

Illustrated Book Study: Digital Conversion Requirements of Printed Illustrations

**Anne R. Kenney and Louis H. Sharpe II
with Barbara Berger, Rick Crowhurst, D. Michael Ott, and Allen Quirk**

July 1999

RELIEF



Wood Engraving

INTAGLIO



Photogravure

PLANOGRAPHIC



Collotype

**Report to the Library of Congress
Preservation Directorate
Contract # IN97C-22/97CLCCT7021**

Contractors

*Cornell University Library, Department of Preservation and Conservation, Ithaca, NY
Picture Elements, Inc., Berkeley, CA*

Cover illustration: The authors intentionally chose representative eyes from the various book illustrations as a tribute to James M. Reilly, Director of the Image Permanence Institute. In his "Flowchart for Identification Guide," Care and Identification of 19th Century Prints (1996), Reilly used photomacrographs of eyes to illustrate the differences among various 19th century photographic prints. This chart is available at: <http://www.rit.edu/~661www1/FRAMESET.html>.

Table of Contents

TABLE OF CONTENTS	3
ABSTRACT	5
1.0 INTRODUCTION	6
2.0 PROJECT METHODOLOGY	6
2.1 SELECTING SAMPLE PAGES	6
TABLE 1. COMMON BOOK ILLUSTRATIONS OF THE 19TH AND 20TH CENTURIES	7
2.2 CHARACTERIZING THE ATTRIBUTES OF DIFFERENT ILLUSTRATION PROCESS TYPES AT VARIOUS LEVELS	10
2.3 MAPPING ILLUSTRATION PROCESS TYPES TO ELECTRONIC CONTENT TYPES	10
2.4 DIGITIZING SAMPLE PAGES	11
2.5 EVALUATING SAMPLE IMAGES	11
2.5.1 <i>Structure</i>	12
2.5.2 <i>Detail</i>	12
2.5.3 <i>Essence</i>	13
2.5.4 <i>Print Evaluation</i>	14
2.6 INVESTIGATING METHODS FOR AUTOMATIC DETECTION OF ILLUSTRATION CONTENT REGIONS	15
2.6.1 <i>Introduction</i>	15
2.6.2 <i>Background</i>	15
2.6.3 <i>A Possible Strategy for Identifying Illustration Regions</i>	15
2.6.4 <i>Document Understanding</i>	16
2.6.5 <i>Detecting Text Regions</i>	16
2.6.6 <i>Conclusion</i>	16
2.7 INVESTIGATING AUTOMATIC METHODS TO DISCRIMINATE DIFFERENT ILLUSTRATION PROCESS TYPES.....	17
2.7.1 <i>Introduction</i>	17
2.7.2 <i>Necessary Distinctions</i>	17
2.7.3 <i>Possible Characteristics for Classifying Illustration Processes</i>	17
2.7.4 <i>The Special Case of Halftones</i>	18
2.7.5 <i>Conclusions on Discriminating Illustration Process Types</i>	18
2.8 INVESTIGATING METHODS FOR PROCESSING DIFFERENT ILLUSTRATION TYPES.....	20
2.8.1 <i>Introduction</i>	20
2.8.2 <i>Possible Image Processing Steps</i>	20
2.8.3 <i>General Concepts</i>	21
2.8.4 <i>Common Processing Steps</i>	21
2.8.5 <i>Halftone Illustration Regions</i>	21
2.8.6 <i>Soft Process Illustration Regions</i>	22
2.8.7 <i>Fine-Featured Hard Process Illustration Regions</i>	22
2.8.8 <i>Hard Process Illustration Regions</i>	22
2.9 AN EXAMPLE UTILITY FOR HALFTONE PROCESSING.....	23
2.10 TESTING AND VERIFYING THE PROCESS	25
2.10.1 <i>Full Resolution View</i>	25
2.10.2 <i>Derivative Creation</i>	26
2.10.3 <i>Printing</i>	27
2.10.4 <i>Other Observations</i>	27
2.10.5 <i>Compound Documents</i>	27
2.10.6 <i>Desirable Enhancements to the Utility</i>	28
3.0 CONCLUSIONS	28

REFERENCES.....	30
APPENDIX 1. ADVISORY COMMITTEE MEMBERS.....	31
LIBRARY OF CONGRESS AFFILIATION.....	31
CORNELL UNIVERSITY AFFILIATION	31
APPENDIX 2. HALFTONE UTILITY USER'S MANUAL	31

Abstract

The Cornell University Library Department of Preservation and Conservation and Picture Elements, Incorporated undertook a joint study for the Library of Congress to determine the best means for digitizing the vast array of illustrations used in 19th and early 20th century commercial publications. This work builds on two previous studies. A Cornell study[1] characterized a given illustration type based upon its essence, detail, and structure. A Picture Elements study[2] created guidelines for deciding how a given physical content region type should be captured as an electronic content type. Using those procedures, appropriate mappings of different physical content regions (representing instances of different illustration processes) to electronic content types were created. These mappings differed based on the illustration type and on the need to preserve information at the essence, detail, or structure level. Example pages that are typical of early commercial illustrations were identified, characterized in terms of the processes used to create them (e.g., engraving, lithograph, halftone), and then scanned at high resolutions in 8-bit grayscale. Digital versions that retained evidence of information at the structure or process level were derived from those scans, which for many illustration types required high resolution to capture. A general consensus was reached, however, that 400 dpi 8-bit capture could serve to preserve the essence and detail information present in all the illustration types studied, regardless of the production process used to create the published originals. This recommendation represents a good cost-benefit requirement for imaging when process identification is not an absolute requirement and in circumstances where mass-produced books, containing both illustrations and text, are to be converted. Project staff investigated the available means for automatic detection of illustration content regions and methods for automatically discriminating different illustration process types, and for encoding and processing them. A public domain example utility was created and tested, which automatically detects the presence and location of a halftone region in a scan of an illustrated book page and applies special processing to it.

1.0 Introduction

In 1998, Cornell University Library's Department of Preservation and Conservation and Picture Elements, Incorporated conducted a joint study for the Library of Congress to determine the best means for digitizing the vast array of illustrations found in 19th and early 20th century commercial publications. This project was intended as the first step in the development of automated means for detecting, identifying, and treating each illustration process type in an optimal manner to create electronic images that can rival the quality of analog capture. The technology does not currently exist to do this. In fact, no thorough attempt has been made to even characterize all of the features of importance for illustrations produced by a given production process, from the point of view of high-fidelity digital image capture.

