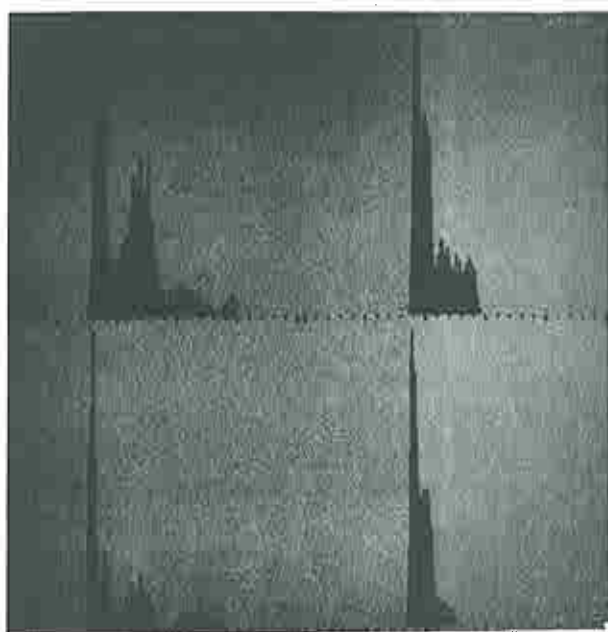
a-d
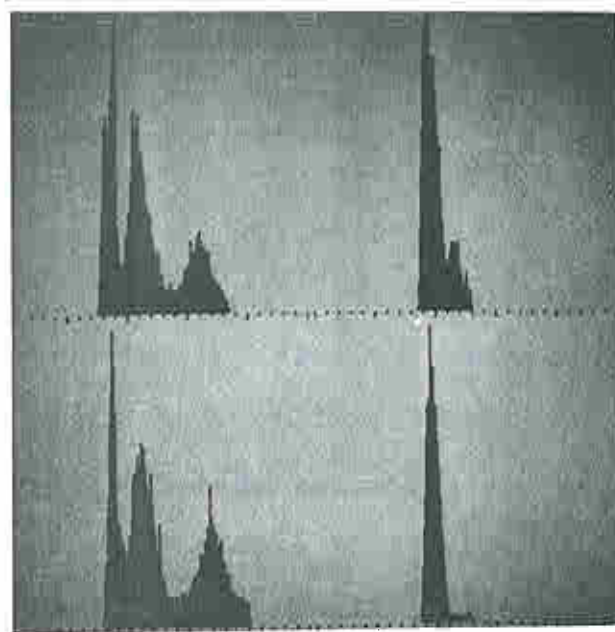
e-h
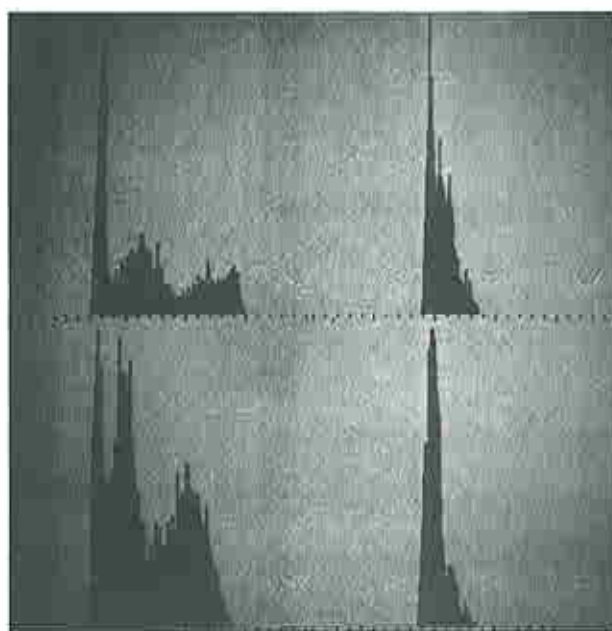
i-l

Figure 5.3 Six histograms of the original green and blue scanned images (with grid marks every 10 gray levels). a) 3-GN, b) 3-BL, c) 4-GN, d) 4-BL, e) 5-GN, f) 5-BL, g) 6-GN, h) 6-BL, i) 7-GN, j) 7-BL, k) 8-GN, l) 8-BL.

$$C = K(BL-GN),$$

$$D = K(GN+BL).$$

Figure 5.4 shows the whitening transform for frames 6, 7 and 8. Note that it did not work very well for frame 6 but did seem to work (to some extent) for frames 7 and 8. The lack of discrimination seems to be due to overlapping RBC's which smear out the blue contributions in a continum (depending on the extent of overlap). This overlap problem is especially common with marrow smears but does not occur as frequently with the spinner-produced smears of peripheral blood used by Bacus.

Because of the problem of overlap, a non-linear technique was developed in order to eliminate as much as possible of the different RBC populations (normal RBCs, rouleaux, RBCs overlapping WBCs).

5.1.2 Image normalization

In order to eliminate local density effects, the two images are normalized by their density extrema. The extrema are computed during histogram computation for the two images. Figure 5.3 shows the histograms for the green and blue images. The normalized blue and green image transforms are defined by

$$gmin = MIN(g(x,y) \mid g(x,y) \text{ in } GN).$$

$$gmax = MAX(g(x,y) \mid g(x,y) \text{ in } GN).$$

$$bmin = MIN(g(x,y) \mid g(x,y) \text{ in } BL).$$

$$bmax = MAX(g(x,y) \mid g(x,y) \text{ in } BL).$$
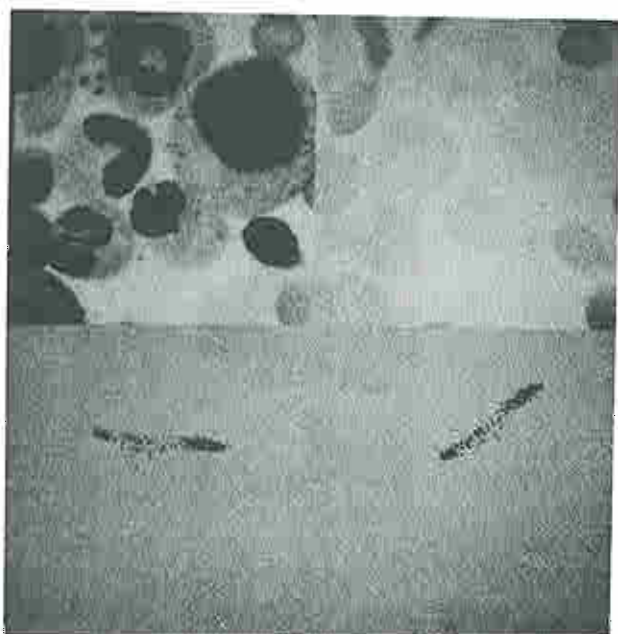
a-d
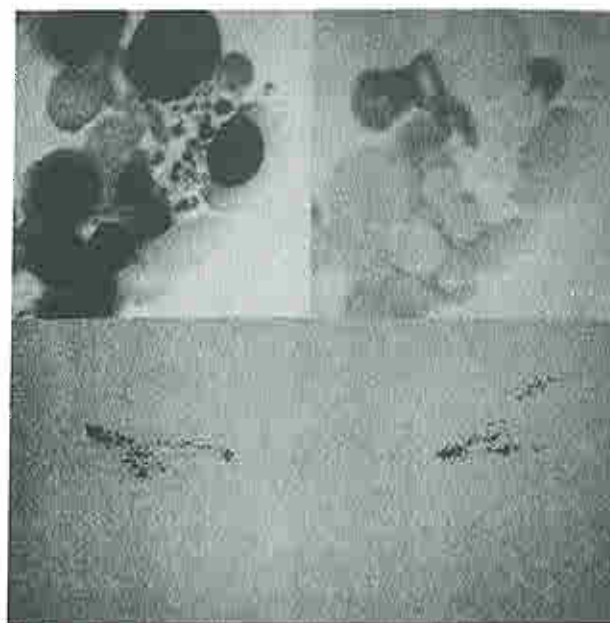
e-h
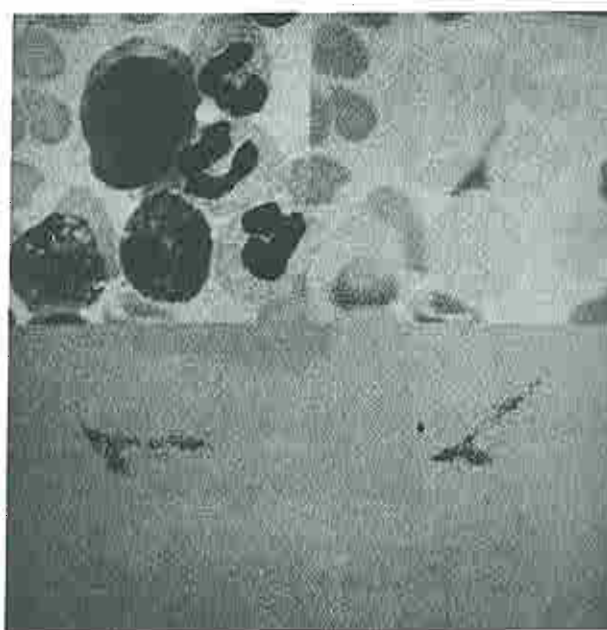
i-l

Figure 5.4 Comparisons of joint gray scale distributions of (GN.vs.BL) and (C.vs.D). The x-axis (left=0) of the joint distribution is GN (C) while the y-axis (top=0) is BL (D).a) 6-GN, b) 6-BL, c) 6-GN.vs.BL, d) 6-C.vs.D, e) 7-GN, f) 7-BL, g) 7-GN.vs.BL, h) 7-C.vs.D, i) 8-GN, j) 8-BL, k) 8-GN.vs.BL, l) 8-C.vs.D.

Then, the normalized images GN' and BL' are computed using linear interpolation to fill an 8-bit dynamic range:

$$GN'(x,y) = \frac{(gmax - gmin)(GN(x,y) - gmin)}{255}.$$

$$BL'(x,y) = \frac{(bmax - bmin)(BL(x,y) - bmin)}{255}.$$

Figure 5.5 shows the GN' and BL' normalized images which have better contrast than the original images shown in Figure 5.1 because of the contrast stretching effect of the normalization.

5.1.3 Positive difference transform

The positive difference transform is an asymmetric arithmetic operation which forces negative gray values to zero. This has the effect of losing some of the information. Figure 5.6 illustrates the positive differences of GB=(GN'-BL') and BG=(BL'-GN') computed using the equations below. As can be seen, the second difference (BG) is not very useful for white cell extraction, whereas the first difference (GB) contains the desired subtraction of hemoglobin-formed objects. The BG difference seems to contain hemoglobin-formed objects as well as background.

$$POSDIFF(x,y) = \text{If } x-y < 0$$
$$\text{Then } 0$$
$$\text{Else } x-y.$$
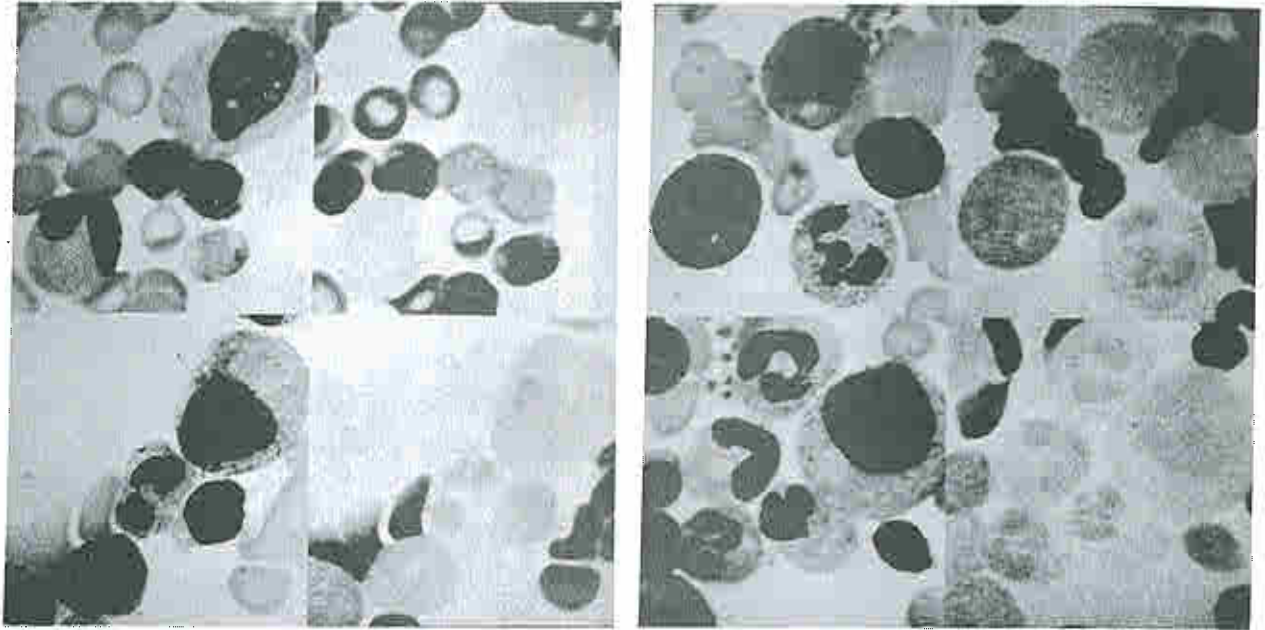
a-d

e-h
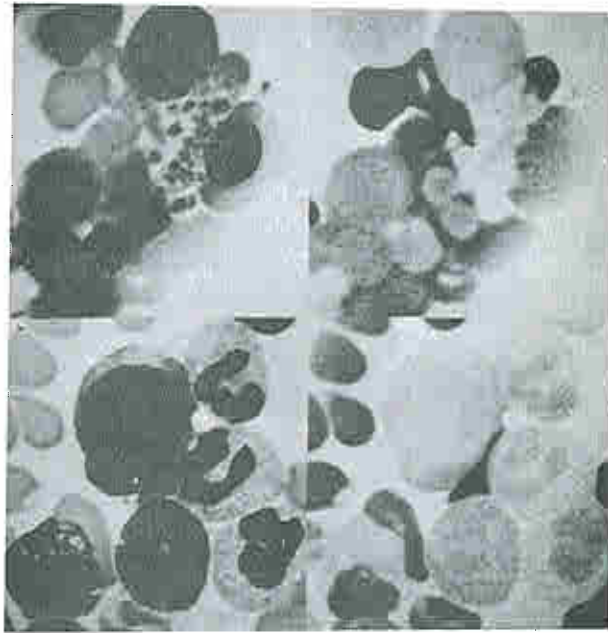
i-l

Figure 5.5 Six normalized green and blue images. a) 3-GN', b) 3-BL', c) 4-GN', d) 4-BL', e) 5-GN', f) 5-BL', g) 6-GN', h) 6-BL', i) 7-GN', j) 7-BL', k) 8-GN', l) 8-BL'.
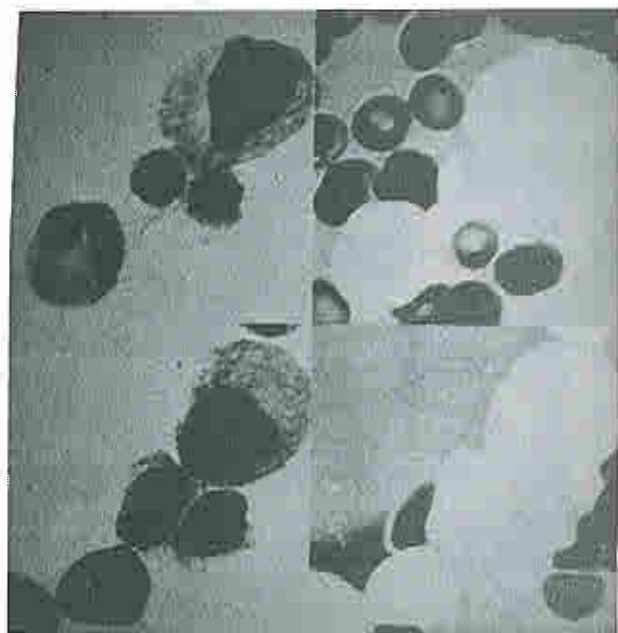
a-d
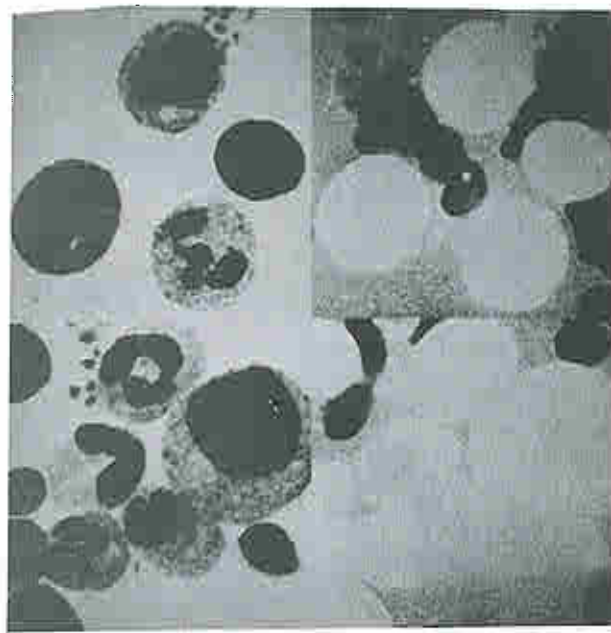
e-h
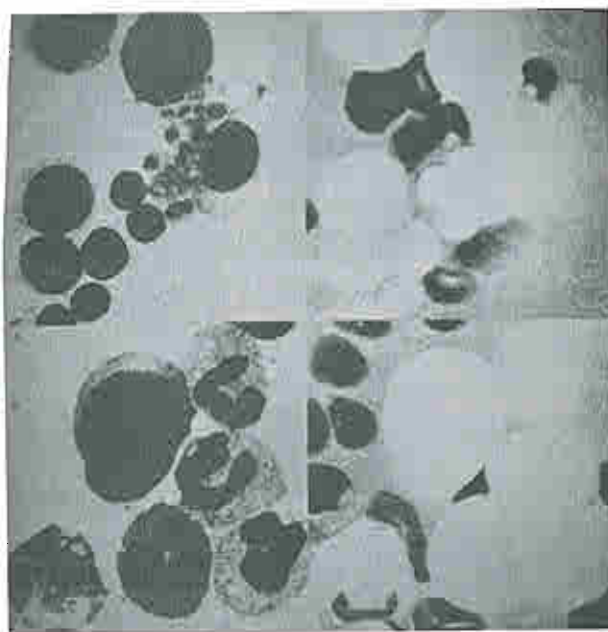
i-l

Figure 5.6 Positive differences. a) 3-GB, b) 3-BG, c) 4-GB, d) 4-BG, e) 5-GB, f) 5-BG, g) 6-GB, h) 6-BG, i) 7-GB, j) 7-BG, k) 8-GB, l) 8-BG.

$$GB = POSDIFF(GN',BL').$$

$$BG = POSDIFF(BL',GN').$$

5.1.4 Results of RBC extraction

Figure 5.7 shows the histograms and the smoothed histograms of the GB images. The GB histograms may be interpreted as follows: 0 to Tcyto is the range for plasma proteins and background, Tcyto to Tnuc is the range for cytoplasm and parts of platelets, and Tnuc to Tmax is the range for nuclei.

If (as it turns out) Tcyto and Tnuc are fairly constant, then no histogram analysis of the GB histograms is necessary. This is helpful since these histograms are noisy because of normalization. Looking at the GB histograms one notes that Tnuc centers around 60 fairly independently of the content of the image and that Tcyto centers around 10 also fairly independently of the image content. Because the cytoplasm intersects the background density distribution, Tcyto is dropped to 5 to keep down fragmentation. Figure 5.8 shows the six images sliced at [Tnuc:255] with very good extraction of the nuclei. The failures in Figure 5.8f were caused by the morphology of the basket cell which was heavily vacuolated and had overlapping RBCs, and the heavy granulation of the mature eosinophil. Both had problems with the defintion of the nuclear
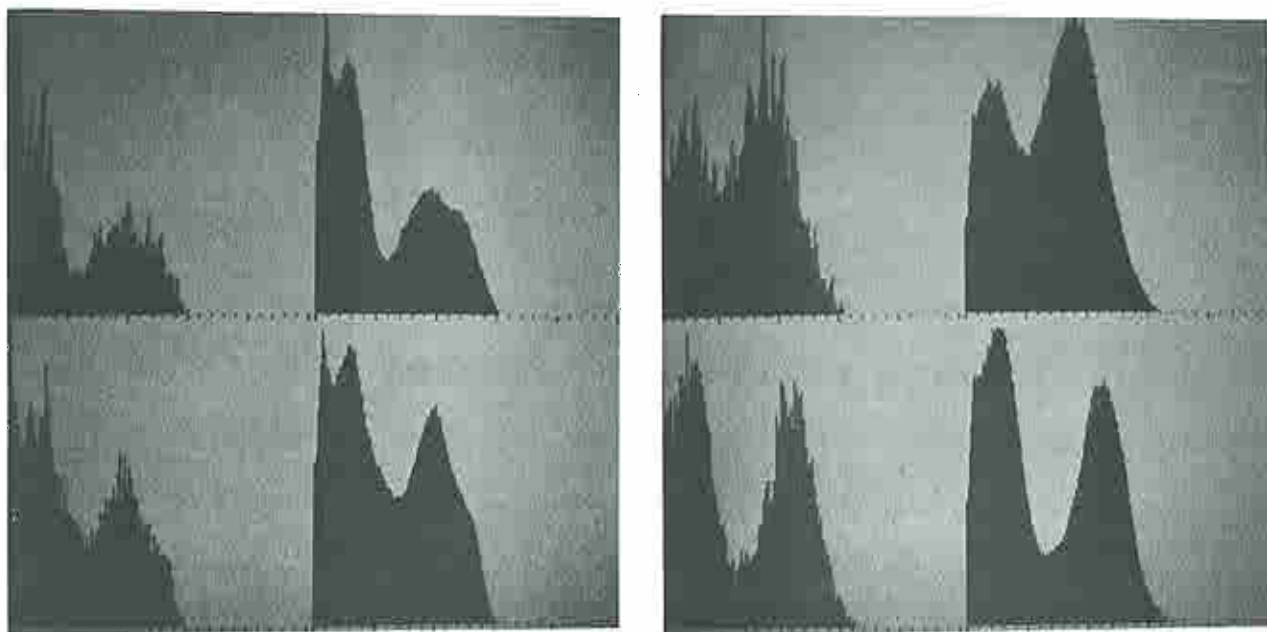
a-d

e-h
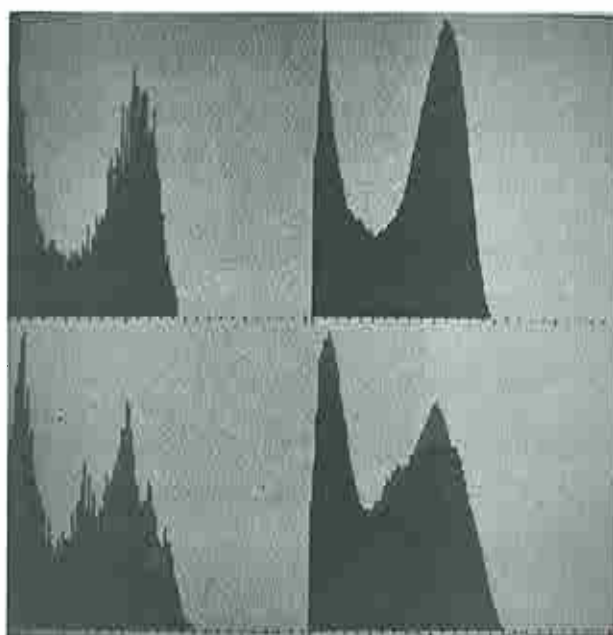
i-l

Figure 5.7 Histograms of positive differences (with grid marks every 10 gray levels). a) 3-H(GB), b) 3-SMOOTH(GB), c) 4-H(GB), d) 4-SMOOTH(GB), e) 5-H(GB), f) 5-SMOOTH(GB), g) 6-H(GB), h) 6-SMOOTH(GB), i) 7-H(GB), j) 7-SMOOTH(GB), k) 8-H(GB), l) 8-SMOOTH(GB).

border.