The Illustrated Book Study had a number of key objectives:

- Select representative samples of relief, intaglio, and planographic illustration processes prevalent in book production in the 19th and early 20th century.
- Characterize the key attributes of different illustration process types by subjective examination, identifying significant informational content for each type at three levels: essence, detail, and structure (see 2.2).
- Develop appropriate mapping of illustration content types to electronic content types that preserve their essential features to an appropriate degree.
- Investigate methods for automatic detection of illustration content regions.
- Investigate automatic methods to discriminate different illustration process types.
- Investigate methods for processing different illustration process types.
- Create an example utility for halftone detection and processing.
- Report project results to the Library of Congress and to the broader preservation community.




As a result of this study, recommendations on digital capture have been advanced for use in preservation reformatting of the range of book illustrations typically found in commercial publications from the past two centuries. The groundwork has also been laid for further development that could lead to fully automated processing of such illustrations to ensure high fidelity to the original. Such automated processing will be exceptionally useful during the next decade as cost-effective, high-quality production scanning will be needed to capture these materials for inclusion in electronic libraries.




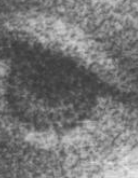
2.0 Project Methodology



2.1 Selecting Sample Pages

An Advisory Committee of Cornell University and Library of Congress curators, faculty, and other experts in printmaking and the graphic arts played a critical role in the selection process (see Appendix 1 for names). Project staff consulted a number of extremely useful publications [3-6] in assembling a group of books and journals containing known illustration types from Cornell University Library's circulating collection. From this grouping, the Advisory Committee chose nine examples that represented the range of printmaking techniques prevalent in the 19th and early 20th century commercial book trade. These included a wood engraving, a halftone, steel and copper engravings, an etching, a mezzotint, a photogravure, a lithograph, and a collotype. All examples appeared in bound volumes, either as separate plates or as illustrations on a text page, and they varied in size, level of detail, and sophistication of technique. The following table characterizes the attributes of these illustration types.

Table 1. Common Types of Illustration in 19th and 20th Century Books

ILLUSTRATION TYPE	PREVALENCE	ILLUSTRATION CHARACTERISTICS	CHARACTERISTICS OF SPECIFIC EXAMPLE
RELIEF PRINTING	Most Common Method of this Era	<ul style="list-style-type: none"> • Raised printing surface; image created subtractively • “White” surface removed to leave “black” printing surface • Matte or glossy paper; separate plate or presented in text; little to no tonal variation 	
Wood Engraving[7] 	1400s- 1890s; most prevalent illustration process in letterpress until introduction of halftones	<ul style="list-style-type: none"> • Created along the end grain of wood • Technique permits finer detail than woodcut • Line width varies. • Ink appears darker around edges • Illustration quality varies depending upon where it is in the press run. 	<ul style="list-style-type: none"> • Typical of 1860s school • Carefully tooled • Presented on text page • Matte paper • .04 mm feature size
Halftone[8] 	1880s-Present	<ul style="list-style-type: none"> • Photo-mechanical reproduction process • Regularly spaced dots of variable sizes • Ridges of ink along dot edges • Poor reproduction of detail • Common screen rulings, 110-200 	<ul style="list-style-type: none"> • Halftone of a painting • Presented on text page • Glossy paper • 166 screen ruling at 45 degrees
INTAGLIO PRINTING	1400s-Present	<ul style="list-style-type: none"> • Recessed printing surface: “black” areas removed to create grooves to hold ink • Tonal variation created by groove size and depth; separate plate or presented in text • Illustration quality depends on print run 	
Steel Engraving[9] 	1820s-Present	<ul style="list-style-type: none"> • Metal removed to create lines; finer detail than wood engraving and larger print runs than copper • Lines are fine, uniform, smooth, and parallel, with crisp edges that tend to be tapered at the end; cross hatching represent mid-tones. 	<ul style="list-style-type: none"> • Typical example; includes some etching • Presented on text page • Matte paper • .02-.04 mm feature size

INTAGLIO PRINTING	1400s-Present	<ul style="list-style-type: none"> • Recessed printing surface: “black” areas removed to create grooves to hold ink • Tonal variation created by groove depth; separate plate or presented in text • Illustration quality depends on print run 	
Copper Engraving[10] 	1700s-1880s	<ul style="list-style-type: none"> • Lines are fine, uniform, smooth, and parallel, with crisp edges that tend to be pointed at the end; cross hatchings represent mid-tones • Difficult to distinguish from steel engraving. • Softer than steel; large print runs show signs of plate wear, with loss of fine lines 	<ul style="list-style-type: none"> • Topographical scene • Separate plate; no plate mark • Matte paper, covered with protective sheet • .04 mm feature size
Etching[11] 	1600s-1880s	<ul style="list-style-type: none"> • Illustration drawn with needle on wax or gelatin covered plate which is etched by dipping in acid. • Lines characterized by blunt ends, width varies as result of acid dips • More free-form than engravings 	<ul style="list-style-type: none"> • Separate plate; plate mark is present • Matte paper • Etching with dry point • .02-.06 mm feature size; most .04 mm
Photogravure[12] 	1880s-Present	<ul style="list-style-type: none"> • Virtually continuous tone photo-mechanical reproduction • Varied amounts of ink on page offer excellent reproduction of detail, mimics tonal variation • Extremely fine grid screen of soft, ragged dots or irregular grain, like confectioner’s sugar 	<ul style="list-style-type: none"> • Representation of photograph • Separate plate; plate mark present • Matte paper, covered with protective sheet • Under .01mm feature size
Mezzotint[13] 	1780s-1870s	<ul style="list-style-type: none"> • Plate surface is roughened to a texture of fine sand paper • Surface then burnished to produce lighter tones • Irregular sandy grain structure; occasional linear pattern detected. 	<ul style="list-style-type: none"> • Typical example • Separate plate; plate mark present • Matte paper, covered with protective sheet • .01 mm feature size

PLANOGRAPHIC PRINTING	1820s-Present	<ul style="list-style-type: none"> • Flatness of both paper and ink, no plate marks • Wide tonal appearance possible. • Matte or glossy paper 	
Lithograph[14] 	1820s-Present	<ul style="list-style-type: none"> • Image transferred or drawn directly on the printing surface • Drawing substance must be greasy • Irregular pebbly grain structure; appears as crayon on coarse paper 	<ul style="list-style-type: none"> • Good example of process • Separate plate • Matte paper • .04 mm feature size
Collotype[15] 	1870s-1910	<ul style="list-style-type: none"> • Virtually continuous tone photo-mechanical reproduction • Telltale irregular and fine cracks (reticulation). • Process is used where accuracy of tone is important; excellent detail rendering 	<ul style="list-style-type: none"> • Collotype of an engraving • Printed on separate paper, trimmed, and pasted into the book • Glossy paper, covered with protective sheet • .01 mm feature size

2.2 Characterizing the Attributes of Different Illustration Process Types at Various Levels

Determining what information in an original artifact should be represented in a digital reproduction is a subjective decision that must be based on a solid understanding of the nature and significance of the material to be converted. Advisory Committee members characterized the key attributes of commercially reproduced versions of the different illustration processes, and assessed the significant informational content that must be conveyed by an electronic surrogate to support various research needs. They articulated the telltale characteristics of the various relief, intaglio, and planographic production processes reviewed, and their descriptions have been summarized in Table 1. Finally, the Advisory Committee was also asked to reflect on the intended uses of the sample documents in the context of their having been issued as part of larger published works rather than as individual pieces of art.

Three levels of presentation were determined:

- **structure**: representing the process or technique used to create the original. The level required for a positive identification of the illustration type varies with the process used to create it. For instance, it is easy to make a positive identification of a woodcut or a halftone with the unaided eye. The telltale reticulation of a collotype, however, may only be observable at magnification rates above 25x.
- **detail**: representing the smallest significant part typically observable close up or under slight magnification, e.g., two times, again a psycho-visual determination.
- **essence**: representing what the unaided eye can detect at a normal reading distance. This view is based on the psycho-visual experience of the reader rather than any feature associated with the source document.

2.3 Mapping Illustration Process Types to Electronic Content Types

Once the various illustration process types had been characterized subjectively, project staff then sought to represent these attributes objectively, e.g., by measuring the spatial extent of the finest lines. The next step involved translating the objective measurements of the original illustrations into similar assessments that pertain directly to the electronic version.

Digital imaging is a process of representing an original document by sampling and mapping it as a grid of uniform dots or picture elements (pixels). Each pixel is assigned a tonal value and represented as a digital number. Conventional wisdom regarding full capture is to have one to two pixels span the finest feature.

Digital requirements to reflect the structure view were predicted by measuring the finest element of the various print processes, which was easy to do for those characterized by well defined, distinct edge-based features, including the engravings, the etching, and the halftone. Despite differences in their identifying characteristics, project staff measured features ranging from .02 mm to .06 mm in size, with the majority of them measuring .04 mm. Evidence of the collotype structure was found in microscopically thin reticulation lines, measuring .01 mm or finer. For those items that were continuous tone-like, exhibiting soft grainy, dotted, or pebbly structures (e.g., the photogravure, mezzotint, and lithograph), feature details were hard to characterize and measure. Feature size estimates ranged from .04 mm to below .01 mm.

Based on the feature size measurements taken at Cornell, we quickly concluded that the resolution required to faithfully represent the structural characteristics would overwhelm any scanning project involving commercially produced publications. At a minimum, the resolution needed to preserve structural evidence in the digital surrogate, calculated at one pixel/feature, ranged from 635 dpi to over 2,500 dpi.

Predictions of digital requirements for the essence view were based on what a person with 20/20 vision could expect to discern under normal lighting at a reading distance of 16 inches. According to optometrists, such a person can distinguish a small letter “e” subtending 5 minutes of arc at that distance. The “e” comprises five parts, each represented by 1 minute of arc. A minute of arc equals 1/60 of a degree, or .01667 degrees. To determine the size of the smallest feature discernible at 16 inches, the following formula is used: $x/16 = \tan (.01667)$; so $x = .004656$ inches. This means that a person with 20/20 vision can detect features as fine as 1/215th of an inch (118 micrometers) at a 16 inch distance. Brian Wandell makes

reference to studies showing that, at high ambient light levels, the highest detectable spatial frequency is 50 to 60 cycles per degree (cpd). Since there are 60 minutes of arc in one degree of arc, this says we need two digital samples (one for the black bar of the cycle and one for the white bar of the cycle) in one minute of arc or 120 of them in one degree of arc. This is reasonably consistent with the optometrists' metric, which is based on visual perception under normal light conditions.[16]

These human visual capabilities suggested that a reasonable digital requirement for an on screen view representing the essence of a page would be 215 dpi. Predictions of digital requirements for the detail view were pegged at 2x magnification, which would require a digital resolution of 430 dpi. The resolution required to produce a print equivalency was estimated to be higher because printing is a notorious "quality sink." [17]

2.4 Digitizing Sample Pages

Each sample page was scanned at a variety of resolutions with 8-bit grayscale data captured. Grayscale data is essential to reproduce the subtleties of perceived tonality inherent in many of the illustration types. It also permits accurate representation of fully bitonal features (having little tonality) when the feature size decreases toward the size of the image sampling function. Grayscale images allow various techniques used by skilled illustration artisans to have the intended tonal effects. For example, grayscale can preserve the modulation of the acid bite in an etching or the variation of the depth of a gouge in an engraving. Grayscale further permits the production of reduced-resolution images from a high-resolution original by means of accurate scaling algorithms.

All illustrations were captured at a fixed spatial resolution of approximately 24 dots per millimeter (600 dots per inch) with an attempt made to capture the entire page that contained the illustration. These full view images were captured on a PhaseOne PowerPhase camera back having a 7,072 pixel moving tri-linear color CCD array. A Hasselblad camera body and Zeiss lenses were used, with a TG-1 filter intended to produce a photopic-like response from the array's wavelength characteristics and the tungsten lighting (using ENH-type reflector bulbs). A color balanced grayscale output was created by the PhaseOne system from the red, green, and blue inputs. For those finely inscribed or continuous tone-like illustrations, a high magnification was used to capture close-up views of their structure. These zoom images were captured on a Kodak Ektron 1400 series camera, having a 4,096 element moving linear grayscale CCD array. Nikon 35mm enlarging lenses and extension tubes were used.