In summary, a non-linear gray scale transform for multispectral bone marrow images can be used to successfully eliminate formed elements containing hemoglobin. This transform also has the property that nuclear and background thresholds are fairly constant and independent of image content. This algorithm is thus a useful pre-processing step in segmenting images of bone marrow smears.

a-d

e-f
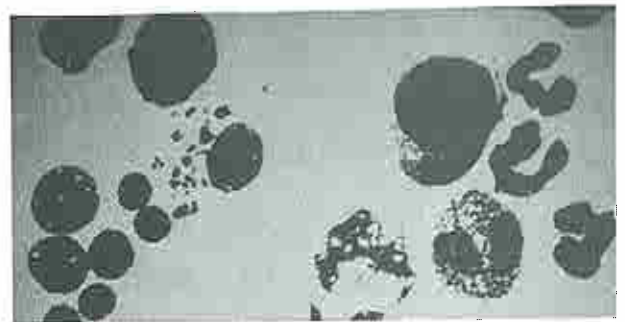
Figure 5.8 Six GB images sliced at [60:255]. a) 126-3, b) 126-4, c) 126-5, d) 126-6, e) 126-7, f) 126-8.

5.2 Erythrocyte extraction

Erythrocytes, because of their hemoglobin content and its characteristic absorption at 419 nm, are segmentable as a group using a threshold discriminant on the blue scan. This threshold is at the valley of the blue density histogram. This discriminant divides the frame into two regions: (leukocytes, platelets, background) and (RBCs, nucleated erythrocytes). Figure 5.9 shows the blue histograms of the thirteen frames in the data base with the RBC threshold valley obvious in only some of the cases. As Bacus points out [BacJ76], scarcity of RBCs in the image (as in the megakaryocyte in 127-13) might cause problems in finding the valley.

As was noted in the Section 5.1, there are sometimes extra valleys in the histogram due to overlap of RBCs. with RBCs (called rouleaux) or with leukocytes, as well as to the occurrence of smashed RBCs which tend to have a lower hemoglobin content. The former (latter) causes additional valleys at the higher (lower) denstity histogram than would be due to RBCs occurring alone.

Figure 5.10 shows the blue images sliced at Trbcs determined for each of the blue image histograms.

a-d

e-h

i-l

m
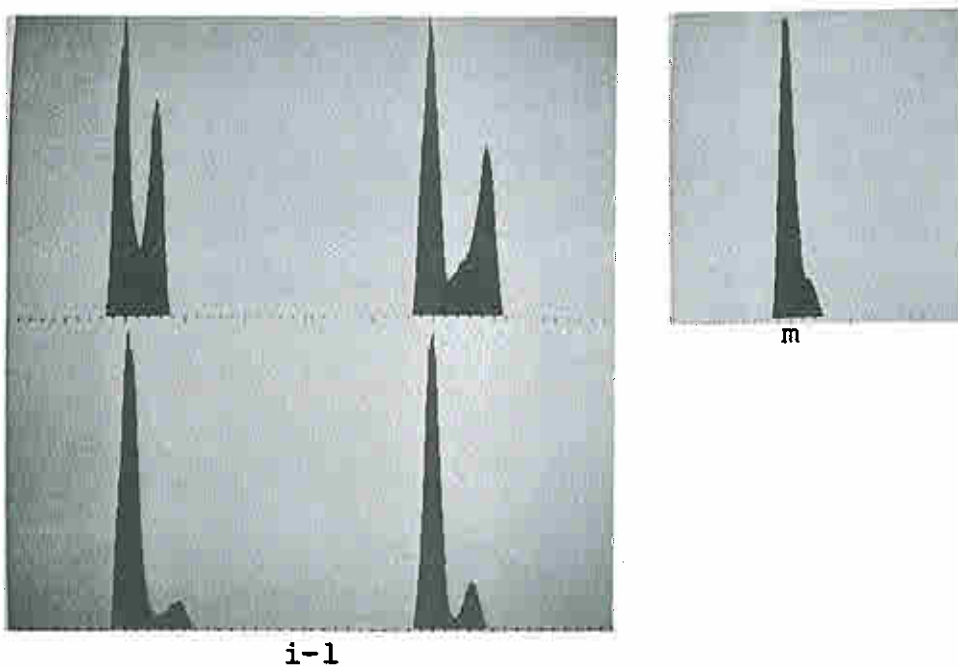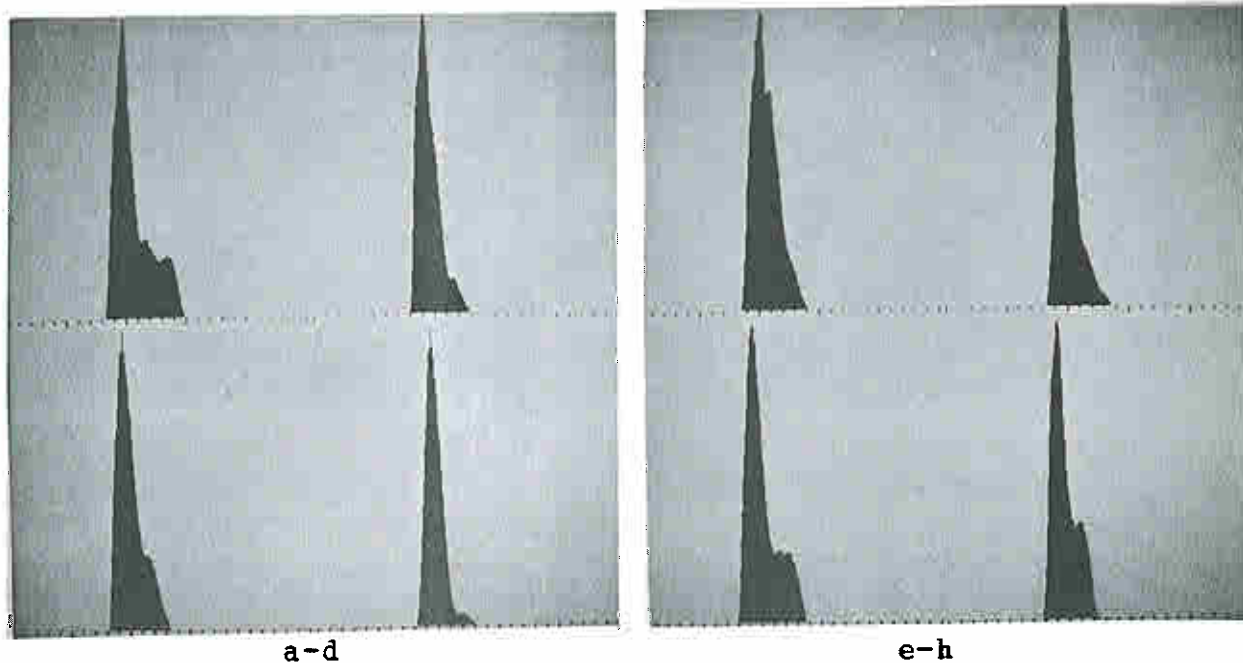
Figure 5.9 Density histograms of the smoothed blue images (with grid marks every 10 gray levels). a) 126-3, b) 126-4, c) 126-5, d) 126-6, e) 126-7, f) 126-8, g) 126-9, h) 126-10, i) 127-11, j) 127-12 k) 127-13, l) 127-14, m) 127-15.
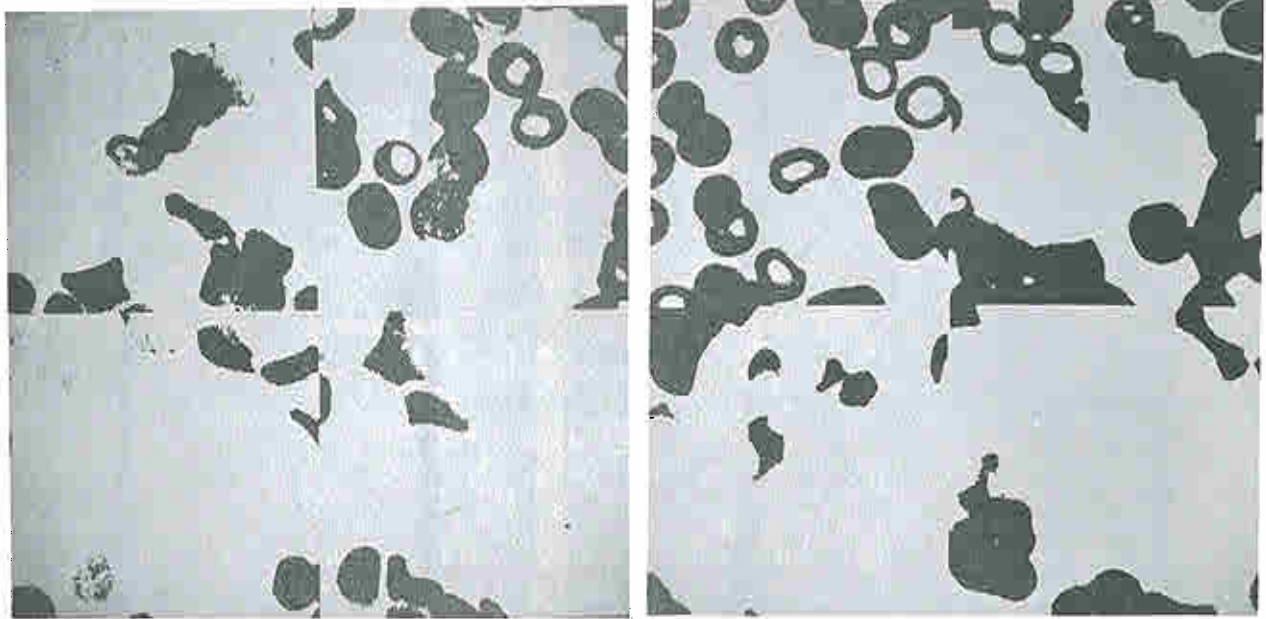
a-d   e-h
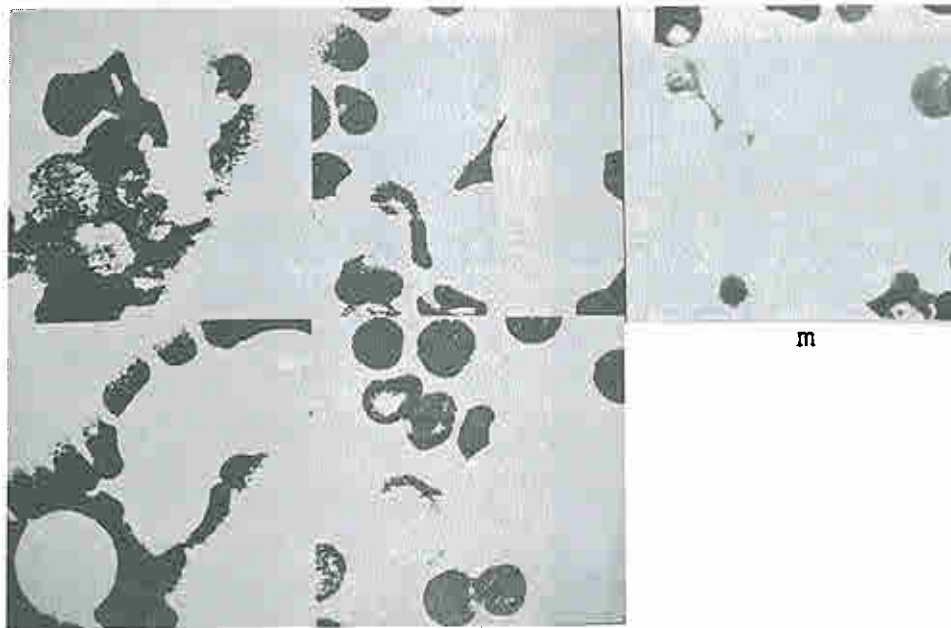
i-l   m

Figure 5.10 The 13 blue images sliced at [Trbc:255]. a)
3-Trbc=110, b) 4-Trbc=110, c) 5-Trbc=110, d) 6-Trbc=108, e)
7-Trbc=110, f) 8-Trbc=110, g) 9-Trbc=110, h) 10-Trbc=105, i)
11-Trbc=110, j) 12-Trbc=110, k) 13-Trbc=125, l) 14-Trbc=114, m)
15-Trbc=110.

5.3 Nucleated erythrocyte extraction

RBC precursor cells (the most common of which are normoblasts) are nucleated cells which appear late enough in the red cell line to contain some hemoglobin. In blue light (419 nm), the hemoglobin in such cells appears less dense than in RBCs but denser than in leukocytes. It often appears as a coarse texture rather than a continuous dense region.

Using the RBC hemoglobin threshold and CCnuc (nucleus connected component image), the following algorithm determines which cells i are possibly nucleated RBCs in CChg.

```
"[ 1] Extract hemoglobin containing formed object."
MSrbc<==If (BL Slice [Trbc:255])
        Then 255
        Else 0;

"[ 2] Test nuclei for hemoglobin content"
For all i in CCnuc Do
        If Area(CCnuc(i) & MSrbc) > epsilon
                Then CCnuc(i) is a nucleated RBC;
```

Figure 5.11 shows (BL&CCnucL) and (BL&CCnuc&MSrbc) for frames 126-6 and 126-7 (the latter frame contains several normoblasts while the former does not). Table 5.1 shows the nuclear area and the hemoglobin area (above threshold Trbc) for each nucleus for cells in these images. For epsilon set to 10 square microns, the normoblasts in frame 126-7 are detected correctly while the nuclei in frames 126-6 are correctly ignored.

Having identified the nucleated RBCs, a modified RBC

mask may be created which has the marked nuclei removed. This is computed over (x,y) as

```
MSrbc' <== If ((MSrbc & CCnuc(i) > 0)
               and (i is nucleated RBC)
               Then 0
               Else MSrbc;
```
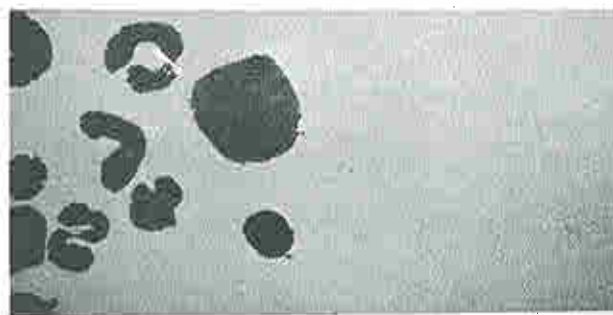

a-b
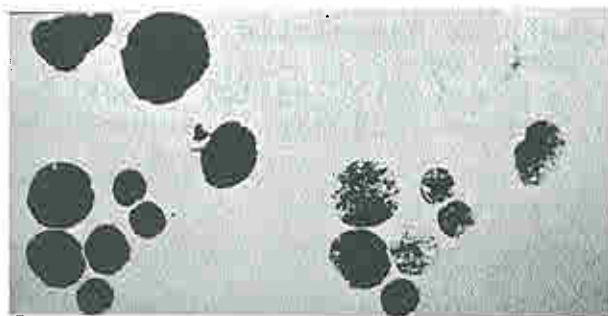
c-d

Figure 5.11 Blue image pixels defined by BL&CCnuc and BL&CCnuc&MSrbc. a) 126-6 (CCnuc&BL), b) 126-6 (CCnuc&MSrbc), c) 126-7 (CCnuc&BL), d) 126-7 (CCnuc&MSrbc).

Frame #	CC# i	Area (CCnuc)	Area (CCnuc(i) &MSrbc))
126-6	2	48.2	0.02
126-6	3	63.8	0.00
126-6	4	160.4	0.10
126-6	5	54.3	0.00
126-6	6	26.2	0.75
126-6	7	40.1	0.00
126-6	8	47.2	0.15
126-6	9	39.5	0.02
126-6	10	32.8	2.37
126-6	11	20.1	0.15
126-7	2	59.3	0.22
126-7	3	112.5	6.22
126-7	4	55.3	31.18
126-7	5	62.0	45.16
126-7	6	17.3	13.66
126-7	7	18.2	14.40
126-7	8	30.8	12.93
126-7	9	47.1	43.39
126-7	10	18.9	18.90

Table 5.1 Hemoglobin content of nuclei. This is measured by the area of the blue image (for each nucleus CCnuc(i)) above Trbc. In order to contrast the nuclei, the areas of the nuclei are given as well. All areas are given in square microns.

6. LARGE OBJECT ANALYSIS

Given the GB image, with RBCs removed, the next step in the segmentation of the frame is to find the nuclear centers of the leukocytes. As these are the darkest large objects in the frame, their segmentation would be trivial except for the fact that leukocytes often clump together in marrow smears. Large dark objects are tested to see whether they are actually clumps of smaller objects and should be broken up into smaller objects. In most cases, the touching dark objects are cell nuclei.

6.1 Touching nuclei: splitting algorithm

The algorithm for splitting touching nuclei starts by performing a boundary-follower segmentation of the GB image. This is done using threshold Tnuc and an area threshold of [10:2000] square microns. This generates the connected component image CCnuc0 and a boundary data file, BD.

Frames 126-6 and 126-7 contain touching nuclei which should be split while frames 126-9, 126-10 contain multilobed poly nuclei which could be mistaken for touching nuclei. The splitting algorithm is demonstrated on both types of objects.

Figure 6.1 shows the trace of the boundary follower segmentation of the nuclei and the connected components CCnuc0 for images 126-6, 7, 9, and 10. Frames 126-6 and 126-7 have
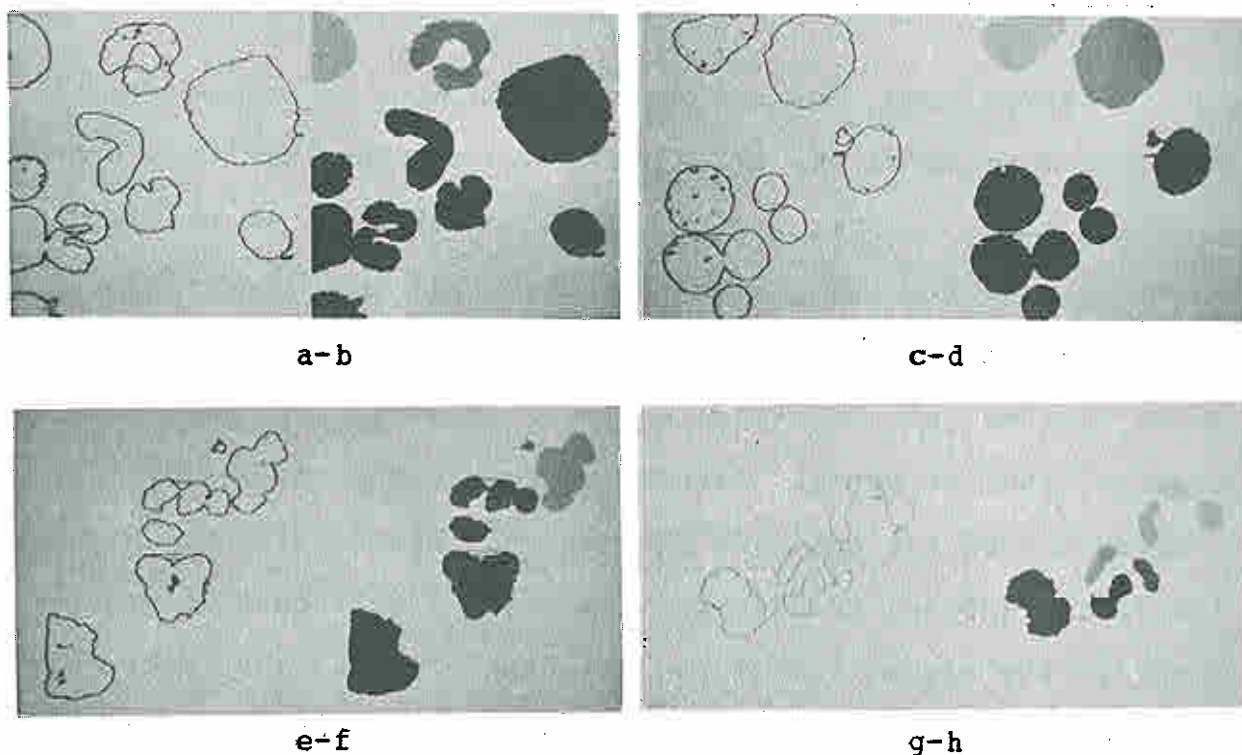
Figure 6.1 Traces and CCnuc0 images produced by the boundary follower segmenter. The CCnuc0 images are scaled by 15 so that they are visible. The threshold was set to Tnuc (60) with area sizing [10:2000] square microns. a) trace 126-6, b) CCnuc0 126-6, c) trace 126-7, d) CCnuc0 126-7, e) trace 126-9, f) CCnuc0 126-9, g) trace 126-10, h) CCnuc0 126-10.

touching nuclei counted as single nuclei. Frames 126-9 and 126-10 do not have any touching nuclei.

The boundary segmentation algorithm, in addition to following the boundary, fills in pixels inside of the object which are below the segmentation threshold. This is done by using a run length map (RLM) representation of the boundary and then testing pixels to determine whether they are inside the object or not. The RLM algorithm is discussed in more detail in [LemP78a].

6.1.1 Testing boundaries for conjugate cracks

After the initial segmentation pass is completed, the boundary data file, BD, is analyzed by testing for conjugate cracks. These cracks are regions which might be created by touching nuclei. The algorithm recognizes the occurrence of conjugate corner pairs. It then creates a list of potential split points of touching nuclei. Figure 6.2 illustrates finding the potential corners which are marked (as 3x3 black rectangles) on the boundary traces to frames 126-6, 126-7, 126-9, and 126-10. These potential corners will then be matched later in the algorithm. Three potential corners were correctly detected in Figure 6.2a and two pairs of potential corners in Figure 6.2b.

6.1.2 Evaluation of texture in potential split region

The texture in the potential split regions is evaluated in the GB image (at the locations of the pixels in CCnuc0) for the cracks specified in the split list. For cracks meeting the split criteria, zeros are painted in a copy of GB in the lightened texture regions of the cracks, thus completing the split procedure. The new image is denoted by GBsplit. Figure 6.3 illustrates the results of applying the texture and zero painting procedures to frames 126-6, 126-7, 126-9, and 126-10. One split in frame 126-6 and two splits in frame 126-7 were performed correctly.

6.1.3 Resegmention of GBsplit

After breaking up nuclear clusters using the splitting algorithm, the GBsplit image is in turn resegmented using the boundary follower segmenter at the same parameters of threshold and area as before to get the new connected component image CCnuc. Figure 6.4 shows the correctly resegmented images for frames 126-6 and 126-7 which required splitting.

a

b

c

d
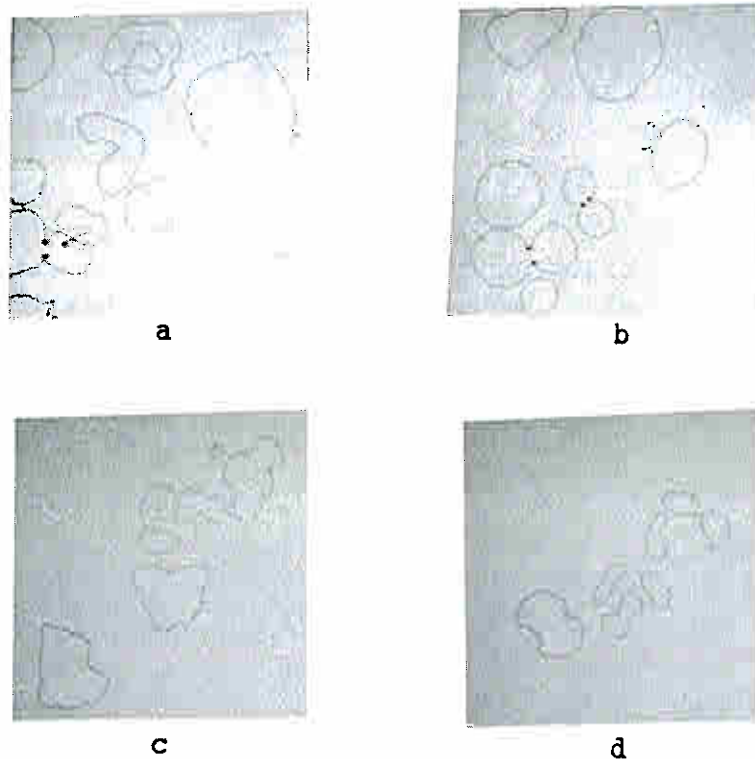
Figure 6.2. Nuclei boundary conjucate crack finding results. The boundary traces for the frames tested are marked with the corners found. The splitting parameters are: range of DAVC=[1:8], window width=7, cornerity threshold=10, split threshold=10. a) marked trace 126-6, b) marked trace 126-7, c) marked trace 126-9, d) marked trace 126-10.

a   b

c   d
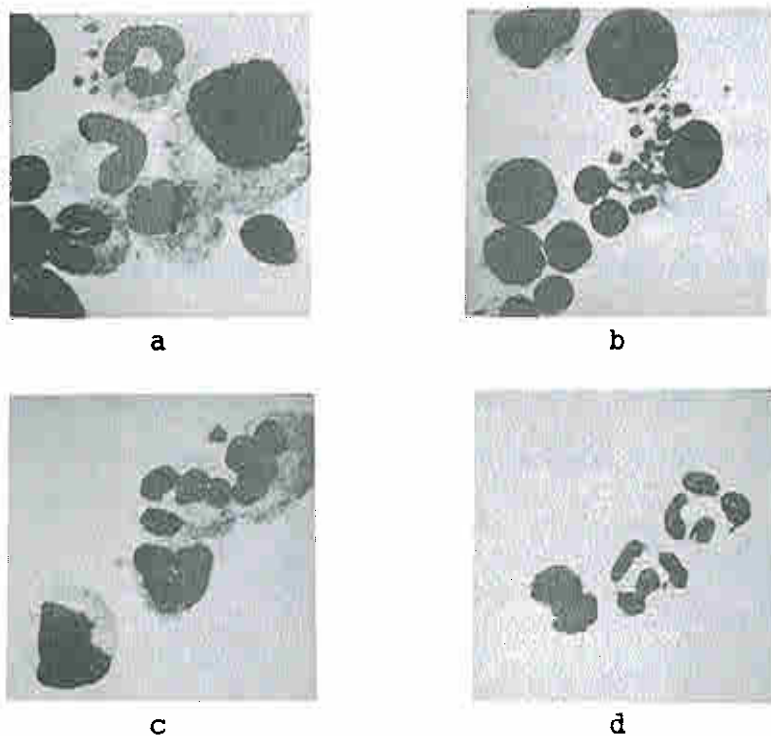
Figure 6.3 Nuclei split after testing the texture of the potential split regions. The texture heuristic threshold=16. a) GBsplit 126-6, b) GBsplit 126-7, c) GBsplit 126-9, d) GBsplit 126-10.

a-b   c-d
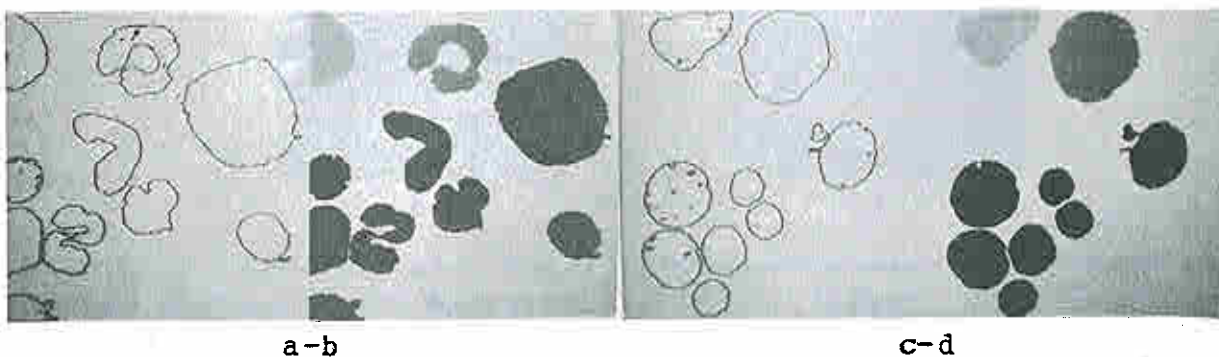
Figure 6.4 Resegmented traces and CCnuc images after splitting of touching nuclei. The CCnuc are scaled by 15 so that they are visible. a) trace 126-6, b) CCnuc 126-6, c) trace 126-7, d) CCnuc 126-7.

6.2 Segmentation of very large objects

The segmentation of very large objects such as megakaryocytes or macrophages cannot be handled with the current segmentation structure since they exceed the frame size and contain multiple nuclei. Some of the nuclei n a macrophage belong to ingested cells as may be seen in frame 127-13 in Figure A.11 in Appendix A. An extended model might contain algorithms to operate over several adjacent frames (as suggested by [SobI76] for tracing neurons across frame boundaries), but this will not be attempted here.

7. SINGLE CELL ARTICULATION

The GBsplit image (which was possibly split using the algorithm discussed in Section 6.1) is now segmented at the single cell level.

7.1 Nucleus articulation

The nuclei are segmented from the GBsplit image using a nuclear threshold Tnuc. It was found that a value of Tnuc=60 gave a reasonable approximation to the valley of the GBsplit density histogram. The area threshold was set to exclude objects less than 10 square microns in size. Fragmented nuclei (such as basophils which are confused because of the granulation) may be excluded using the simple shape feature (perimeter**2/area) for values greater than 80. The nuclei excluded because of fragmentation may be segmented using alternate algorithms (such as the boundary trace transform, BTT(T), to be discussed).

Figure 7.1 shows the boundary trace and CCnuc images for frames which have both touching (frames 126-6, 126-7) and non-touching (frames 126-3, 126-4) nuclei. Table 7.1 lists some of the features computed during the segmentation of the nuclei of frame 126-6.
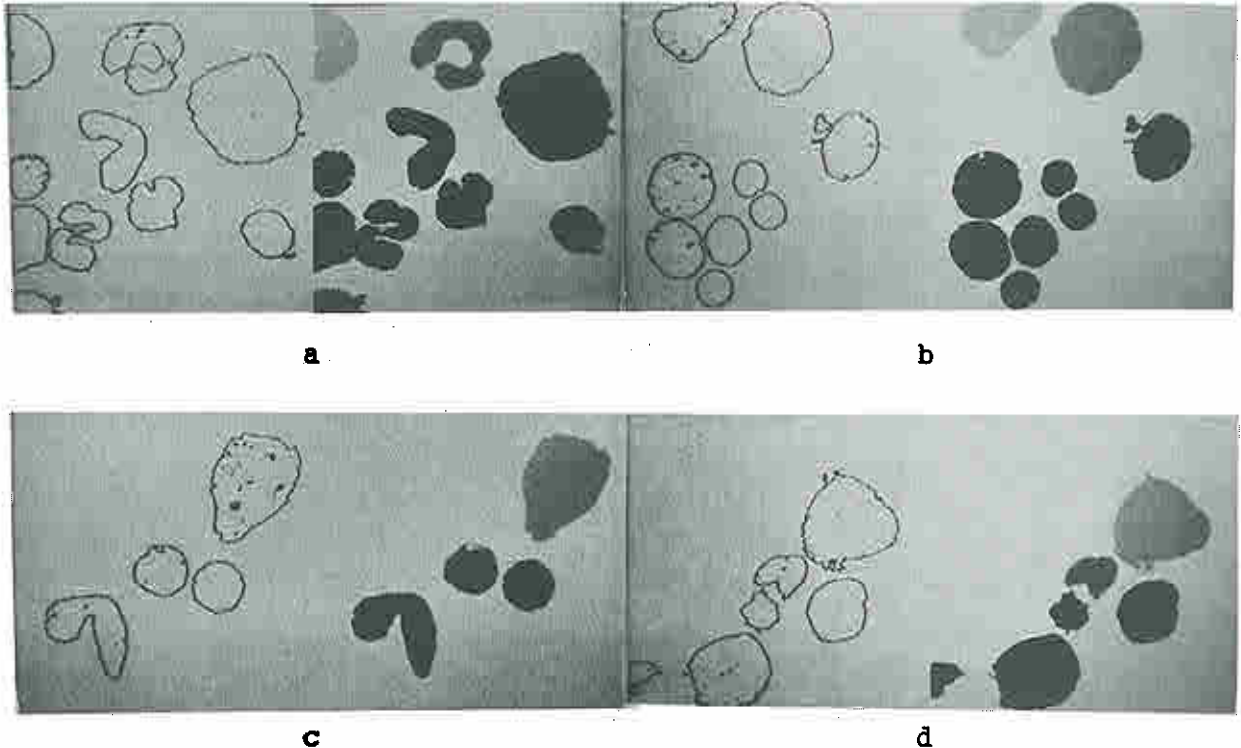
Figure 7.1 The boundary trace and CCnuc images for frames a) 126-6, b) 126-7 which contain touching nuclei. Frames c) 126-3, d) 126-4 which do not contain touching nuclei.

7.1.1 Use of the BTT for nuclear boundary enhancement

The BTT(T) may be used to enhance nuclear boundaries. When very large values of (perimeter**2/area) are found for an object, it is an indication that that object is fragmenting at the current segmentation threshold. The BTT(T) algorithm attempts to find the optimal border of an object at various thresholds so that the composite boundary can be segmented. Figure 7.2 shows BTT(T) and BTT(F) for the thirteen frames. Some of the frames have nuclei with boundary fragmentation (Figures 7.2b,c,f,j,k,l,m) while the other frames have boundaries which are very strong (Figures 7.2a,c,d,e,f,g,h,i,l). In addition holes appear both inside and

CC#	Area	Perim	P*P/A	BndPts	Filled-A	Dens/U	Edges
2	48.2	30.2	18.9	173	0.2	4024	2
3	63.2	54.8	47.1	291	2.5	3700	0
4	160.4	56.6	19.9	301	1.0	4019	0
5	54.3	38.4	27.1	207	0.1	4269	0
6	26.2	22.4	19.1	126	0.7	4244	1
7	40.1	28.0	19.4	153	0.1	5255	0
8'	87.1	73.8	62.4	404	0.5	4169	1
8	47.2	50.2	53.4	272	0.5	4101	0
9	39.5	25.4	16.3	147	0.0	4221	1
10	32.8	26.9	21.9	149	0.0	4678	0
11	20.1	24.0	28.7	141	0.8	3608	2

Table 7.1 Features computed for nuclei during nucleus articulation segmentation pass for frame 126-6. Features are given in microns where applicable. The Edges feature is the number of edges of the object which touch the frame. The BndPts feature is the length of the boundary in pixels. Area sizing is [10:2000] square microns, density threshold=60. CC's 8 and 9 were previously CC# 8 which is listed in the table as 8'.

outside of the boundaries as they fragment as a function of threshold.

The boundary in the case of the eosinophil in 126-8 has filled regions surrounding the real nuclear edge which might be used by the global model as evidence of fragmentation caused by cytoplasm granulation. Note that the polys in 126-8 are correctly detected. An eosinophilic metamyelocyte in 126-5 also fragments for similar reasons. In such cells, two boundaries occur with granules appearing between the two boundaries. This thus leads to a procedure for testing the occurrence of such heavily granulated cells.
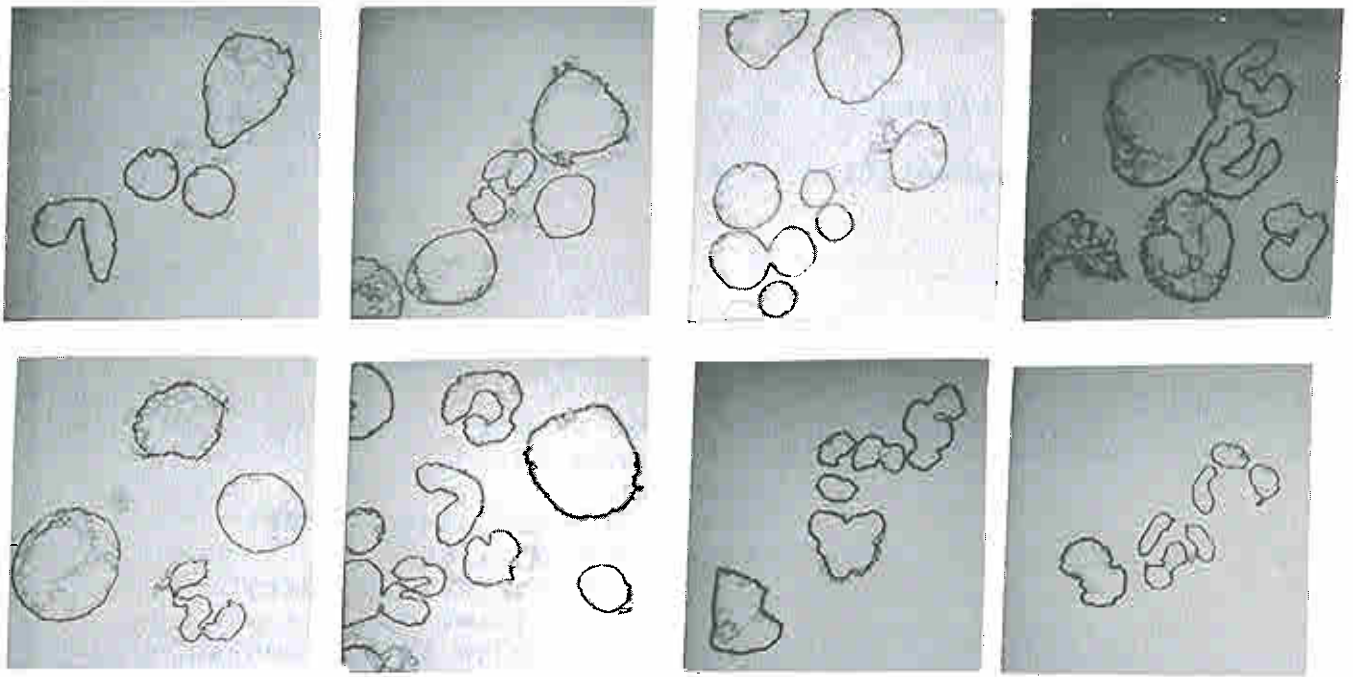
a-d   e-h
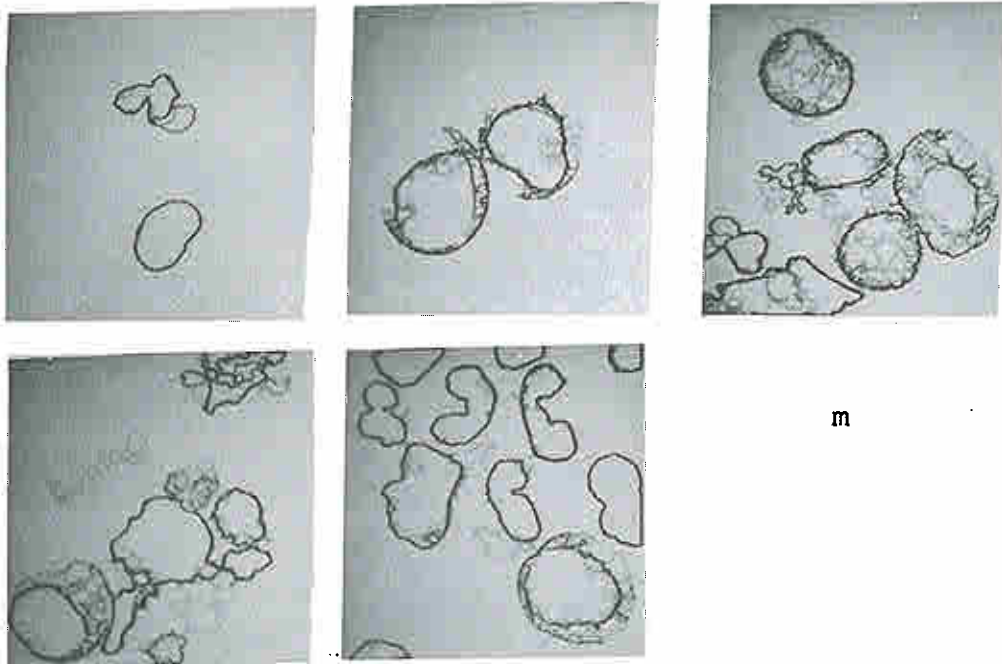
i-l   m

Figure 7.2 BTT(T) for the thirteen frames. a) 126-3, b) 126-4,
c) 126-5, d) 126-6, e) 126-7, f) 126-8, g) 126-9, h) 126-10, i)
127-11, j) 127-12 k) 127-13, 1) 127-14, m) 127-15.

The histogram of BTT(T) can be used to indicate the degree of fragmentation of the boundary using the following algorithm:

```
"[1] Compute the BTT(T) masked by the standard CCnuc."
        BT60 = CCnuc(Tnuc=60) & BTT(T);

"[2] Test each nucleus of frame"
        For each i in CCnuc Do
                Begin
                Compute HIST(CCnuc(i)&BT60);
                If very large values
                        Then strong boundary
                            (i.e. edge is ok)
                        Else edge is fragmented
                        due to cytoplasm
                        granulation.
```


a   b
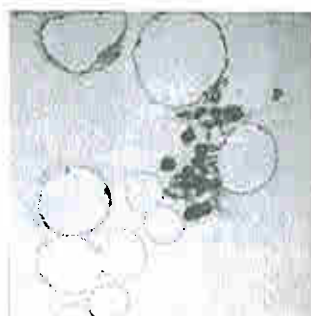
c   d

Figure 7.3 GBsplit' image computed by subtracting the nuclei from the GBsplit image. Frames a) 126-3, b) 126-4, c) 126-6, d) 126-7.

7.1.2 Removal of nuclei from the frame after segmentation

After the nuclei are extracted, they are subtracted from GBsplit, producing the GBsplit' image shown in Figure 7.3 with the nuclei correctly removed except for the case of the eosinophil in the lower right hand corner. Having removed the nuclei, the other regions can now be analyzed.

7.2 Cytoplasm articulation

After the nuclei are removed, the edge nuclei are removed (cf. Section 7.5) and then the platelets are removed (cf. Section 8.1), resulting in image GBsplit'''. The three formed object connected component images are CCnuc, CCedge, and CCplatelet. The cytoplasm is then extracted by generating the cytoplasm mask MScyto. MScyto isolates the cytoplasm of the frame from background, platelets and nuclei. The following algorithm computes MScyto.

```
"[1] Compute the mask of the formed objects."
For all (x,y) Do
        If (CCnuc or CCedge or CCplatelet)
              Then MSformed(x,y)_255
              Else MSformed(x,y)_0;

"[2] Smooth GBsplit."
CYm1 <== AVG8(GBsplit);

"[3] Extract approximate cytoplasm regions."
CYm2 <== If (CYm1 Sliced [Tcyto-5:Tnuc-1]) neq 0
              Then Normalized(GN(x,y))
              Else 0;

"[4] Smooth the extracted cytoplasm region."
CYm3 <== AVG8(CYm2);

"[5] Remove the cytoplasm found in formed object
 regions."
CYm4 <== CYm3 - MSformed;

"[6] Rescale CYm4 to the [0:255] range of MScyto"
        MScyto <== If CYm4 neq 0
                        Then 255
                        Else 0;
```

Figure 7.4 shows the cytoplasm mask generation steps for frames 126-6, and 126-7. The MScyto image in Figure 7.4f is a reasonable approximation of the cytoplasm in frame 126-6 with

a-d

e-h

i-l

m-n
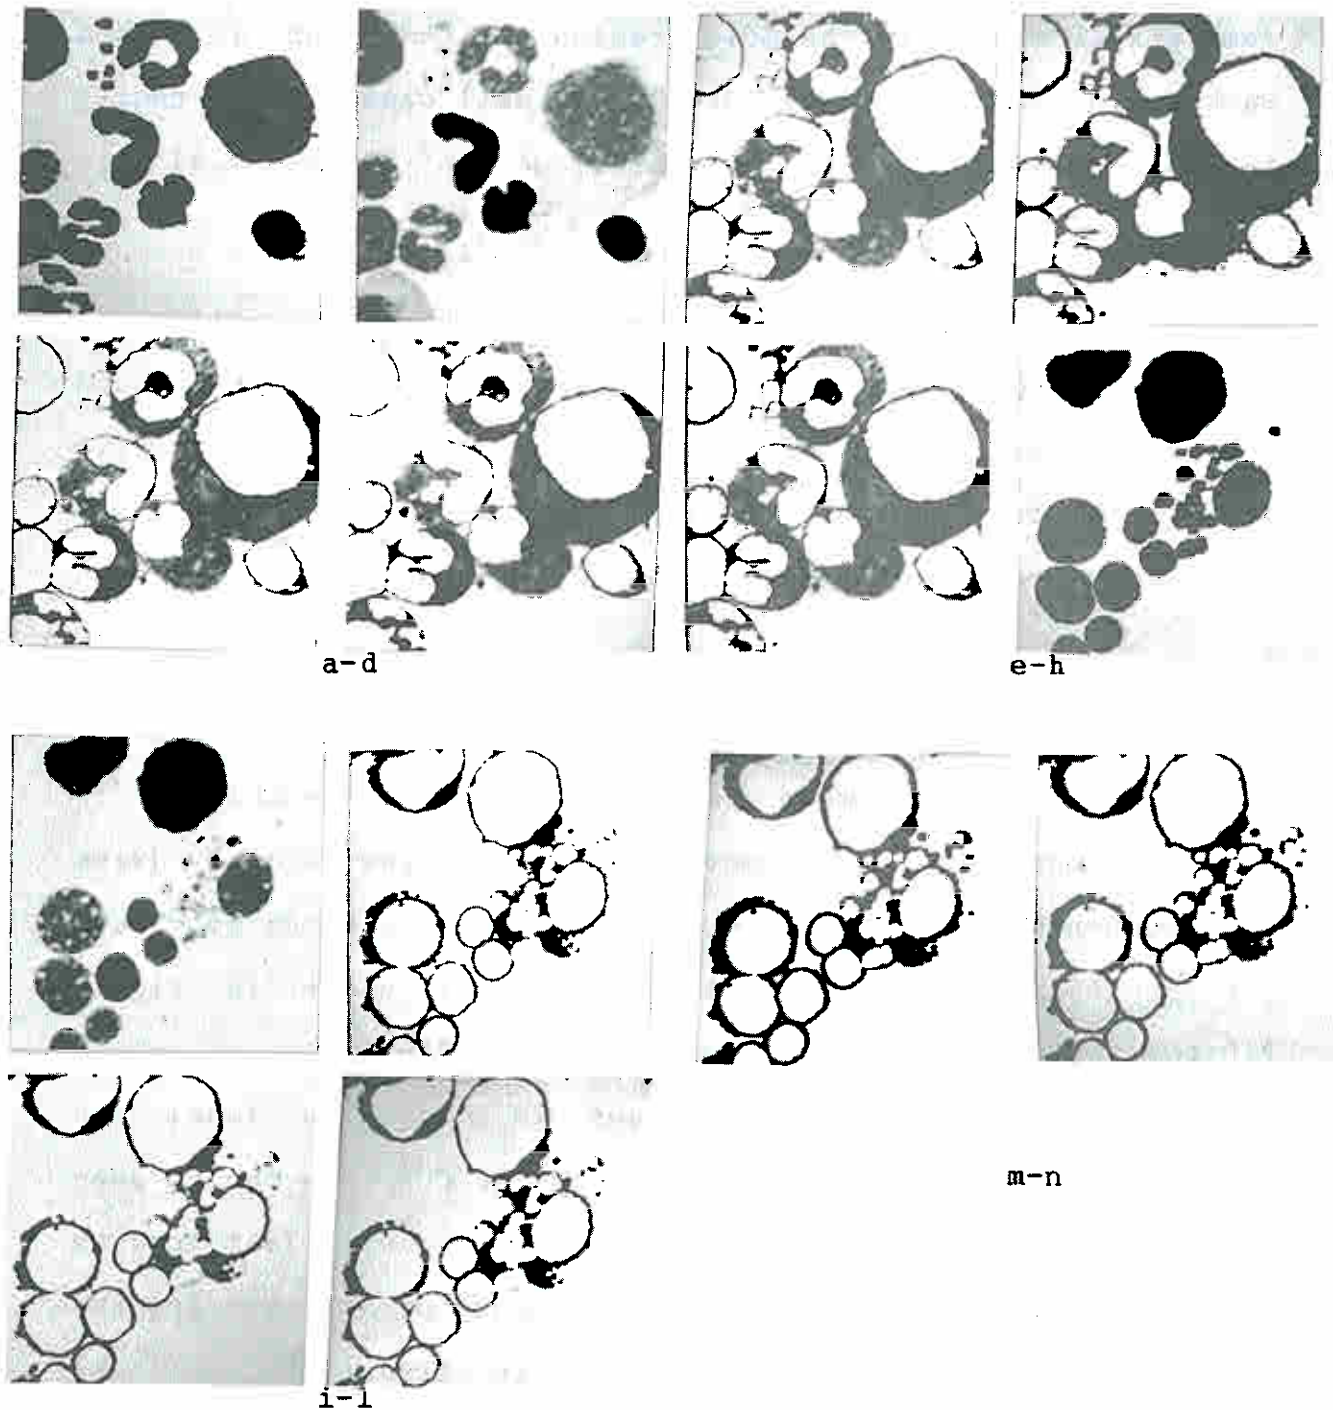
Figure 7.4 Stages in development of the cytoplasm mask for frames 126-6 and 126-7. a) MSformed-6, b) CYm1-6, c) CYm2-6, d) CYm3-6, e) CYm4-6, f) MScyto-6 g) (MScyto-6&GN), h) MSformed-7, i) CYm1-7, j) CYm2-7, k) CYm3-7, l) CYm4-7, m) MScyto-7 n) (MScyto-7&GN-7).

some pixels missed in the Golgi region of the stab and some background pixels at the lower right cell cluster incorrectly labeled as cytoplasm. The MScyto image in Figure 7.41 was approximately correct.