2.5 Evaluating Sample Images

Project staff at Cornell evaluated these images on screen to make a preliminary assessment on resolution requirements for representing the structure of the different illustration processes. They determined that in the case of the relief printing examples (the wood engraving and the halftone) the full view images successfully represented the structure. For the intaglio and planographic illustrations, the zoom images were needed to represent structural evidence. A set of images at lower spatial resolutions (ranging from 200 dpi to 600 dpi) was created from these source images by a process of bi-cubic scaling. Project staff prepared two views of these images to make preliminary judgements regarding the essence and detail representation. View 1 presented image segments at their native resolutions in a 100% view (1:1). For the second view, the lower resolution images were resampled up to 600 dpi using bi-cubic interpolation, a scaling procedure that predicts a new value between two real pixels based on more than the immediately adjacent pixels. The resampled images allowed reviewers to assess images that were the same size on screen. The staff concluded that the 200 dpi versions represented the essence view in all cases, and that the detail view was represented somewhere between 300 and 500 dpi.

Sample images are located at <http://www.library.cornell.edu/preservation/illbk/AdComm.htm>.

The Advisory Committee met several times, both in Ithaca, NY and Washington, DC, and assessed the digital surrogates at the three levels of view, comparing them to the original illustrations with and without magnification, and to printouts created from the essence and detail images.

2.5.1 Structure


In most cases, the Advisory Committee agreed with the project staff's judgement regarding structure representation, but noted that the concept could represent two meanings. The first interpreted "structure" as a view that allowed for identification of process type; the second required a view that faithfully replicated the sample under review. The resolution demands for the latter are much higher. For instance, it is easy to identify a halftone, even at relatively low resolutions. In the examples presented on the Web site, the halftone pattern is evident in the 300 dpi view. Representing the exact dot shape and ruling of the original 166 lines per inch (lpi) halftone placed at a 45 degree angle, however, required a 600 dpi—or perhaps a 900 dpi representation. The Advisory Committee also noted that it may be difficult to differentiate between similar process types even at high resolution without additional testimonial evidence conveyed by the original artifact. These include date of publication, creator's name, whether the illustration appears on a separate plate or paper stock, and whether there was evidence of a plate mark. Finally, committee members felt that process identification for the softer edged images required both close examination and a pull back view to reflect on the nature of the overall composition. For instance, identification of the lithograph process relied on assessing the crayon-like appearance of the representation as well as examination of the pebbly grain structure revealed at higher resolutions or under magnification.

In conclusion, most members of the Advisory Committee determined that digital images could provide good evidence of structure, but at the price of very high-resolution image files. A number suggested that while this might be justified for individual artwork or selective samples, this was an impractical expectation in digitizing most commercially produced monographs and journals. One member suggested that a sample of the higher resolution image (e.g., 2,000 x 2,000 pixel clip) could be produced for identification purposes when necessary.

2.5.2 Detail

Advisory Committee members generally agreed that the 400 dpi on-screen view sufficiently captured the detail present in the original when viewed close up or under slight magnification, using a magnifying glass. With two examples—the copper engraving and the etching—some committee members were divided between the 400 dpi and 500 dpi views. Both cases represented intaglio printing with characteristic hard-edged details, which seem easier to judge in terms of accurate representation than the softer featured illustrations. Nonetheless, the committee's judgement regarding the on-screen detail view was remarkably consistent, and varied little with the illustration type.

Detail Represented at 400 dpi



Committee members agreed that 400 dpi 8 bit capture represented a good cost-benefit requirement for imaging when process identification was not an absolute requirement. The value of this approach is that it represents an assessment of close reading requirements that are based on visual perception, not on the informational content of the original materials. This is an important distinction, and suggests a uniform approach to determining conversion requirements for items that contain a broad range of illustration types or that are difficult to quantify objectively. It also represents a reasonable conversion requirement for mixed items, containing both illustrations and text. The complete work can be imaged at the same level, and files post-processed to reflect the best presentation of the informational content—on screen to support various views or printed out to meet readers' needs or used to create an equivalent to a preservation photocopy or microfilm. Where analyzing the print process of the original source is critical to an understanding of the work, the artifact itself should be preserved.

2.5.3 *Essence*

There was broad consensus from the Advisory Committee on the adequacy of the 200 dpi on-screen view to represent the essence of the original. Lower resolution versions—say 70-100 dpi—will provide a fair likeness of the general image content of the original, but will not match the psycho-visual perception of the original at normal viewing distances. Some tradeoff of perception, however, may be justified in cases where the original can be viewed completely on-screen, particularly for users with lower resolution monitors. For instance, a reader could display the complete image at 200 dpi on an 800 x 600 monitor, only if the dimensions of the original illustration did not exceed 4 inches by 3 inches. At 100 dpi, the complete image could be displayed for illustrations whose dimensions did not exceed 8 inches by 6 inches. **In the future, as monitor resolutions increase, the 200 dpi view may become a practical standard for presenting the essence of original graphic illustrations.**

2.5.4 Print Evaluation

After the first Advisory Committee meeting in which consensus was reached regarding the essence and detail views, project staff prepared a variety of printouts at the two resolution levels for review. Prints were created on a Hewlett Packard 4Mv (HP) laser printer at 600 dpi, using the printer's default dithering algorithm to translate grayscale into the bitonal halftone print. Prints were also produced on the Tektronix Phaser 440, a 300 dpi dye sublimation printer, which offers continuous tone printing rather than halftoning. Project staff created prints two ways: first at the native size of the original, and second in an enlarged mode to review the detail present in the file (e.g., simulating a 400 dpi print on a 300 dpi printer). For comparison purposes, project staff also created photocopies of the original illustrations on the 6085 Canon copier used at Cornell to produce preservation photocopies.