If the cytoplasm regions for each cell are disjoint, then by propagating the nuclear connected component values (CCnuc) into adjacent cytoplasm, it is possible to label cells with the correct nuclear component number. In the cases where this is not true, additional steps need to be taken.

7.2.1 Cytoplasm isolation algorithm

The MScyto image (ranging 0 or 255) is rescaled to (0 or 1) to construct the connected component image CCcyto (with one component value 1). CCcyto is ORed with CCnuc to create the pre-isolation image CNcell. This image is shown in Figure 7.5 for frames 126-3, 126-4, 126-6, and 126-7. The cytoplasm will later acquire nuclear CC values in order to label the regions (nucleus+cytoplasm) as cellular CCs. As can be seen from all the CNcell images in Figure 7.5, the cytoplasm touches multiple nuclei. Therefore, some means of distributing the cytoplasm to the appropriate nuclei is required.

The cytoplasm isolation algorithm finds (if possible) the cytoplasm-cytoplasm edges along the lines between the centers of all nuclei in CCnuc. Given the original estimate of the edge as being perpendicular to the line between the centers, the algorithm tries to optimize this between +/- 22
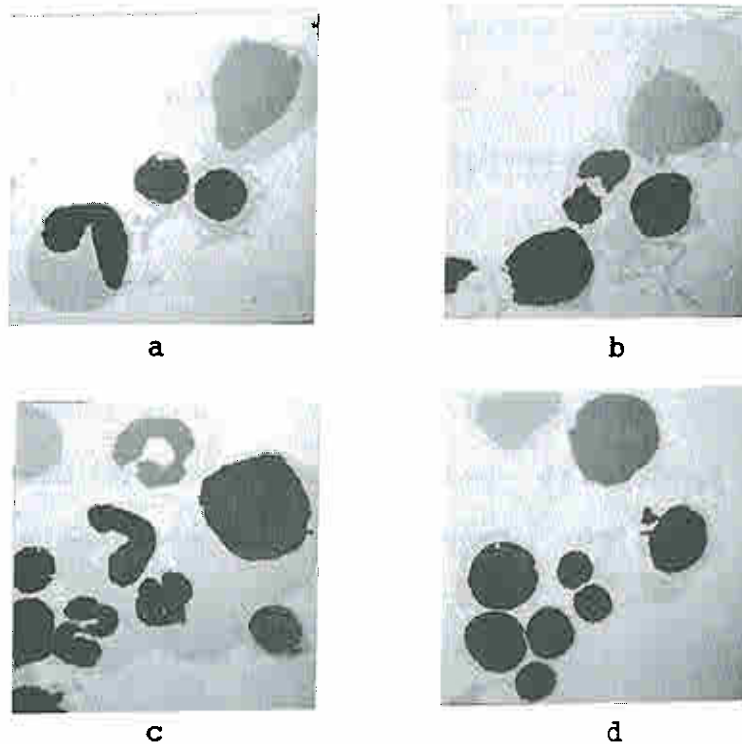
Figure 7.5 CCcyto ORed with CCnuc to create the pre-isolation images CNcell. The images are scaled by 15 so they are visible. a) 126-3, b) 126-4, c) 126-6, d) 126-7.

degrees of the perpendicular. (The angle correction was restricted to prevent it from getting lost on another local minimum.) It then finds the (suboptimal) direction of the intersection and paints a zero line through the touching region to isolate the cytoplasm regions.

Figure 7.6a shows the connectivity graph for all nuclei centers for frame 126-6. Figure 7.6b superimposes this on CCnuc (scaled by 15 so as to be visible). Figure 7.6c shows the reduced connectivity graph (thin lines) with painted regions (thick lines) used to effect isolation. Notice the drastic reduction of nucleus-nucleus pairs required to be tested in the second half of the algorithm. Figure 7.6d shows the result of

the isolation algorithm, the pre-propagation image, CPcell. Note that the cytoplasm regions are labeled with 1's, the background with 0's, and nuclei with values greater than 1. Several of the cytoplasm-cytoplasm edges were correctly isolated. The failure of the large myelocyte/polychromatic normoblast is due to the lact of an adequate light region in the cytoplasm interface. Alternative isolation algorithms might be used in such cases. Figure 7.7 shows the reduced connectivity graphs with painted regions for frames 126-3, 126-4 and 126-7 where the algorithm met with variable success.

a

b

c

d
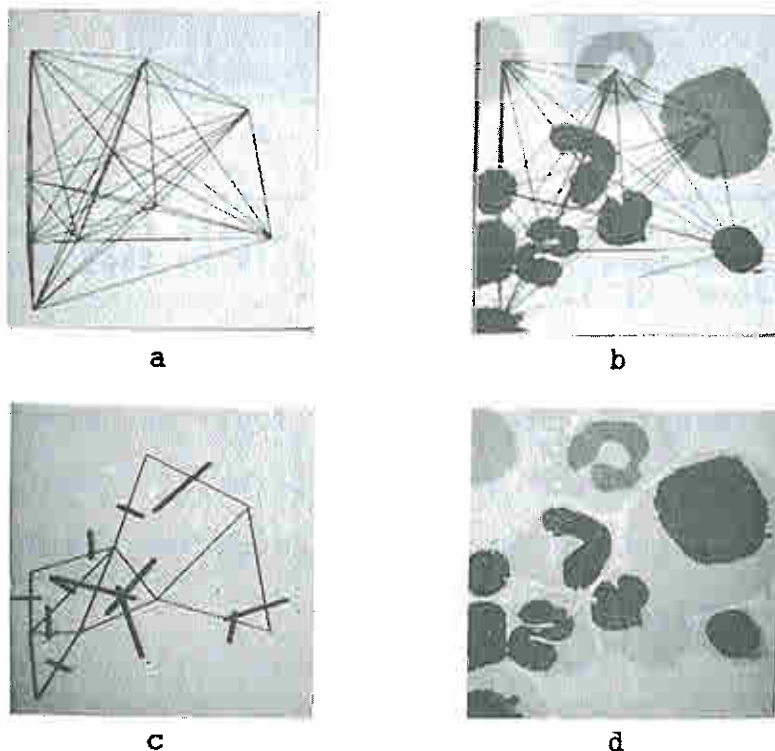
Figure 7.6 Connectivity graphs for cytoplasm isolation for frame 126-6. a) Complete nuclear center connectivity graph, b) superposition of (a) with CNcell scaled by 15 so visible, c) reduced connectivity graph (thin lines) with painted regions (thick lines), d) CPcell.

After CPcell is computed, the nuclei are propagated

into the cytoplasm regions using either of two methods. The first method assumes that all cytoplasm regions were

a

b

c
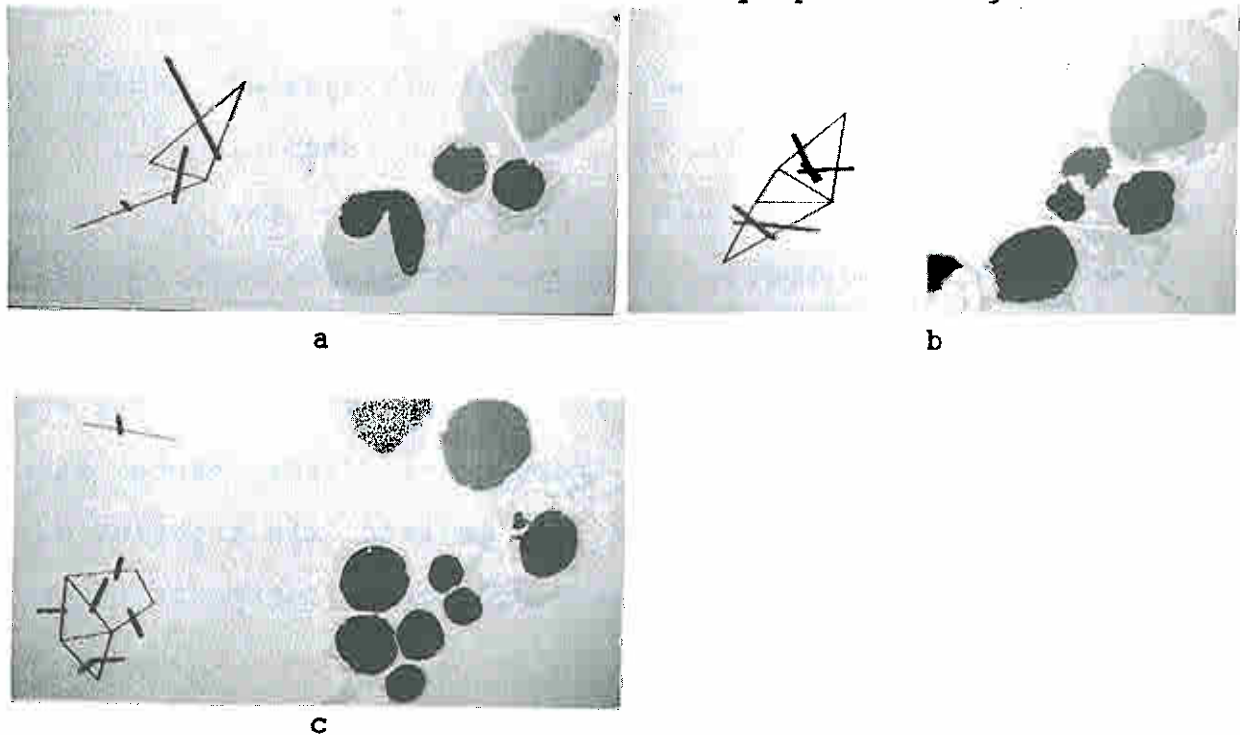
Figure 7.7 Reduced connectivity graph (thin lines) with painted regions (thick lines) used to effect the isolation for frames with CNcell images painted with zeros a) 126-3, b) 126-4, c) 126-7.

successfully isolated (which is not always the case). The algorithm works by iteratively moving in a clockwise contracting spiral clockwise to the midpoint of the image, and then in an expanding counterclockwise spiral back to the outside edge, propagating $g(x,y)=1$ values into 8-neighbor adjacent values of g that are greater than 1 as the image is traversed.

The second method propagates one pixel/pass taking considerably longer to compute the final image CCcell which results when there are no propagations in a pass. (The first propagation algorithm operates until it too has a null pass;

however, it usually converges after one or two passes.) The first method was used in the present study.

Figure 7.8 shows the results of propagating CPcell in order to compute CCcell. As can been seen in 7.8, the propagation mislabeled some pixels because the cytoplasm isolation was incomplete. This was especially true for frame 126-3 (7.8a) where the cytoplasm boundary was indistinct and the isolation algorithm split the cell in the middle. The same type of failure is observed in frame 126-6 (7.8c) where there was not enough texture edge to isolate the touching cell region. Other cases in 7.8a-d were propagated correctly because

a   b

c   d
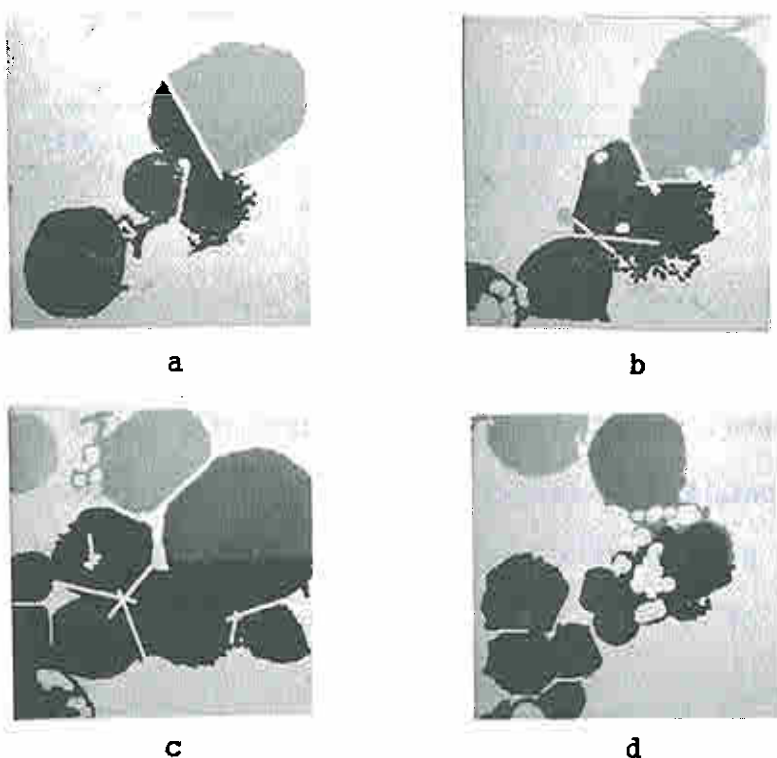
Figure 7.8 CCcell computed by propagating CPcell using the spiral propagation algorithm for frames a) 126-3, b) 126-4, c) 126-6, d) 126-7.

of adequate cytoplasm isolation.

Figures 7.9, 7.10, 7.11 and 7.12 show individual cells extracted from the respective green images for frames 123-3,
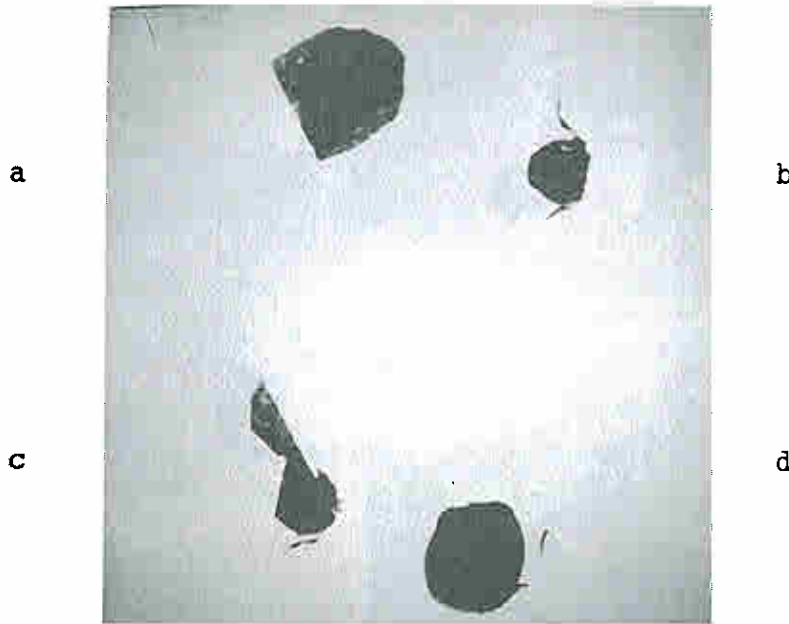
Figure 7.9 Extracted cells from contrast stretched green image using CCcell for frame 126-3. a) CC#2, b) CC#3, c) CC#4, d) CC#5.

123-4, 123-6, and 123-7. The results are what would be expected given the CCcell image in Figure 7.8. Many of the cells were isolated fairly well (Figures 7.9b,d; 7.10a,c,e; 7.11a,b,d,e,g,h,i,j; 7.12a,b,c,d,e,g,h,i). Others had difficulties caused by poor nuclear definition, or by cytoplasm-cytoplasm isolation (Figures 7.9a/c; 7.11c/f, 7.11g/j). Figures 7.10b,d are the two lobes of a poly which had teardrop nuclear fragments. These were segmented as separate cells. The platelets and background in frame 126-7 caused problems with the cytoplasm mask generation in several of the cells in Figure 7.12.

a-d

e-f
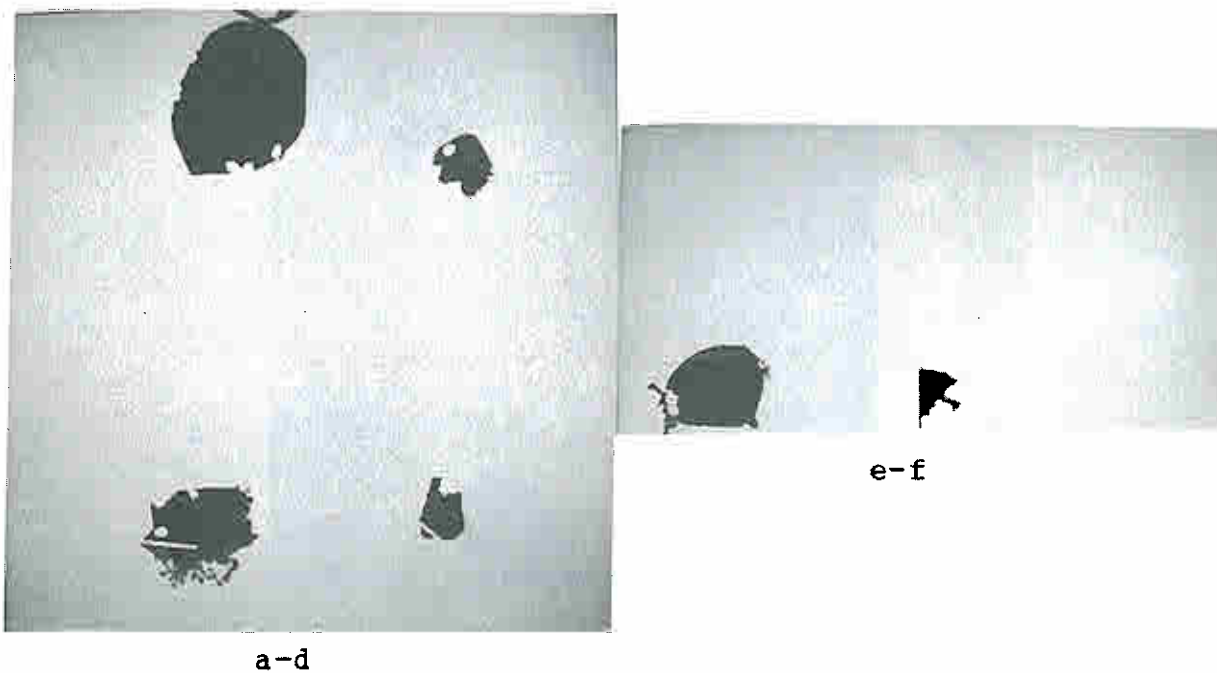
Figure 7.10 Extracted cells from contrast stretched green image using CCcell for frame 126-4. a) CC#2, b) CC#3, c) CC#4, d) CC#5, e) CC#6, f) CC#7.
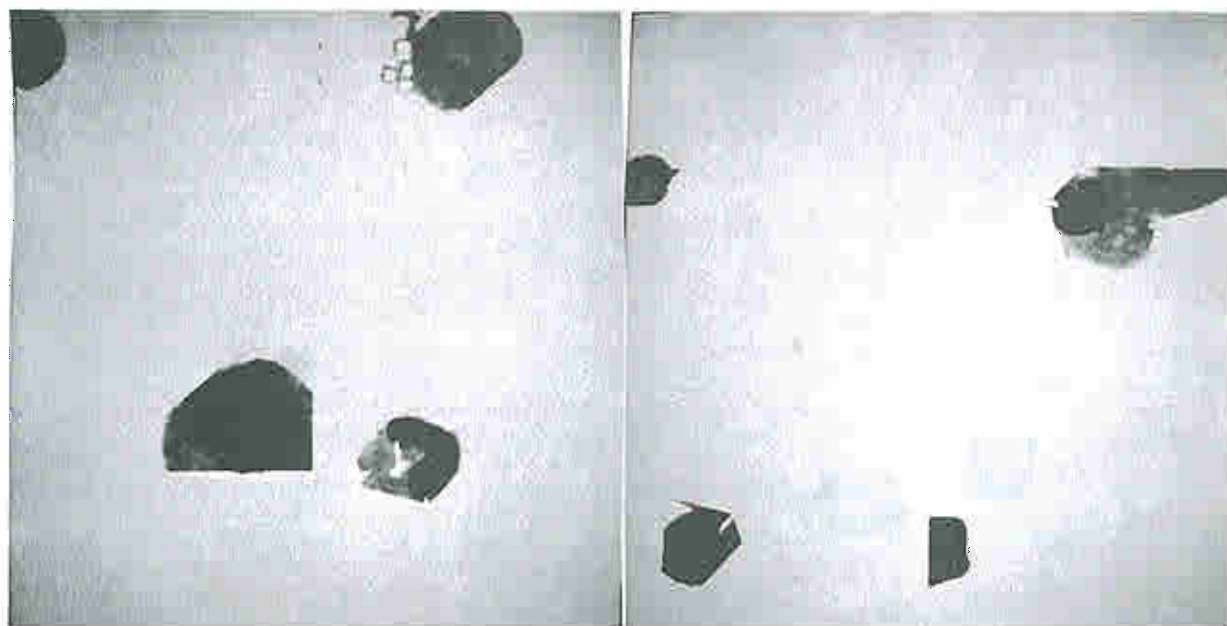
a-d

e-h
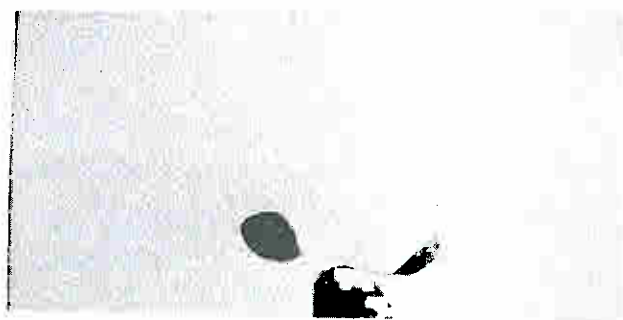
i-j

Figure 7.11 Extracted cells from contrast stretched green image using CCcell for frame 126-6. a) CC#2, b) CC#3, c) CC#4, d) CC#5, e) CC#6, f) CC#7, g) CC#8, h) CC#9, i) CC#10, j) CC#11.
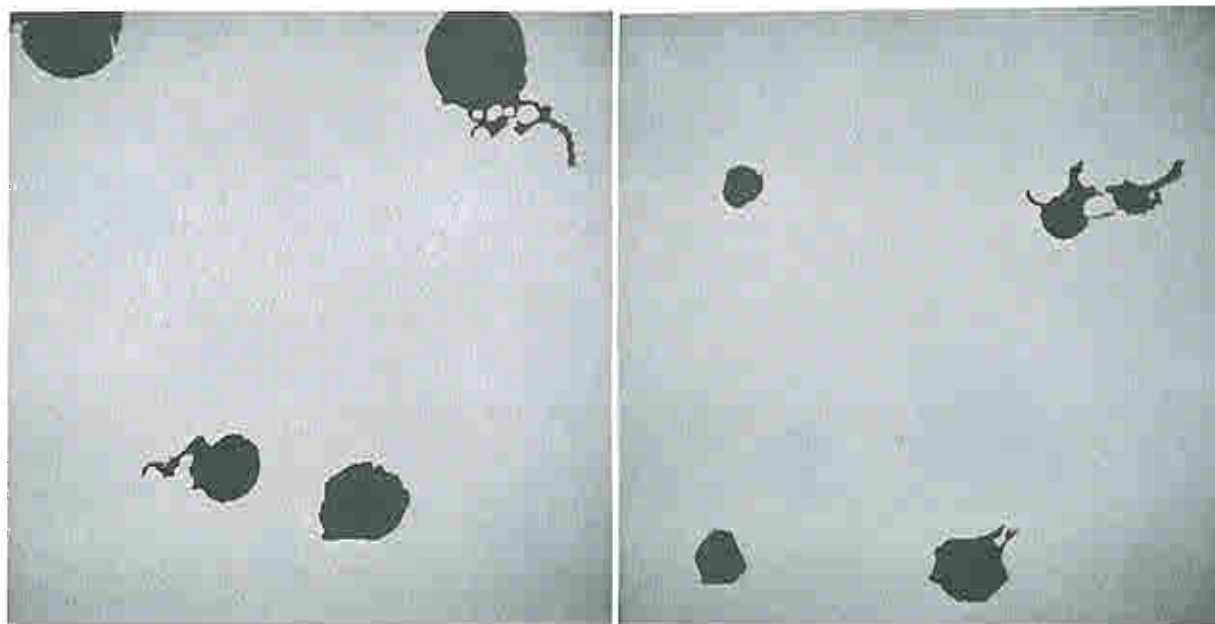
a-d   e-h

i

Figure 7.12 Extracted cells from contrast stretched green image
using CCcell for frame 126-7. a) CC#2, b) CC#3, c) CC#4, d)
CC#5, e) CC#6, f) CC#7, g) CC#8, h) CC#9, i) CC#10.