Comparative evaluation of the prints generated by the two printers varied depending on the process used to create the originals. The structure of many of the originals was so fine that when viewed without magnification they appeared to contain shades of gray. Their underlying structure—dots, grains, and lines—became obvious only under magnification. In the photogravure, the detail is so fine (evident only at 50x magnification) that the deposit of ink appears translucent, perhaps allowing some of the light of the paper support to shine through, thus introducing a gray appearance to the black medium. Although the Tektronix printer had half the resolution of the HP printer, its ability to produce actual grays resulted in superior print quality. The laser printer relied on a halftoning process to simulate the gray, at a comparatively low 106 lpi enabling the representation of only 33 gray levels. True grayscale representation proved to be most advantageous in generating prints for those illustration types with soft-edged features that appear to have continuous tonal variation. The photogravure, which conveys tonality rivaling photographic prints, was well reproduced even at 150 dpi on the Tektronix. To rival this quality through halftoning would have required a 2,400 dpi printer capable of producing a 150 lpi screen with the 256 gray values fully represented. On the other hand, distinct, hard-edged representations, such as the wood engraving, rely more on resolution than apparent tonal range in conveying information, as demonstrated by the enlarged Tektronix view.

Advisory Committee members found the printed versions noticeably inferior to the on-screen views, but adequate representations of the originals, when considered in the context of a preservation reformatting program for brittle books. With one possible exception, prints generated from the 400 8-bit files produced on either the dye sublimation or laser printers were judged superior to the preservation photocopies made directly from the original illustrations. In the case of the very fine, uniformly inscribed copper engraving, a noticeable moiré appeared in the sky region—nonetheless, project staff favored it over the photocopy. Even the 200 8-bit image produced a better quality print on the laser printer than was obtained via photocopy. A complete set of all prints generated has been supplied to the Library of Congress.

2.6 Investigating Methods for Automatic Detection of Illustration Content Regions

2.6.1 Introduction

The goal of this part of the project was to develop some approaches for the detection of illustrations of any type, not for the discrimination of one type of illustration from another. This is a most useful general goal, especially if similar processing is to be applied to other illustration types.

This part of the project was aimed at developing general approaches to the problem rather than at developing actual working algorithms, with the expectation that future work will tackle the creation of software for this purpose. This is in contradistinction to the portion of the project that developed the halftone utility. It is worth pointing out, however, that the halftone utility assumes it to be known that at least one halftone region is present in the page image on which it is run. For this reason, the methods of this and the next section will be necessary precursors in a truly automatic processing tool chain for illustrated book pages.

The steps in such a tool chain could be as follows:

1. Perform basic common processing steps, for example:
 - Conservative brightening
 - Deskewing
2. Detect/locate illustration content regions (Section 2.6—this section)
3. Identify illustration region process types (Section 2.7)
4. Apply processing steps specific to each illustration type (Section 2.8)
(The halftone utility developed in this project is an instance of a tool fitting into step 4.)

2.6.2 Background

It is typically desirable to handle illustrations in a different manner than text. One reason for this is to permit higher spatial resolutions to be used on the text, since careful rendering of fine character features is key to its legibility. The human eye is quite sensitive to the uniformity of style inherent in a carefully designed typeface across all its characters. If text is rendered with too few samples per stroke, this uniformity is destroyed, with one stroke being one pixel wide and the next being two pixels wide. When a larger number of samples occur across strokes of characters, these variations might be, for example, 10 pixels across one stroke and 11 across the next, a nearly imperceptible difference. The accurate rendition of fine serif features also requires high spatial resolution.

Another reason is that illustration regions more often require grayscale or color (what we will call “multitonal”) electronic representations in order to be reproduced with fidelity than text regions do, for which bitonal data often suffices, especially when the spatial resolution is high enough.

The cost of preserving multitonal data rises sharply as the spatial resolution rises, even when moderate compression ratios (of 8:1 or so) are introduced. Since we want to store these multitonal illustration regions at more moderate spatial resolutions, this also argues for separating the regions, allowing separate treatment.

2.6.3 A Possible Strategy for Identifying Illustration Regions

Given the rich variations seen across illustration process types, one approach is not to detect every possible illustration type, but only to detect regions containing textual content and regions containing background. Then, by exclusion, the remaining non-background regions are declared to be illustration regions. This is the approach we suggest.

2.6.4 Document Understanding

Many approaches to segmentation of mixed content pages have been published in the open literature as part of the domain of research referred to as document understanding. Methods for understanding documents attempt to parse a page image into layout blocks or layout elements. These are objects or groups of related objects having a single purpose in the original layout of the page [18]. Several good surveys of this field exist [19,20].

A variety of top-down or bottom-up methods exist for breaking the various regions of the image into layout blocks. Once this is done, the content within the bounding rectangles of these blocks may be analyzed and classified.

2.6.5 Detecting Text Regions

Many methods for identifying a layout element as textual exist. In most cases, the image is first deskewed and thresholded to yield a bitonal image.

Text regions have a variety of distinctive characteristics. They have a relatively predictable ratio of white to black pixels, often on the order of 10 to 1 in regions mostly containing characters. The statistical distributions of run lengths (counts of adjacent pixels of the same color) for white and black pixels are quite different, with long runs dominating for white pixels and relatively shorter, more consistent lengths (corresponding to strokes) prevalent in black runs.

The correlation (or degree of similarity) between adjacent horizontal scan lines is quite high in text regions, this being the principal way in which the Group 4 MMR (modified modified READ) compression algorithm achieves its results.

Block text regions of Roman characters have a very distinctive texture. They have horizontal spacings between the centroids of black objects which have dominant peaks corresponding to inter-character and inter-word gaps, and vertical spacings which correspond to inter-line gaps. The distribution of inter-word gaps relates to the statistics of the language used. Text lines have highly consistent character baselines and top-lines, with well-known frequencies of occurrence of ascenders and descenders poking through these boundaries (which incidentally can allow the detection of upside-down pages).

2.6.6 Conclusion

To automatically distinguish illustration regions, the best approach seems to be first to perform a general document understanding operation to identify all the layout elements in a page image. Next, each non-background region is examined to determine if it is a primarily textual region. If not, it is classified as an illustration region.

2.7 Investigating Automatic Methods to Discriminate Different Illustration Process Types

2.7.1 Introduction

The next step in automating the processing of illustrated book pages is to discriminate among those illustration types requiring different electronic treatments. By studying the statistical and morphological details of the various illustration types that must be discriminated, characteristic signatures in the electronic images of the captured example pages were sought that allow for the classification of a given illustration's type with some degree of accuracy. Although these methods have not been implemented in software during this project, the methods described are designed with such a future implementation in mind.

As discussed in the conclusions section below, a surprising result has been that very few distinctions among process types are needed, making some of the following discussion of academic interest.