7.3 Nucleolus articulation

After CCnuc is computed, it is possible to articulate nucleoli. Nucleoli are detected using either a threshold slicing algorithm or a BTT(F) operation to detect large lightened texture regions inside the nucleus. Nucleoli in frames 126-3, 4, 5, and 127-12 are analyzed using the slicing algorithm here. The same frames were also analyzed using the BTT(T). The two algorithms appeared to do equally well in detecting nucleoli.

The threshold slicing algorithm for segmenting nucleoli uses a second segmentation of the nucleus to look for light round regions only inside of the nucleus. It is given as

```
"[1] Get nucleus region i gray scale image."
     GBnuc <== GBsplit & CCnuc(i);

"[2] Threshold the light nuclear regions."
     GBnuc' <== GBnuc Sliced at [Tnuc+10:255];

"[3] Get potential nucleolus regions in CCnuc(i)."
     GBnucleoli <== GBnuc' - GBnuc;

"[4] Segment with sizing set for nucleoli."
     CCnucleoli <== Segment(GBnucleoli, Thr=60,
                               area > 0.75 sq. microns,
                               (10<(P*P/A)<21));
```

Figure 7.13 shows nucleoli segmented using the slicing secondary segmentation algorithm for frames 126-3, 126-4, 126-5 and 126-12. The algorithm detected nucleoli in Figure 7.13a although it was fooled by strong chromatin texture in Figure 7.13m. The nucleolus in the early stab was not detected but this might be because it was not well defined in the image.
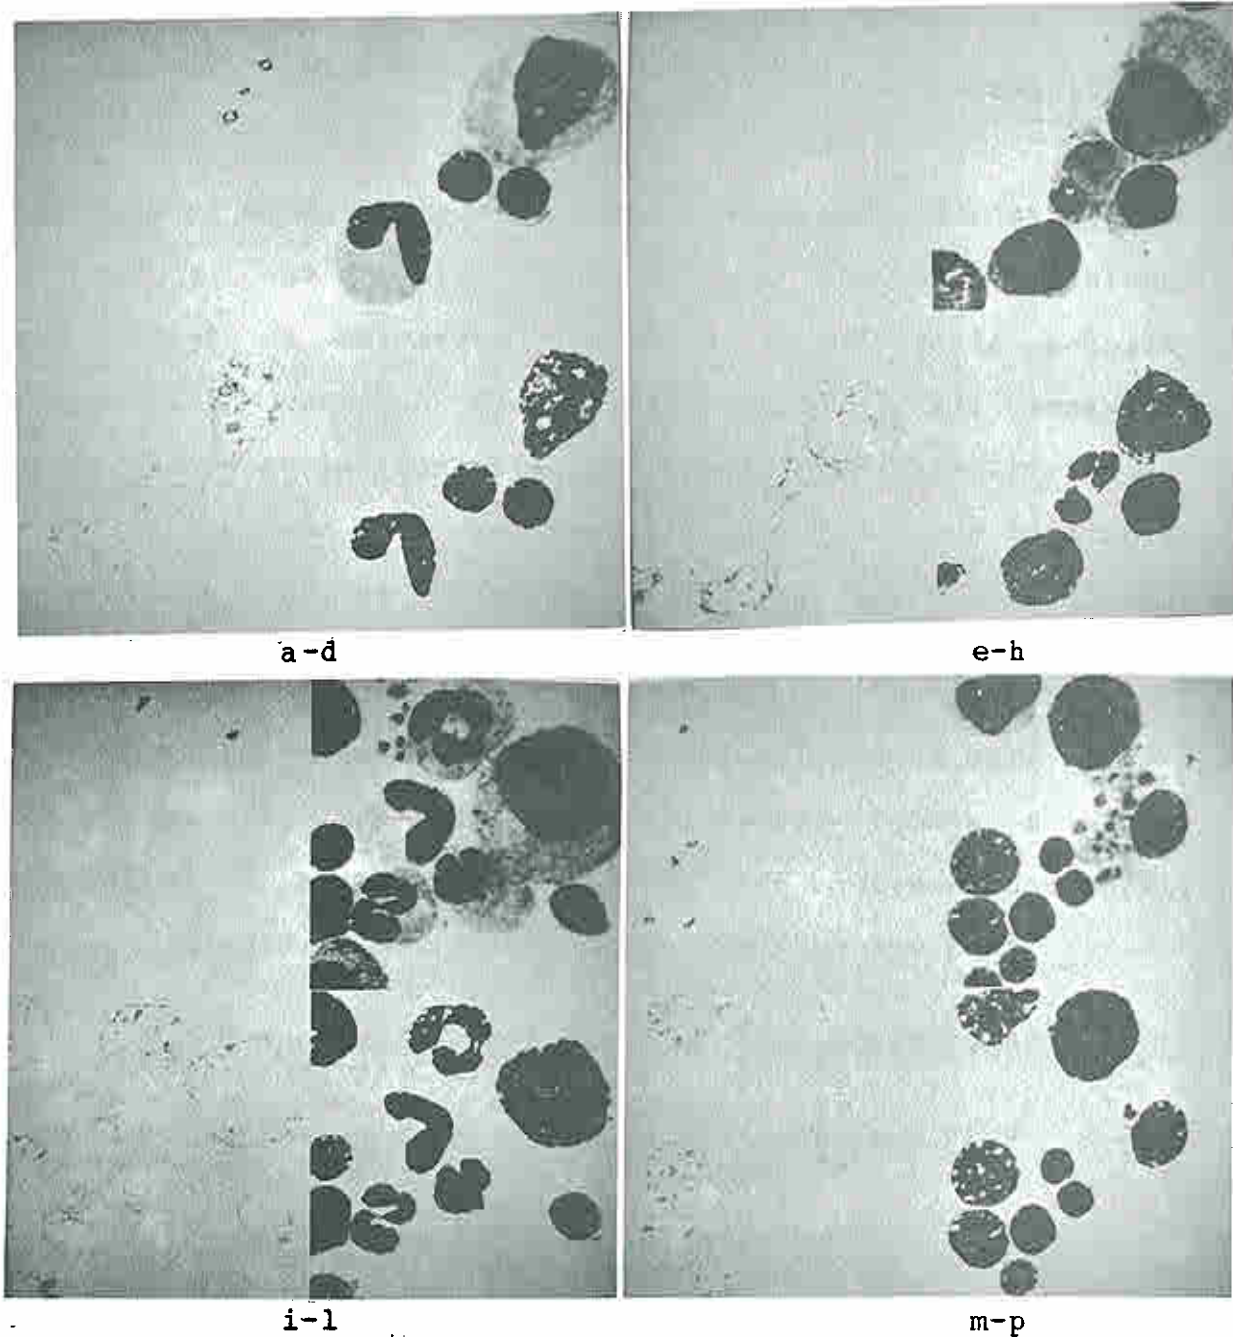
a-d

e-h

i-l

m-p

Figure 7.13 Nucleoli segmented using the threshold slicing secondary segmentation algorithm. The boundary trace is produced when GBnucleoli is segmented with T=60, area > 0.75 square microns, and [10<P2A<21]. a) boundary trace-3, b) GBsplit-3, c) GBnucleoli-3, d) GBnuc'-3, e) boundary trace-4, f) GBsplit-4, g) GBnucleoli-4, h) GBnuc'-4, i) boundary trace-5, j) GBsplit-5, k) GBnucleoli-5, l) GBnuc'-5, m) boundary trace-12, n) GBsplit-12, o) GBnucleoli-12, p) GBnuc'-12.

The BTT(F) algorithm may be used to enhance nucleoli so that they may be segmented.

Figure 7.14 shows the BTT(F) for the thirteen frames. Note the enhancement of the nucleoli in the frames which have them. Notice the dark regions where the nucleoli occur. Using either algorithm should result in potential nucleoli. Additional work needs to be done in refining the procedure.

a-d

e-h

m

i-l
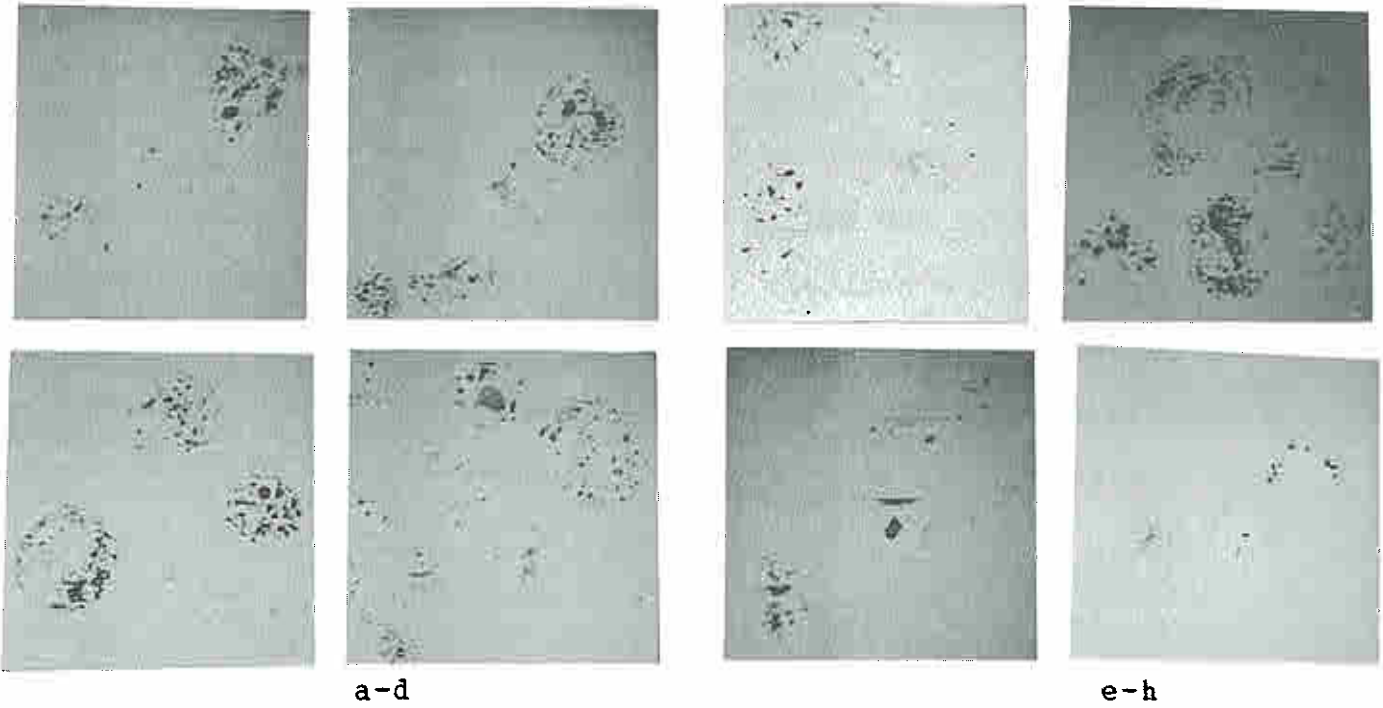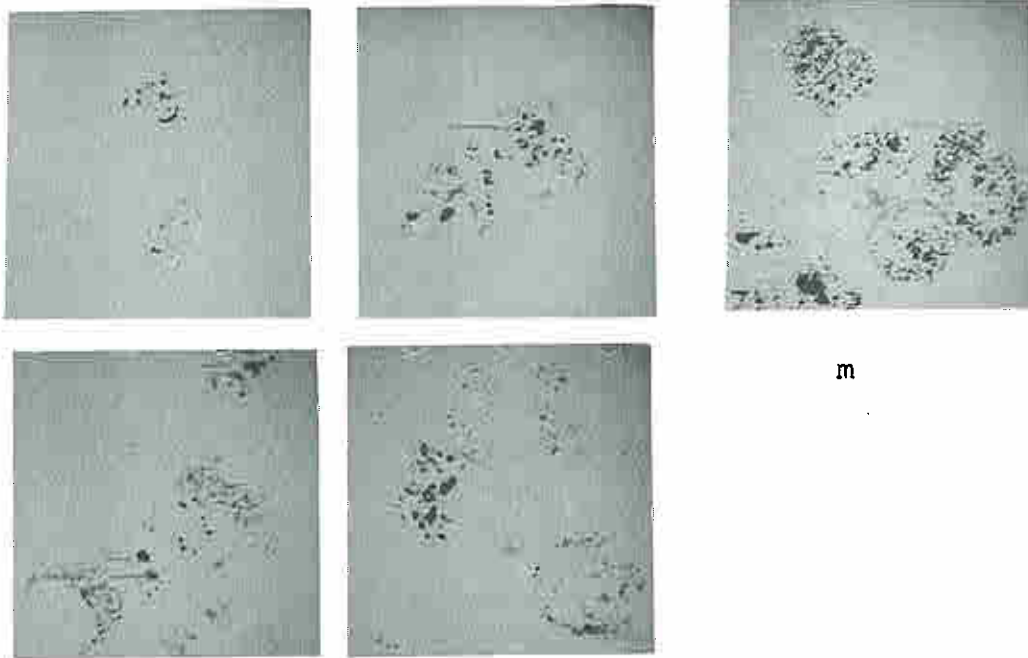
Figure 7.14 BTT(F) for the thirteen frames. a) 126-3, b) 126-4, c) 126-5, d) 126-6, e) 126-7, f) 126-8, g) 126-9, h) 126-10, i) 127-11, j) 127-12 k) 127-13, l) 127-14, m) 127-15.

7.4 Golgi region articulation

The Golgi region is seen in myelocytes up to the stab stage of maturation. It appears as a lightened region adjacent to where the nucleus indents. No operational algorithm is given here to articulate the Golgi region, but indications of positive results from several experiments are discussed.

Although the experiment has not been performed, it seems reasonable to use the BTT(F) on the cytoplasm region of the cell to detect lightened regions.

The Golgi region in the stab in frame 126-6 caused problems with the isolation algorithm in Figure 7.6c. The Golgi region was misidentified as the lightened touching cytoplasm region. Another possible method might be to sweep a line filter around the center of the cell (only in the concavity, if it can be adequately detected) and look at the variance of cytoplasm density.

A third possible method might be to look at the clumping of pixels for thresholds less than Tcyt inside of the cytoplasm region.

7.5 Small nuclear fragments touching the frame

The segmentation of large nuclear fragments which touch the frame is performed after CCnuc is computed and was used to remove the nuclei from the frame. Nuclear fragments less than 10 microns in area can be extracted by segmenting GBsplit' at the nuclear threshold Tnuc using the conditions of area less than 10 microns and one or more frame edges touching the object. Figure 7.15 shows several edge fragments. Figure 7.15a contains two small dark artifacts while 7.15b contains a nucleus fragment successfully extracted.

| a | b |

Figure 7.15 Small edge fragments extracted by segmenting GBsplit' with threshold Tnuc, area sizing less than 10 square microns, and restricting all objects to touching at least one edge. a) CCedge-6 consisting of fragments from nuclei which were touching edge but were greater than 10 microns., b) CCedge-7 consisting of nucleus split by frame.

7.6 Plasma protein background articulation

Plasma protein is background material which tends to accumulate near cells. The plasma and background image is computed using the following algorithm.

```
MSplasma <== If (CNcell > 1) and (0 < GBsplit < Tcyt)
             Then 255
             Else 0;
```

Figure 7.16 shows the plasma background of Frames 126-6 and 126-7 scaled to black. Note the increased density of the background in the lower part of Figure 7.16a which caused problems in the MScyto definition.

a   b
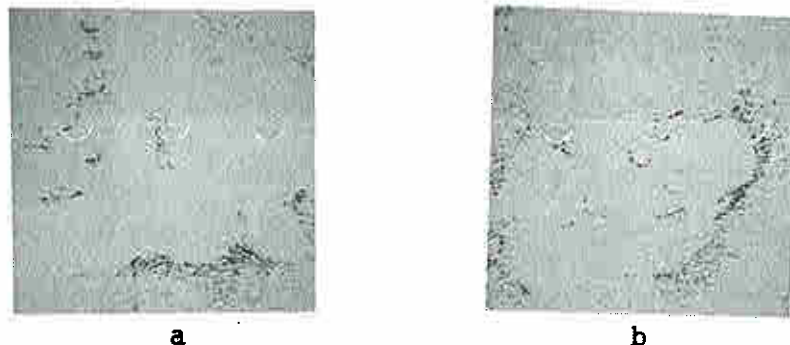
Figure 7.16 Plasma background defined by MSplasma. a) 126-6, b) 126-7.

8. SMALL OBJECT ANALYSIS

Small objects are defined to be objects other than cells or nuclei which are not touching the frame. These include platelets, cell fragments and artifacts.

8.1 Platelet articulation

Platelets are segmented from the GB image using the boundary follower with sizes between 0.25 and 9.9 square microns at threshold Tnuc. The resulting connected component image, CCplatelet0, is then expanded by two pixels to get a better estimate of the platelets, called CCplatelet. Figure 8.1 shows platelets correctly extracted from some of the frames that have large numbers of platelets (frames 126-5, 126-6, and 126-7). Figures 8.1a, 8.1e and 8.1i show GBsplit''. Figures 8.1b, 8.1f and 8.1j show GBsplit''' with the platelets removed, Figures 8.1c, 8.1g, and 8.1k show CCplatelet0, and Figures 8.1d, 8.1h, and 8.1l show CCplatelet. Figures 8.1c, 8.1f, and 8.1i show GBsplit''' computed by subtracting CCplatelet from GBsplit'' (i.e. GBsplit - CCnuc - CCedge). The platelet detection worked fairly well except for a few dark granules in frame 126-5 which were misidentified as platelets.

Not all of the platelets are segmented well. This is because they vary in their density. As will be seen in [LemP78b], platelets are normally 2 to 4 microns in diameter. The reason for using the rather high area size threshold is that the platelets are sticky and tend to clump together. This

122

a-d

e-h

i-1
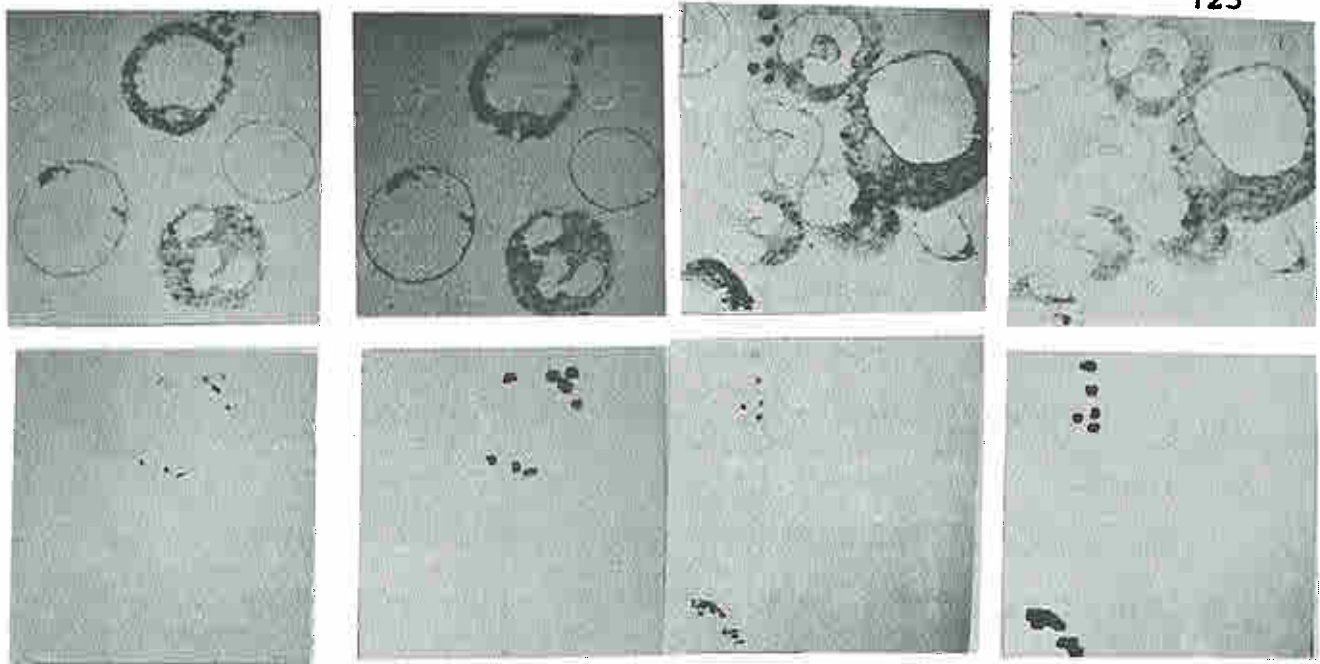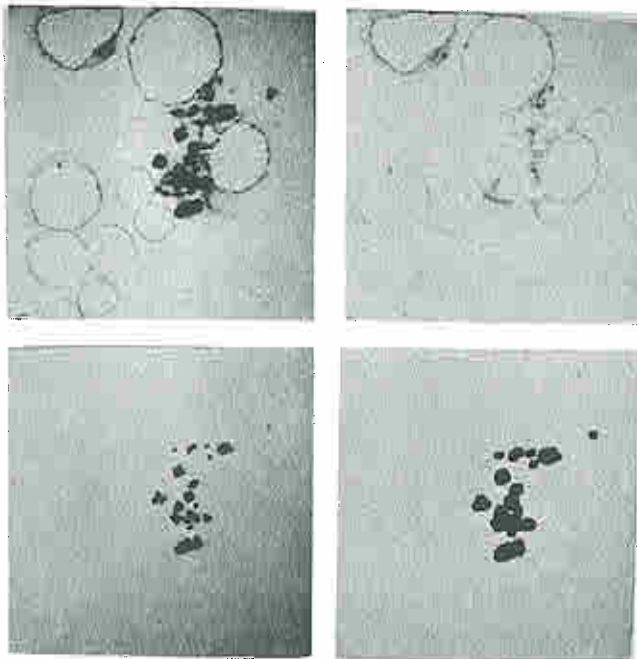
Figure 8.1 Platelets segmented from frames with many platelets.
a) 126-5 GBsplit'', b) 126-5 GBsplit''', c) 126-5 CCplatelet0,
d) 126-5 CCplatelet, e) 126-6 GBsplit'', f) 126-6 GBsplit''',
g) 126-6 CCplatelet0, h) 126-6 CCplatelet, i) 126-7 GBsplit'',
j) 126-7 GBsplit''', k) 126-7 CCplatelet0, l) 126-7 CCplatelet.

means that one or more platelets are being segmented as one object.

Figure 8.2 shows some cytoplasm vacuoles misidentified as platelets in frames without any platelets. Additional tests may be able to prevent this if valid cytoplasm boundaries can be found. The latter is sometimes difficult and the problem is compounded by the fact that platelets tend to cling to white cell boundaries.
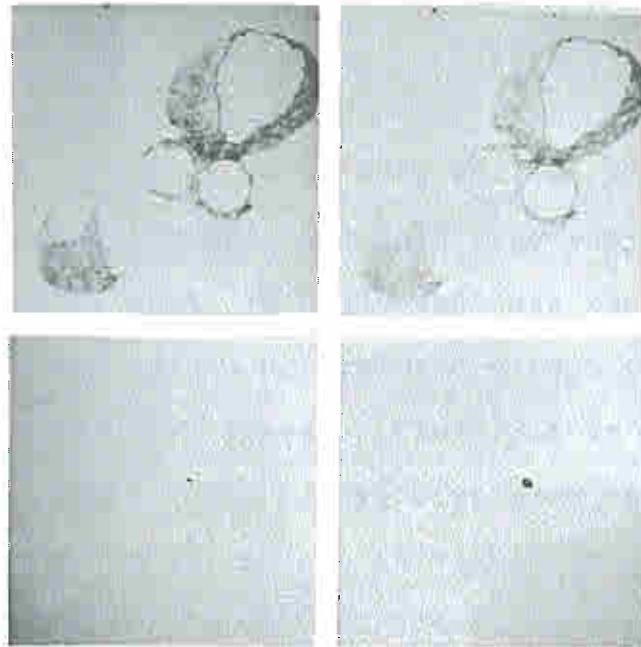
a-d
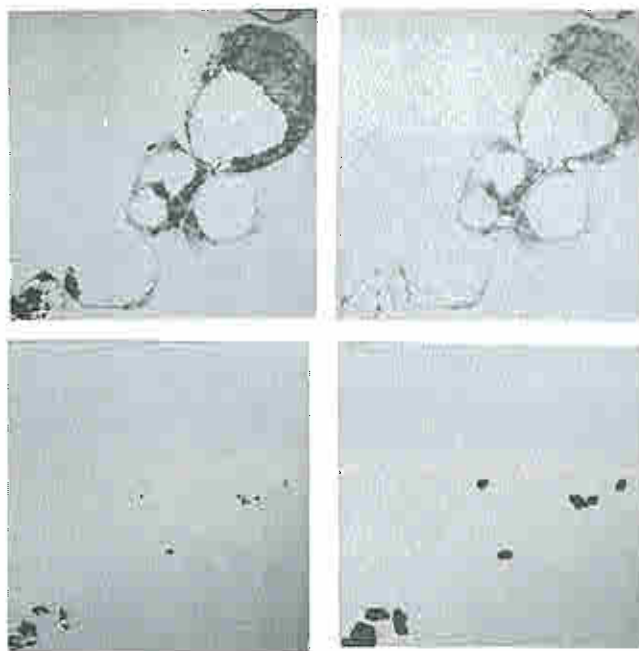
e-h

Figure 8.2 Objects misidentified as platelets. a) 126-3
GBsplit'', b) GBsplit''', c) CCplatelet0 and d) CCplatelet for
dark region where cytoplasm comes together from several cells.
e) 126-4 GBsplit'', f) GBsplit''', g) CCplatelet0 and h)
CCplatelet caused by granules in cytoplasm, dirt and
fragmentation of the eosinophilic late metamyelocyte in the
lower left corner.

8.2 Articulation of artifacts

Artifacts are defined here to be different from plasma protein which was left over from the segmentation discussed in Section 7.6. Artifacts are detected by high density, small area, high $p**2/a$, square corners and sides, etc. Figure 8.3 shows some examples of artifacts. Figure 8.3a is the GB image of a distorted late metamyelocyte in frame 126-9. Figure 8.3b is the GB image of a distorted large lymphocyte in frame 127-14.
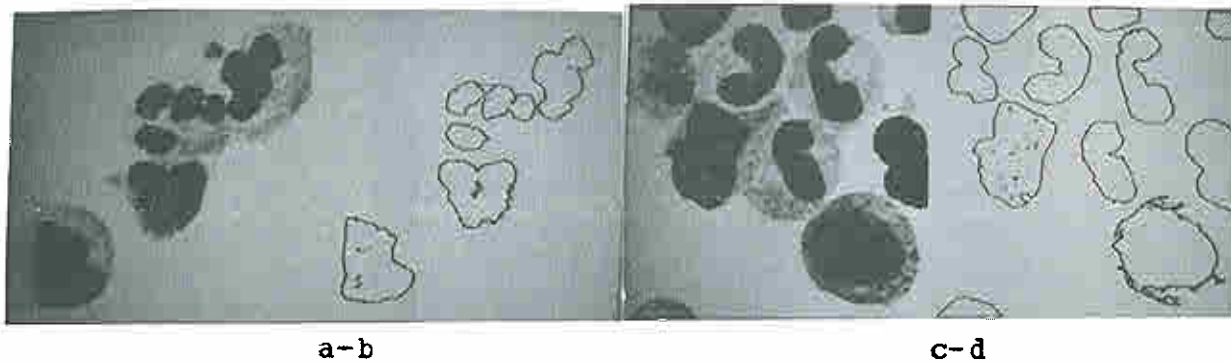

a-b   c-d

Figure 8.3 Examples of artifacts. a) GB image of distorted late metamyelocyte in 126-9, b) boundary trace of 126-9, c) GB image of distorted large lymphocyte in 127-14, d) boundary trace of 127-14.

8.3 Articulation of cell fragments

Cell fragments are formed objects (or what appear to be formed objects) which once belonged to cells. They appear in the smear as a result of both normal and abnormal events. Two common cell fragments are smashed cells and naked nuclei. The former is produced during slide preparation while the latter is the result of normal erythrocyte cell line maturation.

Figure 8.4a shows smashed red cells in frame 126-4. The variation in hemoglobin density would normally cause problems if a binary image of the blue scan were used at the T_{rbc} value. However, the green-blue positive difference algorithm, because of the non-linearity of the operator, still removes the smashed RBCs from the GB image. Figures 8.4b-c shows extruded "naked" nuclei (which have no cytoplasm) in frames 126-5, 126-6. The segmentation algorithm still segments these formed objects. Later when the cell as a whole is analyzed (not done here) the lack of cytoplasm will be noted. A possible problem would arise if the naked nuclei were adjacent to other cells with cytoplasm in which case some of their cytoplasm might be mistakenly counted as belonging to the naked nucleus.
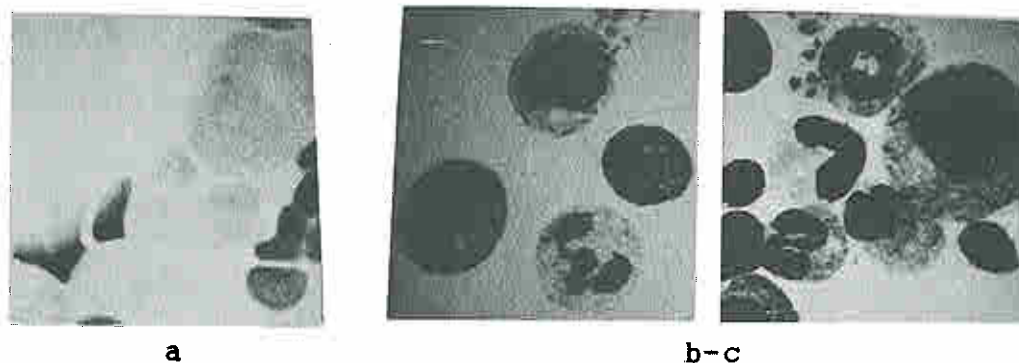
a   b-c

Figure 8.4 Examples of articulation of cell fragments. a) smashed red cells in contrast stretched blue image of frame 126-4 (see Figure 5.10k for results of slicing at Trbc). b) Extruded "naked nuclei" which have no adjacent cytoplasm in contrast-stretched green scans for frames 126-5, and c) 126-6.

9. CONCLUSIONS AND RECOMMENDATIONS

This report has described a system for performing segmentation of marrow images taking problem domain knowledge into account. Several difficult frames typical of marrow smears were segmented to test some of the segmentation algorithms. The algorithms worked well in most cases, but had problems in others. The problems can be explained in terms of required extensions to the algorithms. A morphologic data base presents additional problem domain semantics which could be incorporated into the segmentation model.

9.1 Contributions

The major contributions are listed in Section 2.3 and are divided into two groups, problem domain knowledge driven segmentation and image processing tools. The former includes: compilation of a morphologic data base for bone marrow; RBC elimination algorithm using the positive difference transform; cytoplasm mask generation algorithm; touching-nuclei splitting algorithm; cytoplasm isolation algorithm; run length map algorithm; and boundary trace transform algorithm. The tools include: the Buffer Memory MONitor system (BMON2); the Real Time Picture Processor (RTPP) hardware; the PRocedural Description Language (PRDL); and the PDP10 image processing system (PROC10).

9.2 Possible future extensions of the work

The desirable extensions to this work are primarily in two areas, extensions to the segmentation algorithm and implementation of a high level language connection between the segmentation and the model.

9.2.1 Extensions to segmentation algorithms

Some of the following extensions to the segmentation algorithms would correct several of the problems seen with the current algorithms.

(a) Extend the use of texture in the cytoplasm and nuclear splitting algorithms.

(b) Use one of several optional splitting algorithms for cytoplasm isolation based on success of (and indications from) the isolation method used. Where possible, use touching cytoplasm edge tracking if the split region is strong enough to track; another algorithm might be to find cytoplasm notches in the touching cytoplasm envelope as split junctions when the split region cannot be seen (this is similar to [BrenJ77]).

(c) Develop a backtracking boundary follower which would follow texture gradients of the .cytoplasm/background or cytoplasm/cytoplasm edges.

(d) Implement teardrop continuity of curvature algorithm for poly nuclei fragment articulation.

(e) Develop RLM global shape features discussed in [LemP78].

(f) Develop rouleaux (overlapping RBCs) processing methods incorporating shape and overlap analysis. This might also be extended to handle RBC overlap of leukocytes.

(g) Develop a Golgi region finding algorithm based on a lightened region in cytoplasm conditioned by a stab nucleus shape.

(h) Develop a texture based segmentation algorithm (possibly based on a BTT(F) extension) for finding granules in cytoplasm.

(i) Develop a similar BTT(T) algorithm for finding vacuoles between cell and nuclear boundaries (Note that [BrenJ77] detected vacuoles touching the inside of the cell membrane but did not detect those not touching it.).

(j) Develop a BTT(T) algorithm (possibly also using BTT(F)) to segment difficult granulocytes such as basophils based on modelling texture regions between the cell boundary and the potential nuclear boundary found by BTT(T).

(k) Develop an ability to work on multiple windows larger than 256x256 in order to being to handle megakaryocytes and macrophages.

9.2.2 High level language connection

The connection between the marrow model and the low level segmentation process could be bridged by a high level model. The model could be extended as is suggested in Section 2 by combining different levels of processing (control structure and low level image operations).

(a) Implement a connection of the high level marrow model with the low level image segmentation.

(b) Incorporate some of the marrow morphologic and meta-morphologic data base into a high level model and use it to drive the segmentation.

(c) Extend the marrow morphologic data base to abnormal cells and situations which would include disease states.

(d) Work on a methodology for use in a long term cost effective practical system.

(e) Extend the method to other biological image domains.

10. REFERENCES

AgrA77. Agrawala A K, Kulkarni A V: A Sequential Approach to the Extraction of Shape Features. Comp Graph Image Proc, Vol 6, 1977, 538:557.

AggR77. Aggarwal R K, Bacus J W: A Multi-spectral Approach for Scene Analysis of Cervical Cytology Smears. J Hist Cyto, Vol 25, 1977, 668:680.

ArceC71. Arcelli C, Levialdi S: Picture Processing of Overlapping Blobs. IEEE Trans Comp, Vol C-20, 1971, 1111:1115.

AusH77. Aus H M, Ruter A, Ter Muelen V, Gunzer U, Nurnberger R: Bone Marrow Cell Scene Segmentation by Computer-Aided Color Cytophotometry. J Hist Cyto, Vol 25, 1977, 662:667.

BacJ71. Bacus J: An Automated Classification of the Peripheral Blood Leukocytes by Means of Digital Image Processing. Thesis, Univ Illinois, 1971.

BacJ72. Bacus J W, Gose E E: Leukocyte Pattern Recognition. IEEE Trans Syst Man Cyber, Vol SMC-2, 1972, 513:526.

BacJ73. Bacus J W: The Observer Error in Perhiperal Blood Cell Classification. Amer J Clin Path, Vol 59, 1973, 223:230.

BacJ74. Bacus J: Erythrocyte Morphology and Centrifugal "Spinner" Blood Film Preparations. J Hist Cyto, Vol 22, 1974, 506:516.

BacJ75. Bacus J W, Aggarwal R K, Belanger M G: Computer Recognition of Microscopic Images. EASCON, 1975, 73-A:73-F.

BacJ76. Bacus J W: A Whitening Transform for Two-Color Blood Cell Images. Pat Rec, Vol 8, 1976, 53:60.

BacJ77. Bacus J W, Weens J H: An Automated Method of Differential Red Blood Cell Classification with Application to the Diagnosis of Anemia. J Hist Cyto, Vol 25, 1977, 614:632.

BartP76a. Bartels P, Bibbo M, Wied G L: Modeling of Histologic Images by Computer. Acta Cyto, Vol 20, 1976, 62:66.

BengE76. Bengtsson E, Holmquist J, Olsen B, Stenkvist B: SCANCANS - An Interactive Scanning Cell Analysis System. Comput Prog Biom, Vol 6, 1976, 39:49.

BengE77. Bengtsson E:On the Design of Systems for Computer Aided Analysis of Microscope Images. Dissertation Inst. Physics, Uppsala Univ and Dept Clin Cyto, Univ Hosp, Uppsala, Report 428, 1977.

BesM73. Bessis M:Living Blood Cells and Their Ultrastructure. Springer-Verlag, NY (translated by R L Weed), 1973.

BrenJ74. Brenner J F, Gelsema E S, Necheles T F, Neurath P W, Selles W D, Vastola E V: Automated Classification of Normal and Abnormal Leukocytes. J Hist Cyto, Vol 22, 1974, 697:706.

BrenJ76. Brenner J F, Dew S B, Horton J B, King T, Neurath P W, Selles W D: An Automated Microscope for Cytologic Research. J Hist Cyto, Vol 24, 1976, 100:111.

BrenJ77. Brenner J F, Necheles T F, Bonacossa I A, Fristensky R, Weintraub B A, Neurath P W: Scene Segmentation Techniques for the Analysis of Routine Bone Marrow Smears from Acute Lymphoblastic Leukemia Patients. J Hist Cyto, 25, 1977, 601:613.

BowJ77. Bowie J E, Young I T: An Analysis Technique for Biological Shape - III. Acta Cyto, Vol 21, Nov/Dec, 1977.

CahR77. Cahn R L, Poulsen R S, Toussaint G: Segmentation of Cervical Cell Images. J Hist Cyto, Vol 25, 1977, 689:695.

CarmG74. Carman G, Lemkin P, Lipkin L, Shapiro B, Schultz M, Kaiser P:A Real Time Picture Processor for Use in Biological Cell Identification - II Hardware Implementation. J Hist Cyto, Vol 22, 1974, 732:740.

CheG74. Cheng G: Color Information in Blood Cells. J Hist Cyto, Vol 22, 1974, 517:522.

Davl74. Davidson I, Henry J B (eds): Todd-Sanford Clinical Diagnosis. 15th edition, W. B. Saunders Co., Phila., 1974, 100:448.

DavL77. Davis L S: Understanding Shape: Angles and Sides. IEEE Trans Comp, Vol C-26, 1977, 236:242.

DEC74. Digital Equipment Corp: OS8 Handbook. Maynard, Mass., 1974.

EccM77. Eccles M J, McQueen M P C, Rosen D: Analysis of the Digitized Boundaries of Planar Objects. Pat Rec, Vol 9, 1977,

31:41.

FelJ72.   Feldman J, Low J, Swineart D, Taylor R: Recent
developments in SAIL. SAI Memo June 5, 1972.

FreH74. Freeman H: Computer Processing of Line Drawing Images.
Comp Surveys, Vol 6, March 1974, 57:97.

FreH77a.   Freeman H: Shape Description Via the Use of Critical
Points. Proc Pattern Recog and Image Process Conf, Troy, NY,
1977, 168:174.

FreH77b.   Freeman H,   Davis L: A Corner Finding Algorithm for
Chain-Coded curves. IEEE Trans Comp, Vol C-26, 1977, 297:303.

FuK71.   Fu K S: Syntactic Methods in Pattern Recognition.
Academic Press, 1974.

GalM75. Galloway M: Texture Analysis using Gray Level Run
Lengths. Comp Graph Image Proc, Vol 4, 1975, 172:179.

GalG74.   Gallus G, Regoliosi G: A Decisional Model of
Recognition Applied to the Chromosome Boundaries. J Hist Cyto,
Vol 22, 1974, 546:553.

GolM69. Golay M: Hexagonal Parallel Pattern Transformations.
IEEE Trans Comp, Vol C-18, 1969, 734:740.

HewC72. Hewitt C: Description and Theoretical Analysis (Using
Schemata) of PLANNER: a language for proving theorems and
manipulating models of a robot. Thesis, MIT, AI-TR-258, 1972.

HolJ77. Holmquist J, Bengtsson E, Ericksson O, Stenkvist B: A

Program System for Interactive Measurements on Digitized Cell Images. J Hist Cyto, Vol 25, 1977, 641:654.

IngM68. Ingram M, Preston K Jr, Norgren P E: Automatic Differentiation of White Blood Cells. In "Image Processing in Biological Science", D M Ramsey (ed), UCLA Press, Berkeley, 1968.

IngM69. Ingram M, Norgren P E, Preston K Jr: Advantages of Topology as a Basis for Automatic Analysis of Blood Cell Images. Ann NY Acad Sci, Vol 157, 1969, 275:284.

KanL72. Kanal L N, Chandrasekaran B: On Linguistic, Statistical and Mixed Models for Pattern Recognition. In "Frontiers of Pattern Recognition". S. Watanabe, Academic Press, 1972, 163:192.

KliA71. Klinger A, Kochman A, Alexandridis N: Computer Analysis of Chromosome Patterns: Feature Encoding for Flexible Decision Making. IEEE Trans Comp, Vol C-20, 1971, 1014:1022.

KnoG73. Knott G, Reese D: MLAB - An On-line Modeling Laboratory. NIH, DCRT, Dec 1973.

LanG77. Landeweerd G H, Gelsema E S: The Use of Texture Parameters in the Automatic Analysis of Leukocytes. Pat Rec, in press.

LedR66. Ledley R S, Jacobson J, Belson M: BUGSYS: A Programming System for Picture Processing - Not for Debugging. CACM, Vol 9, 1966, 79:84.

LeeE77. Lee E T: Similarity Directed Picture Storage and Management. In Proc Workshop on Picture Data Description and Management, April 21-22, 1977, 150:154.

LemP72. Lemkin P: A Simplified Biological Cell World Model for Question-Answering using Functional Description. Univ Md Scholarly Paper 75, May 16, 1972.

LemP74. Lemkin P, Carman G, Lipkin L, Shapiro B, Schultz M, Kaiser P:A Real Time Picture Processor for Use in Biological Cell Identification - I System Design. J Hist Cyto, Vol 22, 1974, 725:731.

LemP75. Lemkin P: A Literature Survey of the Technological Basis for Automated Cytology. Univ Md CMSC TR-384, 1975.

LemP76a. Lemkin P, Shapiro B, Gordon R, Lipkin L: PROC10 - An Image Processing System for the PDP10. NCI/IP Technical Report #8, NTIS PB261535/AS, Dec, 1976.

LemP76b. Lemkin P: Functional Specifications for the RTPP Monitor - Debugger DDTG. NCI/IP Technical Report #2, NTIS PB250726 (or DECUS 8-823), Feb, 1976.

LemP76c. Lemkin P: BMOMNI - Fortran Interface Program for the RTPP Buffer Memory, Quantimet and Control Desk. NCI/IP Technical Report #23, NTIS PB261538/AS, December, 1976.

LemP76d. Lemkin P:MAG10 - A PDP8e File Based Magtape Utility. NCI/IP Technical Report #20, NTIS PB261534/AS (or DECUS 8-879), Dec, 1976.

LemP76e. Lemkin P:Buffer Memory Monitor System for Interactive Image Processing. NCI/IP Technical Report #21, NTIS PB261536/AS, Dec. 1976.

LemP77a. Lemkin P:Buffer Memory Monitor System for Interactive Image Processing. NCI/IP Technical Report #21a, NTIS PB269642/AS, June, 1977.

LemP77b. Lemkin P, Carman G, Lipkin L, Shapiro B, Schultz M:Real Time Picture Processor - Description and Specification. NCI/IP Technical Report #7a, NTIS PB269600/AS, June, 1977.

LemP78a. Lemkin P: The Run Length Map - A Representation of Contours and Regions for Efficient Search and Low Level Semantic Encoding. Univ Md CMSC TR-655, 1978.

LemP78b. Lemkin P: Morphologic Data Base for Bone Marrow. Univ Md CMSC TR-654, 1978.

LipL66. Lipkin L E, Watt W, Kirsch R: Analysis, Synthesis and Description of Biological Images. Annals NY Acad Sci, Vol 128, 1966, 984:1012.

LipL76. Lipkin L E, Lipkin B: Data Type Construction and the Analysis of Complex Situations. In "Physics in Industry", Proc International Conf Held in Dublin, March 9-13, E O O'Mongain and C P O'Toole (eds), Pergamon Press, Oxford, 1976.

MarN76. Markovic N:Blast Cell Morphology, A Possible Parameter for the Estimation of Acute Leukemia Prognosis. Dissertation, Belgrade Medical Faculty. March 1976.

MayB74. Mayall B (ed): Special Issue Third Engineering Foundation Conference on Automated Cytology. J Hist Cyto, Vol 22, 1974.

MayB76. Mayall B (ed): Special Issue Fourth Engineering Foundation Conference on Automated Cytology. J Hist Cyto, Vol 24, 1976.

MayB77. Mayall B (ed): Special Issue Fifth Engineering Foundation Conference on Automated Cytology. J Hist Cyto, Vol 25, 1977.

McDD72. McDermott D, Sussman G: Conniver Reference Manual, MIT AI Memo 259, May 1972.

MerR73. Merrill R D:Representation of Continuous Regions for Efficient Computer Search. CACM, Vol 16, 1973, 69:82.

MenM74. Mendelsohn M, Bennett D E, Bogart E, Mayall B H: Computer-Oriented Analysis of Human Chromosomes. IV. DNA-based Centromeric Index. J Hist Cyto, Vol 22, 1974, 554:560.

MesC70. Mesztenyi C: FORMAL - a Formula Manipulation Language. Univ Md. CMSC TR70-133, Sept, 1970.

MuiJ77. Mui J K, Fu K S, Bacus J W: Automated Classification of Blood Cell Neutrophils. J Hist Cyto, Vol 25, 1977, 633:640.

PlatW69. Platt W R: Color Atlas and Textbook of Hematology. J B Lippincott Company, Phila, 1969.