2.7.2 Necessary Distinctions

There is no need to distinguish between illustration processes that will have the same image processing treatment. The limiting case is that only halftones need to be discriminated from other illustration types. Then we need only classify a given illustration as being halftone or other. A case short of that would further distinguish "hard" and "soft" process illustrations. In our study, "hard" process illustrations included the engravings, etchings, and halftones; "soft" process illustrations included photogravures, mezzotints, lithographs, and collotypes.

At moderate resolutions, the detailed image structures are largely obscured, leaving less information to use in discriminating among illustration types, but also lowering the need to make distinctions, with the softening of features allowing similar treatment for multiple illustration types.

As a general rule, at very high spatial resolutions many more distinctions can be made than at more moderate resolutions. But only the moderate resolution case is of current interest from the standpoint of economic viability and equipment availability. In fact, our investigation into the details of the structure seen in the various illustration process types was hampered somewhat by several related factors: (1) the so-called zoom images acquired with the Ektron camera were not at as high a resolution as desired owing to lens limitations; (2) the poor focus seen in those images (owing to having no live two-dimensional display during setup) made them of little use; and (3) no microscope was available with digital imaging capability.

2.7.3 Possible Characteristics for Classifying Illustration Processes

2.7.3.1 Morphology

Morphology refers to the shape of image structures. Much can be inferred from the subtle details of individual strokes or cuts in an illustration:

- The exact shape of strokes at a crossing of two strokes can tell an observer if the creating process is an additive or a subtractive process; whether one leaves material to hold ink or removes material to hold ink. Similarly, the relative widths of white strokes and black strokes can give an indication of the process type.
- Whether the opposite sides of a stroke remain nearly parallel or subtly convergent or whether they are independently wobbly enables one to tell whether they were made by one cut (as with a gouging or scratching tool) or by two (as with a woodcut).
- Greater wobble in a stroke can indicate a pushed, rather than a dragged, tool. The scale of the small curves in a wobbling stroke can indicate the material type and the speed with which the stroke was made.

- Widths of strokes can indicate the hardness of the material used for a plate (copper versus steel engraving).
- The linear or circuitous nature of a stroke can help establish whether an eroding process (like etching) or a scratching process (like engraving) was used. Smoothness can give related indications.

2.7.3.2 Scale and Texture

Different illustration processes have different dominant scales. The dominant scale is the distance over which its most significant structural feature exists.

A variety of techniques exist in the literature [19,20] for analyzing the scales over which given images have information content; much texture analysis involves comparison of the energy present at different scales.

The periodicity of the structures seen in halftones and machine gravure is very distinctive and amenable to automatic detection. The dimensionality (one- or two-dimensional nature) of the periodicity can assist in distinguishing between these two methods of illustration.

The granular nature of collotypes and etchings produces very distinctive textures. The scale or size of the grains gives clues to the process used. Statistical measures such as the ratio of perimeter to area of closed structures can be used to further characterize one process over another. Reticulating patterns can also be characterized statistically, with measures showing the scale of such a pattern's curvature or the circuitousness of its path.

2.7.3.3 Tonality

Tonality is inherently difficult to quantify; just how many distinguishable shades are present in a given image?

The extent (in terms of the number of distinguishable shades) of swings in tonality and the spatial distance across which these may occur represent very pointed characterizations of the texture and subtlety of a pattern in an image.

The steepness of edges (in terms of the number of distinct transitional tonalities in going from fully on to fully off a stroke) in an image pattern is also a distinguishing feature.

2.7.4 The Special Case of Halftones

Halftones have unique characteristics that are well suited to automated detection. Traditional halftones from the 19th and early 20th centuries (prior to the computer age which opened a wide range of variations [21]) have a periodic structure of varying size dark regions. The orientation of this periodic grid is at an angle to the bottom edge of the book page; this angle is almost always 45 degrees, only rarely are angles of 15, 30 or some other number of degrees used.

More about the method we used for locating halftone regions can be found in the user's manual for the utility found in Appendix 2.

2.7.5 Conclusions on Discriminating Illustration Process Types

Moderate capture resolutions (400 dpi, 8 bit) were recommended as a significant result of this study. This places illustrated book capture into an economically advantageous realm at the edge of today's commercial image capture abilities. While this permits excellent rendition of the essence and adequate rendition of the detail of most commercial book illustrations, it does not provide complete information about their structure.

Without the higher 1,000 or 2,000 dpi resolutions, insufficient information is present to enable most of the automated techniques for discriminating illustration process type discussed in this section to operate. Further, even if a special scanner arrangement could be constructed that permitted small sub-regions of the illustration to be sampled at the higher resolutions for the purpose of characterizing the illustration process type, the 400 dpi data available is insufficient to support the special processing envisioned in most automated techniques.

This sounds grim, so why did we recommend the 400 dpi? Because it works remarkably well. An interesting book by Becker [22] portrays side-by-side many original artists' drawings with the new embodiment they took when they became book illustrations rendered via the methods available in book publishing. Sometimes the original artist, but more often a completely different artisan, rendered the essence and often the detail of the original work into a new medium with a completely different structure.

In an analogous way, the 400 dpi capture performs an optical averaging function that more or less completely obscures the original structure (while also making it unavailable for specialized processing). That averaging function, operating much as does the human eye when presented with excessive detail, does a surprisingly good job of distilling the hidden structural features of the illustration process down into an image having the intended net effect on the viewer. The alternative would be painstakingly capturing a 2,000 dpi image, at huge expense, then running it through a variety of sophisticated algorithms aimed at detecting its structure and inferring its process type, then specially processing it to convert those structures to the best moderate-resolution digital equivalent. Our suspicion is that the latter result would not be significantly better.

In conclusion, at moderate resolutions where no process structure is fully preserved, one approach is that only halftone illustrations need to be treated differently—and therefore reliably detected—with other illustration process types lumped into the “Other” category.

Another possible approach says that the “Other” category needs to be split further into “hard” and “soft” illustration processes. The “hard” category includes only those hard-edged, cutting processes that have feature sizes and dominant scales equal to or larger than the 400 dpi capture resolution limit allowing their edges to remain sharp and well-defined. These images would benefit from an edge-sharpening operation, at least for the 400 dpi master, if not for its derivatives. Very fine-scale hard process illustrations would have poorly defined edges, beyond the roll-off of the modulation transfer function (MTF) and would thus generate a wide range of tonalities or shades, much as the soft process illustration types do. These would be grouped with the soft processes, which would not generally benefit from a sharpening operation.

Techniques exist for distinguishing halftone illustrations from other illustrations, although these have not been implemented in the halftone utility developed under this project. While the current halftone utility runs on 600 dpi data, it is expected that 400 dpi data offers sufficient information both to detect the presence of halftone data in an illustration region and to descreen it. Although the utility's current method for measuring halftone frequency has difficulty at the lower resolutions, it is expected that another approach, perhaps using frequency-domain methods, should be workable with such images.

2.8 Investigating Methods for Processing Different Illustration Types

2.8.1 Introduction

The most appropriate electronic treatment of a given illustration type has two components. The first is the set of parameters that describe the electronic image (such as spatial resolution and grayscale bits per pixel). The second component comprises any additional processing that must be applied to that electronic image to place it into final form for either viewing on a computer monitor or for printing. These two destinations will likely require different processing steps.

When considering how to proceed with this study, it was assumed that very high-resolution images would be available. The first step was to “map” each physical content type (seen embodied on a page) to an electronic content type, i.e., a spatial resolution and a bit depth. Given that the very high-resolution images were not available and that (in any case) it would not be wise to assume they would be available in economically-viable mass conversion projects, the focus shifted to moderate resolution images. At more moderate resolutions, bit depths below 8 bits do not give the payback in space savings to offset the complications they create, so all multitone content was assumed to be 8 bit (color not being considered in the project).

2.8.2 Possible Image Processing Steps

Possible image processing steps can include:

- **Brightening.** Stretching of the brightness values of an image through a look-up table to cause the range of its values to fill the available range possible at its bit depth.
- **Deskewing.** Removal of small angle rotation of the image content relative to the scanning axes.
- **Inverse Halftoning.** Conversion of halftone dots (at some screen spacing) from one electronic image (at some spatial resolution) into multitone pixel values (at possibly yet another spatial resolution). The original electronic image for this process might be either a binary image or a multitone image.
- **Thresholding.** Conversion of a multitone image to a bitonal image based upon either the brightness at a given picture element (global or adaptive level thresholding) or upon the presence of edges of sufficient strength (edge thresholding).
- **Smoothing.** Modification of edge boundaries in a bitonal image to reduce edge raggedness, such as might have been produced by the illustration process, the printing process (paper texture or ink spreading), the digitization process, or the thresholding process.
- **Low-Pass Filtering.** Modification of a multitone image via a convolution kernel to reduce its apparent resolution by removing data at higher spatial frequencies. Sometimes used to remove moiré patterns or other scanning artifacts.
- **Scaling.** Lowering or raising the actual spatial resolution of an image, ideally by using a set of nearby pixels to develop a reasonably accurate prediction of the pixel’s most likely value. Contrast this method with scaling via pixel replication and deletion.
- **Edge Enhancement Filtering.** Modification of a multitone image via a convolution kernel to increase the energy of data at higher spatial frequencies. Sometimes used to enhance “muddy” or blurred originals.
- **Halftoning (or re-halftoning).** Conversion of a multitone image to a bitonal image in such a manner as to retain the impression of multiple tones.
- **Compression.** Insofar as compression may be lossy, it also constitutes a processing step that may affect the image’s appearance in useful or non-useful ways. Appropriate compression types will vary with the electronic image type and its “texture.” For example, halftoned representations of multitone data may require JBIG compression to achieve a sufficiently small compressed size.

- ***Scale-to-Gray.*** When high-resolution bitonal images are destined for a lower-resolution presentation (as on a video monitor), this process moves information from the spatial domain to the tonal domain to preserve legibility. This process is better handled by scaling from a multitone source image.

For each distinguishable class of illustration type (each member of which shares a set of attributes measurable in the originally captured electronic image), a set of appropriate processing steps was identified.

2.8.3 General Concepts

Some processing is applied to the master image. Additional steps are applied to create derivative screen access images and derivative print access images [23].

In the ultimately flexible system—a dynamic scanner—all images are captured at very high resolution and in full color and are analyzed to see how much detail is present, both spatially and tonally. Then a digital master image is developed that captures all that detail without any excess. If color data is present, the color representation is retained. If not, and grayscale content is present, a grayscale representation is retained. If no grayscale content is found, only a bitonal representation is stored. This could occur on a region-by-region basis.

In creating a derivative image, it may be desirable to convert one feature to another, representative one. An example is found in the case of preparing a grayscale image for bitonal printing. The grayscale representation is converted to a bitonal one through halftoning. Similarly, when taking a very high-resolution bitonal image down to a lower spatial resolution, conversion to a grayscale representation may be more appropriate.

Different illustration processes have different dominant scales. The dominant scale is the distance over which its most significant structural feature exists. An appropriate choice of span for a sharpening kernel could depend on knowing this scale.

2.8.4 Common Processing Steps

Conservative brightening, which stretches the tonal range while clipping no values is an appropriate first step for the master page image containing any of the illustration types.

Deskewing using a bilinear or higher-order interpolation is an appropriate next step, particularly if the efficiency of subsequent steps can be improved (document understanding, etc.). It is worth noting that a mild low-pass filtering effect results from this step, so extremely small angle skews are best left untouched.

2.8.5 Halftone Illustration Regions

Halftones have their screen frequency measured and are then descreened, which involves two filtering steps. For more information, see the software user's manual in Appendix 2.

A digital master which is a fully grayscale image is most appropriate. The text regions should stay at the capture resolution or be scaled up, while the descreened halftone region could remain at the capture resolution or could be scaled down. Sharpening should be applied to the textual regions but not to the descreened halftone (which has already experienced a mild sharpening).

Screen access images would probably be prepared by recompositing the page, while scaling the two pieces of content down to the screen resolution, then compressing with a progressive method like progressive JPEG or interlaced GIF.

Print access images could be delivered as PDF files, where the text regions have been scaled up to the native resolution of the printer (600 or 1200 dpi), thresholded by a high-quality process and Group 4 compressed and where the halftone regions have been left as grayscale and compressed using moderate

compression JPEG (8:1 ratio). This allows the print driver or printer controller to decide how best to re-halftone the photo region in light of its knowledge of the print engine's brightening or darkening characteristics. If re-halftoning is performed before sending to the printer, a calibration method should be used to determine the appropriate brightening operation to perform prior to halftoning.