PoulR77. Poulsen R S, Oliver L H, Cahn R L, Louis C, Toussaint

G: High Resolution Analysis of Cervical Cells - A Progress Report. J Hist Cyto, Vol 25, 1977, 689:701.

PresN76. Pressman N J: Optical Texture Analysis for Automatic Cytology and Histology: A Markovian Approach. Dissertation, Lawrence Livermore Lab, UCLA, Report # UCRL-52155, Oct 1976.

PresK71. Preston K Jr: Feature Extraction by Golay Hexagonal Pattern Transformations. IEEE Trans Comp, Vol C-20, 1971, 1007:1013.

PresK72. Preston K Jr: On Determining Optimum Simple Golay Marking Transformations for Binary Image Processing. IEEE Trans Comp, Vol C-21, 1972, 1430:1433.

PresK76. Preston K Jr: Digital Picture Analysis in Cytology. In "Digital Picture Analysis", A Rosenfeld (Ed), Springer-Verlag, NY, 1976, 209:294.

PrewJ66. Prewitt J M S, Mendelsohn M I: The Analysis of Cell Images. Annals NY Acad Sci, Vol 128, 1966, 1035:1053.

QuaL68. Quam L: Standford LISP 1.6 manual. SAI 28.2, Dec 1968.

ReisJ76. Reiser J: SAIL User Manual. SAI Memo AIM-289, August 1976.

RosA69. Rosenfeld A:Picture Processing by Computer. Academic Press, NY, 1969.

RosA76. Rosenfeld A, Kak A: Digital Picture Processing. Academic Press, NY, 1976.

SchB77. Schachter B, Davis L, Rosenfeld A: Some Experiments in Image Segmentation by Clustering of Local Feature Values. Submitted to Pat Rec.

ShapB74. Shapiro B, Lemkin P, Lipkin L: The Application of Artificial Intelligence Techniques to Biological Cell Identification. J Hist Cyto, Vol 22, 1974, 741-750.

ShapB76. Shapiro B:A SLR(1) Parser Generator. NCI/IP-76/01, Technical Report #9, NTIS PD249127/AS, Feb, 1976.

ShapB77. Shapiro B, Lemkin P: PRDL - Procedural Description Language. NCI/IP Technical Report #13, NTIS PB273112/AS, Oct, 1977.

ShapB78. Shapiro B: Dissertation. In preparation.

SmiK75. Smith K T, Solmon D C, Wagner S L: Practical and Mathematical Aspects of the Problem of Reconstructing Objects from Radiographs. Address to Far West Sect Meeting of American Math Soc, Monterey, Calif., April, 1975.

SobI76. Sobel I: Personal Communication, June 23, 1976.

SilR70. Silver R T: Morphology of the Blood and Marrow in Clinical Practice. Grune and Stratton, NY, 1970.

SycJ76. Sychra J J, Bartels P H, Taylor J, Bibbo M, Wied G L: Cytoplasmic and Nuclear Shape Analysis for Computerized Cell Recognition. Acta Cyto, Vol 20, 1976, 68:78.

TalG57. Talbot G D, Hunsicker E R, Li J: Blood and Marrow

Patterns. Grune and Stratton, NY, 1957.

TayJ75. Taylor J, Bahr G F, Bartels P H, Bibbo M, Richards D L, Wied G L: Development and Evaluation of Automatic Nucleus Finding Routines: Thresholding of Cervical Cytology Images. Acta Cyto, Vol 19, 1975, 289:298.

TenJ74. Tenenbaum J M, Garvey T D, Weyl S, Wolf H: An Interactive Facility for Scene Analysis Research. Stanford Research Institute Tech Note 87, Jan 1974.

TycD76. Tycko D H, Anbalagan S, Liu H C, Ornstein L: Automatic Leukocyte Classification Using Cytochemically Stained Smears. J Hist Cyto, Vol 24, 1976, 178:194.

WecH75. Wechsler H, Sklansky J:Automatic Detection of Rib Contours in Chest Radiographs. 4IJCAI, 1975, 688:694.

WiedG70. Wied G, Bahr G F, Bartels P H: Automatic Analysis of Cell Images by TICAS. In "Automated Cell Identification and Cell Sorting", Wied G, Bahr G F (eds), Academic Press, NY, 1970, 195:356.

YouI72. Young I T: Classification of White Blood Cells. IEEE Trans Biomed Eng, Vol 19, 1972, 291:298.

YouI74. Young I T, Walker J E, Bowie J E: An Analysis Technique for Biological Shape I. Inf Contr, Vol 25, 1974, 357:370.

YouI75. Young I T, Paskowitz I L: Localization of Cellular Structures. IEEE Trans Biomed Eng, Vol 22, 1975, 35:40.

Appendix

IMAGE DATA BASE

This Appendix contains photographs of the thirteen multispectral frames of the bone marrow data base. The frames are numbered 126-3 to 126-10 and 127-11 to 127-15. Each frame description contains:

1. Four images/frame (white, green (546 nm), blue (420 nm), and the positive difference (green,blue), denoted GB.

2. Labeled drawings of each of the frames with objects labeled with capital letters from top to bottom and from left to right. In cases where more than one object of a given type occurs in a field, a number is appended.

3. A verbal description of particular features of each of the frames.

The slide used was a Wright's stain normal marrow smear acquired by Markovic from the NIH Clinical Center Hematology department with slide # H-1163 and staining date of 6-25-77. The slide was scanned using the BMON2/Axiomat microscope system on 1-26-77 and 1-27-77. The microscope was set up using the 100 watt mercury arc lamp, Koehler illumination, using the the 100X NA1.32 oil objective lens. The vidicon camera was used with the Quantimet automatic video gain control turned on. Frames are denoted by 126-i or 127-i depending on which date they were scanned.
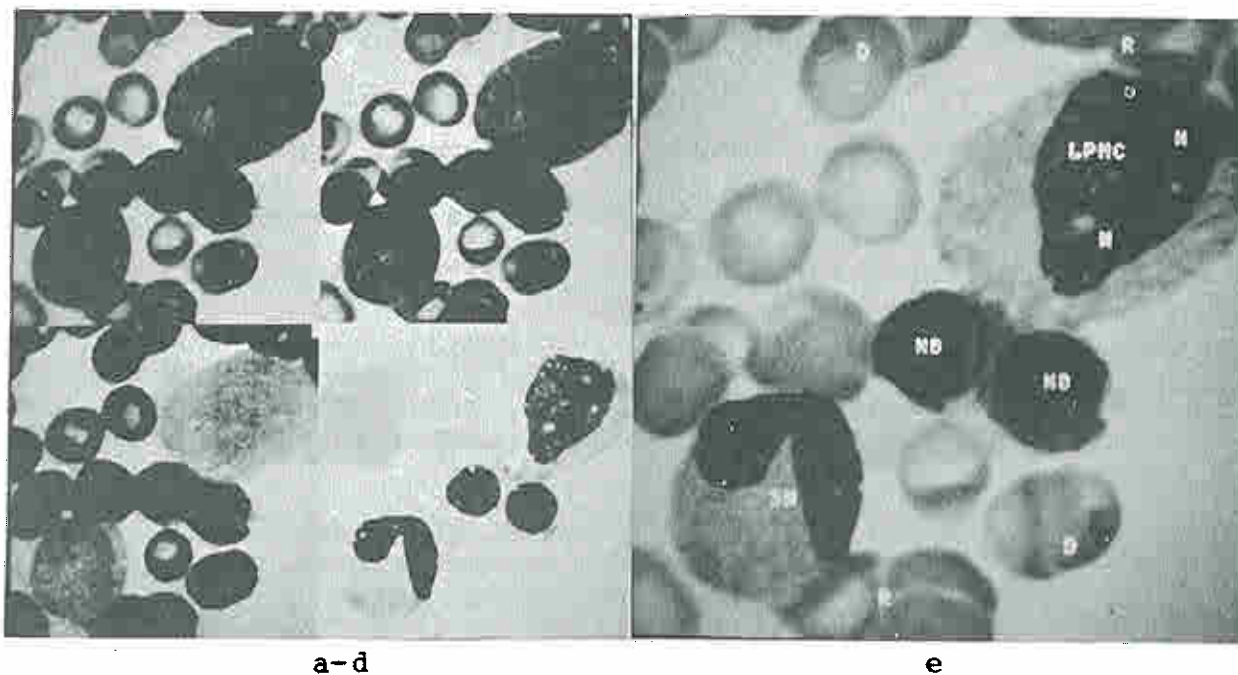
a-d   e
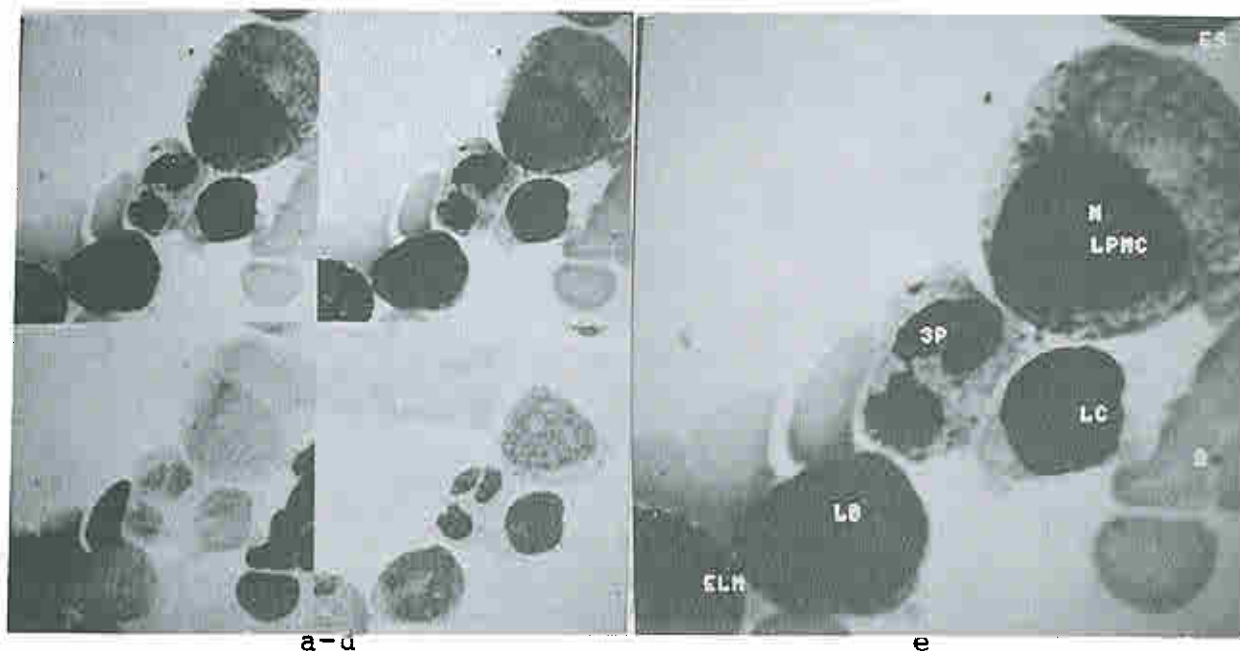
Figure A.1 Frame 126-3. a) white, b) green, c) blue, d) positive difference GB, e) labeled green frame.

Features of interest

Frame 126-3 contains (from left to right, top to bottom): a late promyelocyte (LPMC), two normoblasts (NB), and a stab neutrophil (SN). There are rouleaux (R) adjacent to the SN and one of the NB. The LPMC nucleus is distorted and possibly has two nucleoli (N). There is dirt (artifact) overlying the RBC in the lower right part of the field (D) and the RBC at the top of the field.

Figure A.2 Frame 126-4. a) white, b) green, c) blue, d) positive difference GB, e) labeled green frame.

Features of interest

Frame 126-4 contains (from left to right, top to bottom): an empty region due to wetting effects (E), part of an early stab (ES) (cut by top of frame), a late promyelocyte (LPMC), a 3-lobed bridged poly (3P), a lymphocyte (LC), RBCs in rouleaux (R), a lymphoblast (LB), and an eosinophilic late metamyelocyte (ELM) (in the lower left corner of the frame). The LPMC, 3P, and LB are touching. The rouleaux are adjacent to 3P, LB, and ELM. The LPMC has a nucleolus (N). The 3P has a very fine bridge between the lobes. The LB is sufficiently undifferentiated that some hematologists might prefer to call it an undifferentiated "blast" cell.
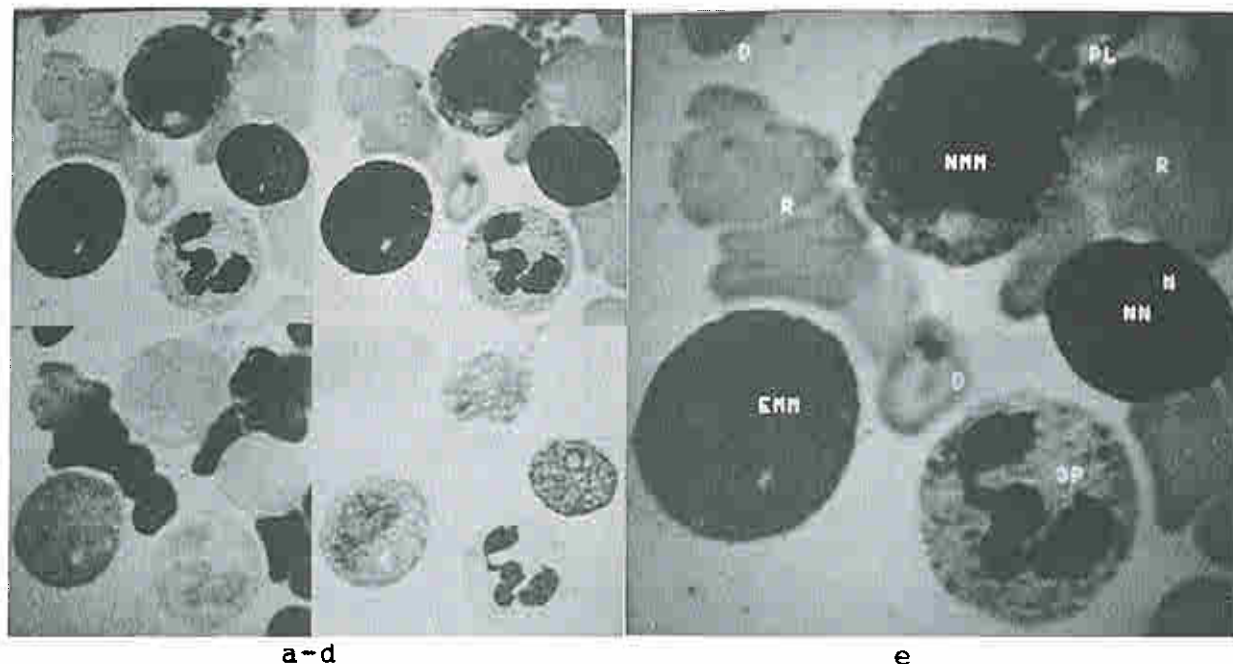
a-d   e

Figure A.3 Frame 126-5. a) white, b) green, c) blue, d) positive difference GB, e) labeled green frame.

Features of interest

Frame 126-5 contains (from left to right, top to bottom): dirt overlying RBC (D), a neutrophilic metamyelocyte (NMM), platelets (PL), rouleaux (R), dirt overlying RBC in the center (D), a naked nucleus (NN) of unspecified class membership, an eosinophilic metamyelocyte (EMM), and a three-lobe poly (3P). The PL is touching the top of NMM. Rouleaux R are between NMM and EMM, and NMM and NN. Another R is touching NN. The NMM has about 30-40 azure granules and its Golgi region is particularly prominent. The NN has a nucleolus N.
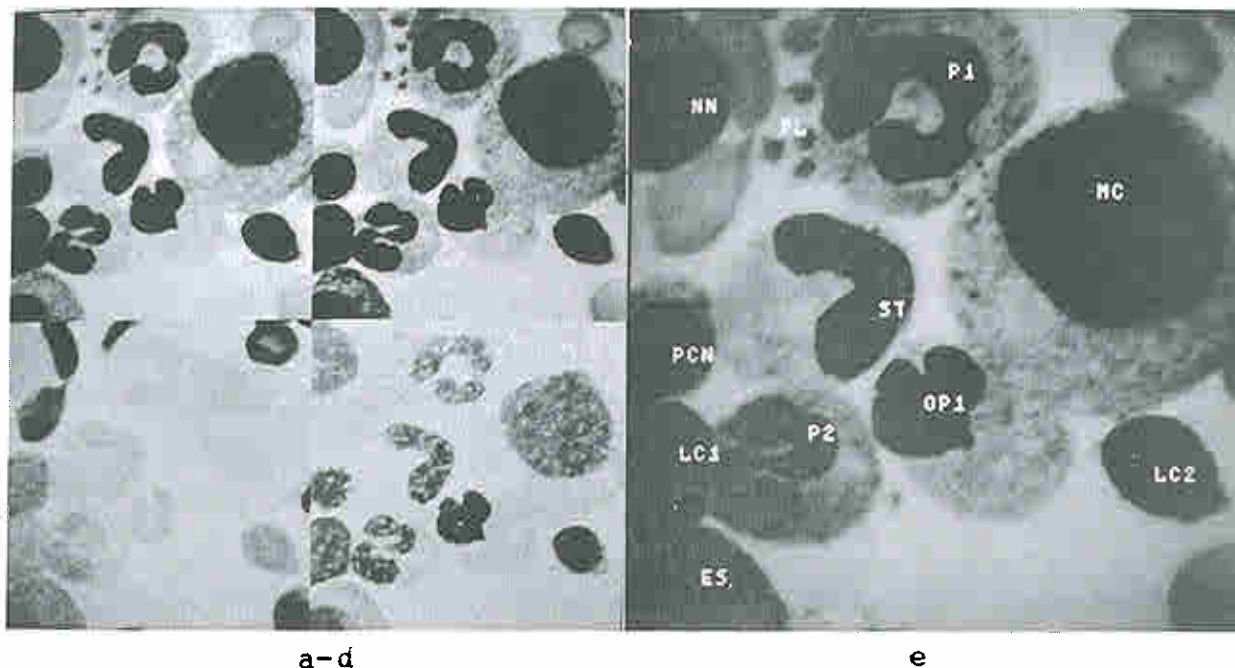
a-d   e

Figure A.4 Frame 126-6. a) white, b) green, c) blue, d)
positive difference GB, e) labeled green frame.

Features of interest

Frame 126-6 contains (from left to right, top to bottom):   part
of a naked nucleus NN (touching the left frame edge), RBCs
overlying NN, platelets (PL), a poly (P1), a myelocyte with too
mature nuclear texture (MC), a stab (ST), an overlapping poly
(OP1), a polychromatic normoblast (PCN), a poly (P2),
a lymphocyte intersecting the frame (LC1), a lymphocyte (LC2),
and an eosiniphil intersecting the left and bottom of the frame
(ES).   The MC has no secondary granules, is large like a
promyelocyte but has no nucleolus. The nuclear texture is out
of phase with respect to the cytoplasm because of the lack of
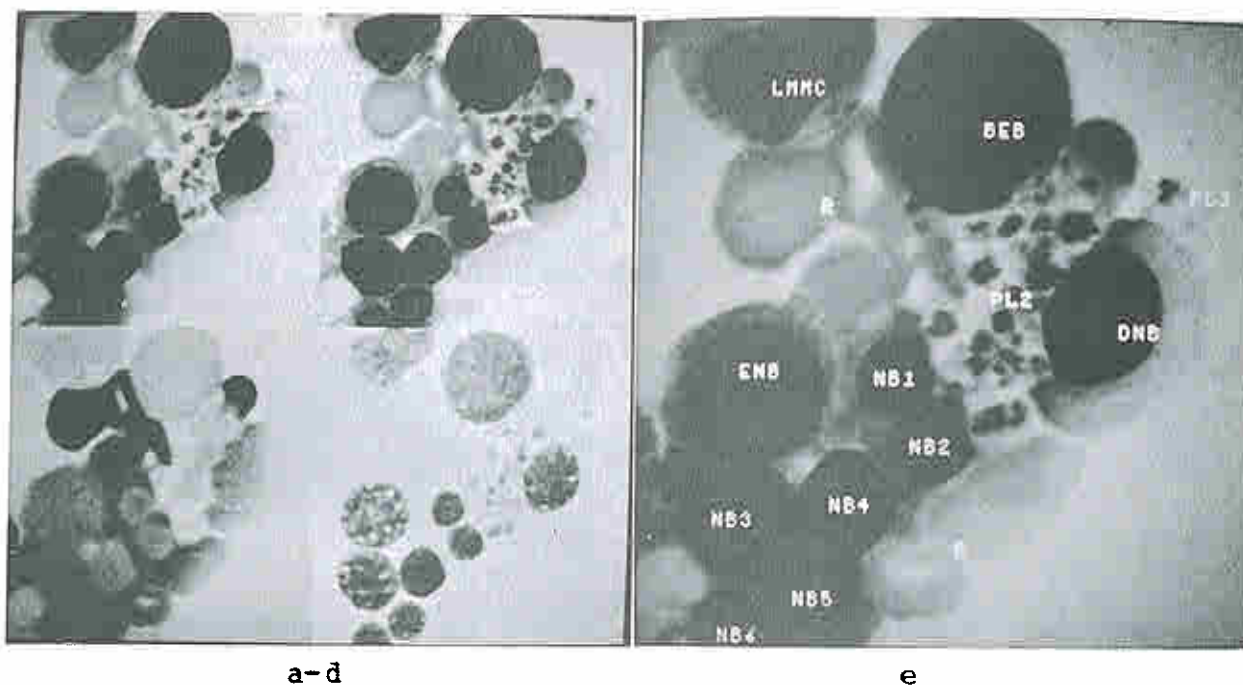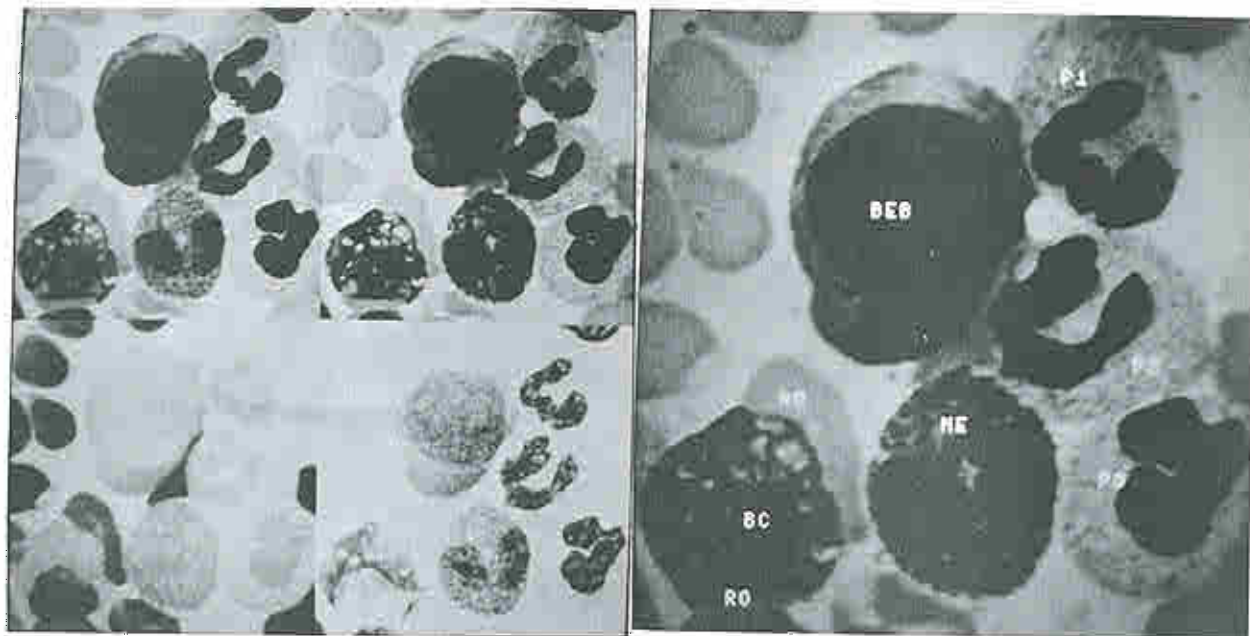secondary granules.

a-d   e

Figure A.5 Frame 126-7. a) white, b) green, c) blue, d) positive difference GB, e) labeled green frame.

Features of interest

Frame 126-7 contains (from left to right, top to bottom): part of a late metamyelocyte (LMMC) intersecting the top of the frame, a platelet (PL1) on the bottom of a basophilic erythroblast (BEB), rouleaux (R), platelets (PL2), a distorted normoblast cell (DNB), an early normoblast (ENB), and normoblasts in various stages of maturation (NB1,NB2), (NB3,NB4), (NB5,NB6) with (NB6) intersecting the bottom of the frame. The PL2 are between BEB, DNB, and NB1 and NB2. The normoblasts are touching one another.

a-d   e

Figure A.6 Frame 126-8. a) white, b) green, c) blue, d) positive difference GB, e) labeled green frame.

Features of interest

Frame 126-8 contains (from left to right, top to bottom): RBCs with some touching each other and leukocytes, a poly (P1), a basophilic erythroblast (BEB) with a light region artifact on top of it, a poly (P2), a mature eosinophil (ME), a poly (P3), a basket cell with vacuoles of degeneration and RBCs overlapping it (BC). The BEB has dark cytoplasm because the basophilic material is pushed to one side which is probably an artifact. The basket cell has eosinophilic cytoplasm.
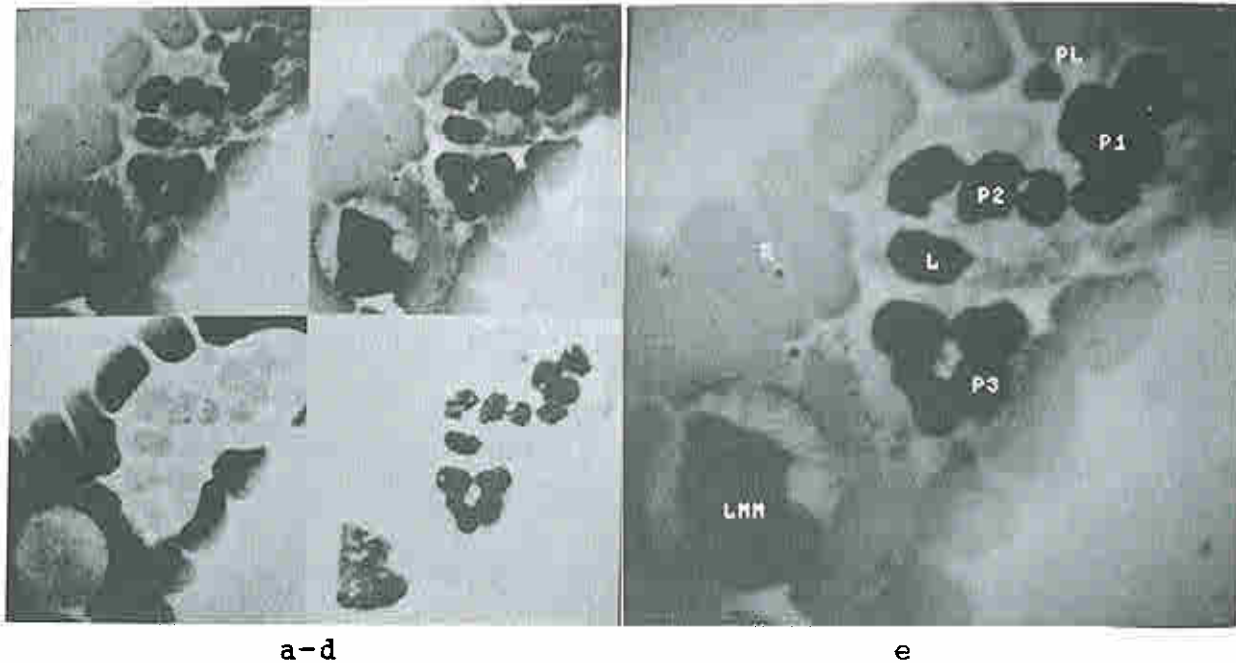
a-d   e

Figure A.7 Frame 126-9. a) white, b) green, c) blue, d)
positive difference GB, e) labeled green frame.

Features of interest

Frame 126-9 contains (from left to right, top to bottom):
leukocytes embedded in rouleaux (R) clusters, an isolated
platelet (PL) between two polys (P1) and (P2), a poly (P3), and
a distorted late metamyelocyte (LMM) at the bottom. The polys
seem to each have four lobes. The lobe (L) between P2 and P3
could be mistaken as belonging to P3 instead of probably
belonging to P2. The LMM is possibly a stab which is compressed
by the RBCs (although the nucleus and cytoplasm distortion
makes the decision between a late metamyelocyte and stab
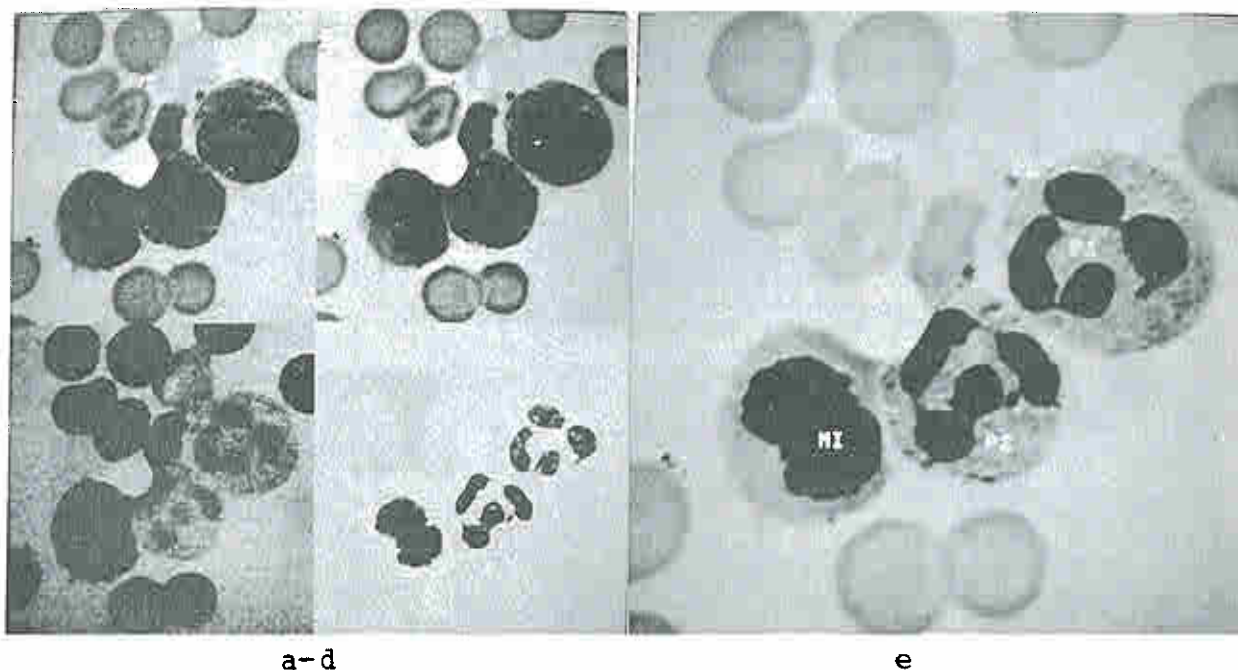difficult). The granules of LMM are probably compressed at the
surface.

a-d   e

Figure A.8 Frame 126-10. a) white, b) green, c) blue, d) positive difference GB, e) labeled green frame.

Features of interest

Frame 126-10 contains (from left to right, top to bottom): a few RBCs with one adjacent to two touching four lobed polys (P1) and (P2), a platelet (PL) adjacent to P1; the cell at the bottom is in mitosis (MI) and is touching P2. The cell in mitosis has teardrop extensions comming out of the cytoplasm at the bottom. Because of the shape, there might be confusion between a cell in mitosis and a stab. However, the cell does not show many granules in the cytoplasm which would be the case in a stab, the nucleus is denser in general, the nuclear texture is different, and the breadth of the nuclear bridge is greater in the case of the cell in mitosis. The rightmost nuclear lobe of P1 has a slight downward extension which may represent a nuclear appendage or fragmenentation of a nuclear bridge.
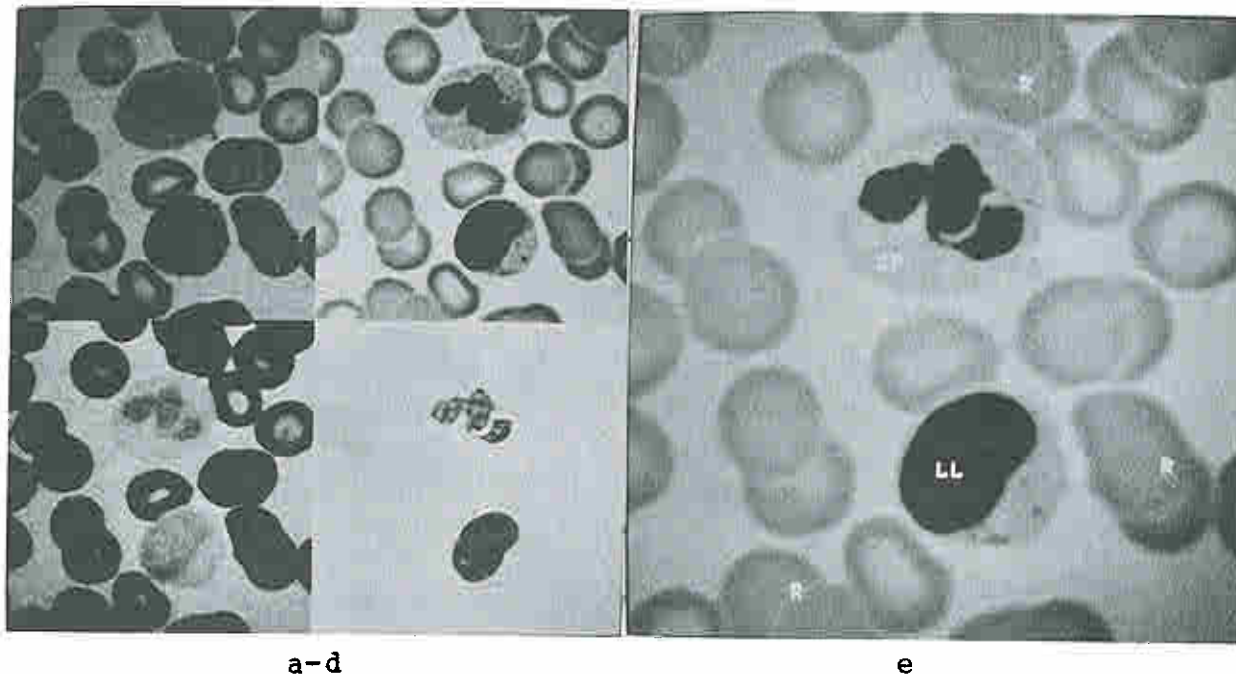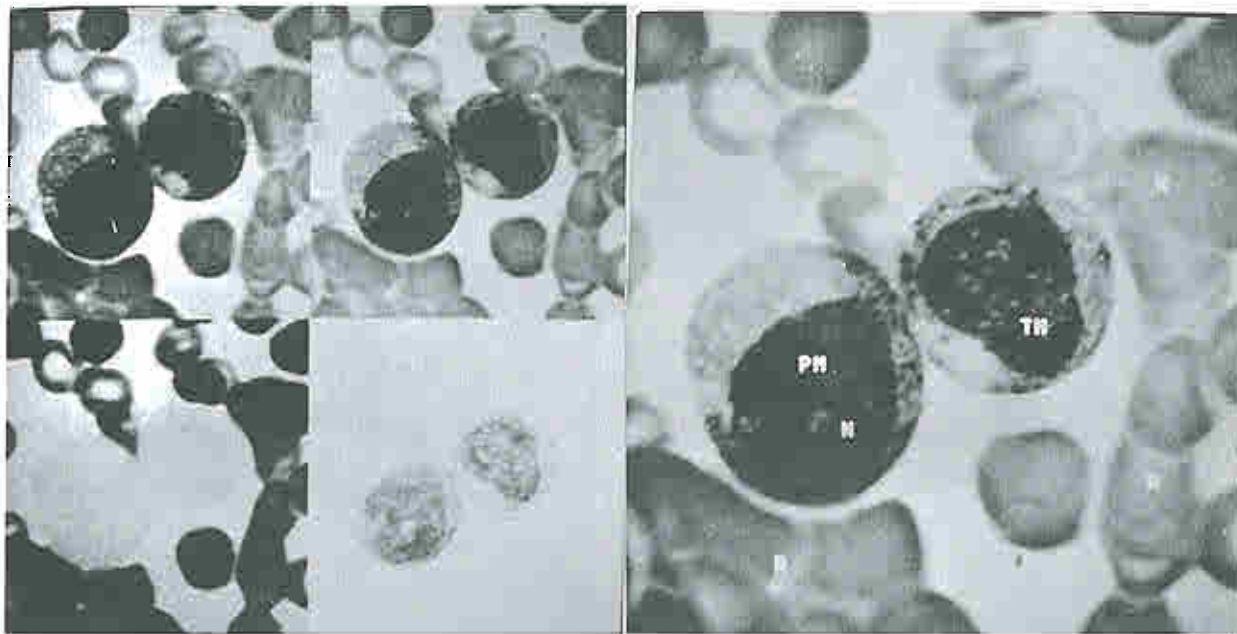
a-d   e

Figure A.9 Frame 127-11. a) white, b) green, c) blue, d)
positive difference GB, e) labeled green frame.

Features of interest

Frame 127-11 contains (from left to right, top to bottom): a

three lobe poly (3P) with two lobes overlapping one another, a

large lymphocyte (LL), and RBCs with some grouped in rouleaux.

The large lymphocyte is not found too often in the marrow.   It

has several azurophilic granules in its cytoplasm.   It might

possibly be classified by others as a small monocyte.

a-d   e

Figure A.10 Frame 127-12. a) white, b) green, c) blue, d) positive difference GB, e) labeled green frame.

Features of interest

Frame 127-12 contains (from left to right, top to bottom): clumped and distorted RBCs, a transition stage between late promyelocyte and myelocyte (TM), and a promyelocyte (PM). The TM has secondary and azurophilic granules. The PM has no secondary granules and has a nucleolus (N).
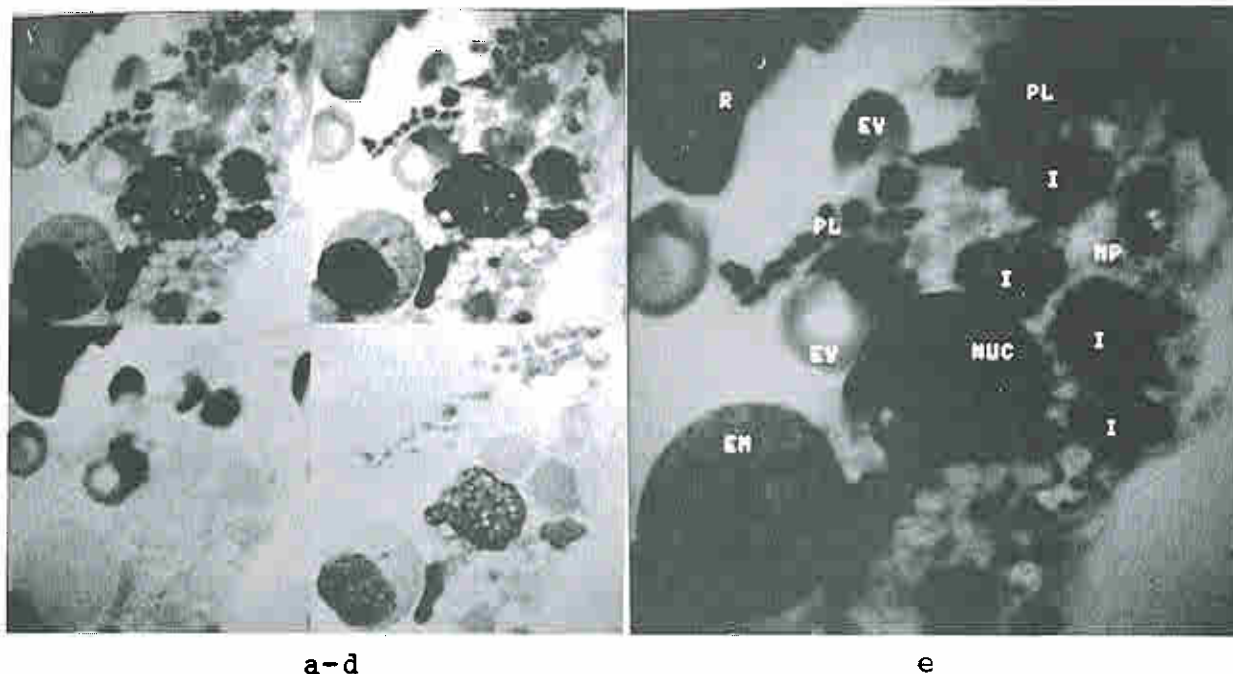
a-d   e

Figure A.11 Frame 127-13. a) white, b) green, c) blue, d)
positive difference GB, e) labeled green frame.

Features of interest

Frame 127-13 contains (from left to right, top to bottom):
rouleaux (R) in the upper left corner, and a macrophage (MP)
which extends outside of the scan field. The cytoplasmic edges
are notably indistinct, and the normally variable texture of
the cytoplasm is futher accentuated by ingested material. It
has platelets (PL) near the top with adjacent RBCs on the upper
left.   Several ingested, partially digested nuclei are visible
(I). Its nucleus (NUC) is larger than the I. An early
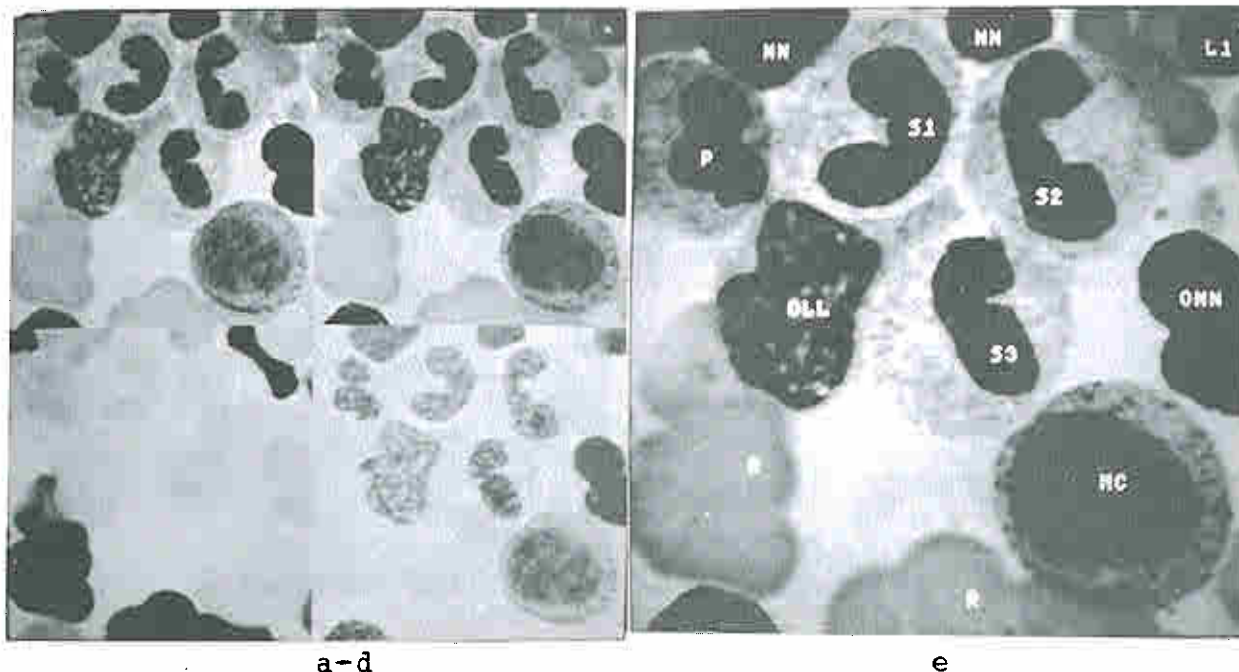metamyelocyte (EM) is partly in the lower left corner of the
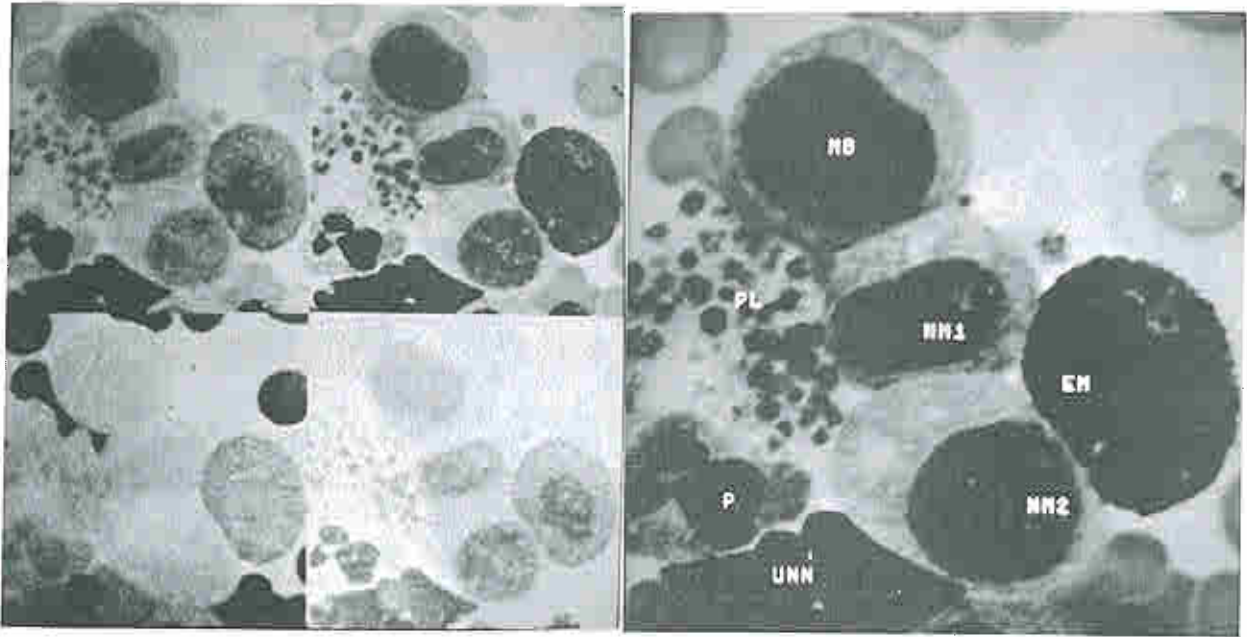frame.

a-d   e

Figure A.12 Frame 127-14. a)   white,   b)   green,   c)   blue,   d)
positive difference GB, e) labeled green frame.

Features of interest

Frame 127-14 contains (from left to right, top to bottom):
naked nuclei (NN) intersecting the top of the frame, a
lymphocyte intersecting the top right corner (L1), a poly (P),
two touching stabs back-to-back (S1) and (S2), a distorted
large lymphocyte (DLL) having cytoplasm between S1 and another
stab (S3) in the center of the frame, two overlapping naked
nuclei (ONN) intersecting the right center of the frame,
rouleaux (R) at the lower left of the DLL, and a myelocyte
(MC).

a-d   e

Figure A.13 Frame 127-15. a) white, b) green, c) blue, d) positive difference GB, e) labeled green frame.

Features of interest

Frame 127-15 contains (from left to right, top to bottom): RBCs at the top center, a myeloblast (MB), neutrophilic myelocytes (NM1) and (NM2), an RBC with overlying dirt (D), an eosinophilic myelocyte (EM), a single platelet (PL1) between NM1 and EM, platelets (PL) on the left of MB and the NM2, a poly (P) at the lower left also touching PL, and unknown naked nuclei (UNN) at the lower left bottom.