2.8.6 Soft Process Illustration Regions

These processes have such small scale or subtly tonal features that a softened representation is key to their fidelity.

A digital master which is a fully grayscale image is most appropriate. The text regions should stay at the capture resolution or be scaled up, while the illustration region could remain at the capture resolution or could be scaled down. Sharpening should be applied to the textual regions but perhaps not to the illustration region that is experiencing the moderate sampling resolution as a means of creating a softened content. It is possible that a different brightening operation could be applied to the illustration region than to the text regions.

Screen access images would probably be prepared by recomposing the page, while scaling the two pieces of content down to the screen resolution, then compressing with a progressive method like progressive JPEG or interlaced GIF.

Print access images could be delivered as PDF files, where the text regions have been scaled up to the native resolution of the printer (600 or 1,200 dpi), thresholded by a high-quality process and ITU Group 4 compressed and where the illustration regions have been left as grayscale and compressed using moderate compression JPEG (8:1 ratio). This allows the print driver or printer controller to decide how best to halftone the illustration region in light of its knowledge of the print engine's brightening or darkening characteristics. If halftoning is performed before sending to the printer, a calibration method should be used to determine the appropriate brightening operation to perform on the illustration region prior to halftoning.

2.8.7 Fine-Featured Hard Process Illustration Regions

This type of illustration is handled just like the soft process illustrations, since the moderate resolution has converted the illustration's hard, fine features to an impressionistic, softer appearance.

2.8.8 Hard Process Illustration Regions

These illustrations have bold-lined primarily bitonal content that has large enough features to not be washed out by the moderate resolution capture process.

A digital master which is a fully grayscale image is most appropriate. The text regions should stay at the capture resolution or be scaled up, while the illustration region could remain at the capture resolution or could be scaled down. Sharpening should be applied to the textual regions and to the illustration region although the span of the sharpening filters may be different for the two regions. It is possible that a different brightening operation could be applied to the illustration region than to the text regions.

Screen access images would probably be prepared by recomposing the page, while scaling the two pieces of content down to the screen resolution, then compressing with a progressive method like progressive JPEG or interlaced GIF.

Print access images could be delivered as PDF files, where the text regions have been scaled up to the native resolution of the printer (600 or 1,200 dpi), thresholded by a high-quality process and Group 4 compressed and where the illustration regions have been left as grayscale and compressed using moderate compression JPEG (8:1 ratio). This allows the print driver or printer controller to decide how best to halftone the illustration region in light of its knowledge of the print engine's brightening or darkening characteristics. If halftoning is performed before sending to the printer, a calibration method should be used to determine the appropriate brightening operation to perform on the illustration region prior to halftoning.

2.9 An Example Utility for Halftone Processing

Halftones are particularly difficult to capture in digital form, as the screen of the halftone and the grid comprising the digital image will often conflict with one another, resulting in distorted digital image files exhibiting moiré patterns at various scales on computer screens or printers. A method for satisfactorily converting halftones has been most pressing, as the halftone letterpress process became one of the most dominant illustration types used in commercial book runs beginning in the 1880s.

This project has resulted in the development of a practical, working utility to detect the location and characteristics of a halftone region on a page (known to contain a halftone) and appropriately process that halftone region independently from its surrounding text.

Since this utility is not embedded inside a specific scanner, but runs externally on a UNIX server, it may be used on data from any scanner that can supply the appropriate raw bit stream (e.g., unprocessed grayscale of a sufficient spatial resolution).


The documentation for the software is found in Appendix 2. The source code distribution is located at: <http://www.picturel.com/halftone>.

Below is an example of the utility locating the bounding rectangles of six different halftone regions on the same page, followed by an enlarged comparison of the unprocessed halftone.

Detecting Halftone Regions



Processing Halftone Regions



Raw Grayscale Capture



Processed Halftone Information

Some Portable Document Format (PDF) files also have been prepared, which show raw and processed halftone images next to one another, allowing easy experimentation with zoom levels and printing results. These PDF files are found at: <http://www.picturel.com/halftone>.

The left-hand page shows the automatically detected halftone region removed from a 600 dpi page, but without any descreening process applied. The right-hand page of each PDF shows the same area with the descreening process applied. For neither image has any scaling or compression been performed.

2.10 Testing and Verifying the Process

The Cornell University Library Department of Preservation tested and evaluated the prototype utility for halftone processing. Observations from that testing are listed in this section. *Italicized paragraphs in this section represent comments by Picture Elements (the designers of the utility) on some of those observations.*

The evaluation was performed using serial and book publications dating from the 1880s to the 1940s that contained a range of halftones, from 110 line to 175 line screens. With very few exceptions, these halftones represented screens rotated to a 45 degree angle. Some examples represented separate plates, and others were presented within a page of text.

2.10.1 *Full Resolution View*

Cornell:

As the example in the previous section illustrates, the halftone processing utility enables one to “see through” the screen of the halftone to the pictorial content beneath. This process worked equally well on a range of screen rulings, suppressing or smoothing the halftone screen but never entirely eliminating its presence. At higher screen rulings, there is a denser information base for the utility to interpret, and the result is less evident halftone “shadow” in the processed image. The following illustrations demonstrate the halftone processing for common halftone screen rulings, placed at 45 degrees.

120 lpi Halftone



Unprocessed



Processed

133 lpi Halftone



Unprocessed



Processed

150 lpi Halftone



Unprocessed



Processed

The suppression process was most successful on halftones placed at a 45 degree angle. Other angles proved more troublesome, but fortunately these occurred very rarely in 19th century commercial printing.

Picture Elements:

The utility currently assumes that all the halftones to be processed have a screen ruling at a 45 degree angle, owing to the overwhelming prevalence of this choice. While the designers have methods for detecting the screen angle, these are not included in the current version of the utility. Some residual diagonal texture is seen in the processed images. While it is not enough to cause moiré patterns, allowing the design objective to be met, it seems possible that additional work on the filter may reduce it still further.

2.10.2 Derivative Creation

Cornell:

Resampling halftone images introduces the likelihood of moiré patterning from screen frequency interference. This was evident when some of the full resolution images were scaled to 100 dpi to create derivatives for Web access, as illustrated in the following examples. Obviously image processing routines can be used to minimize the introduction of moiré patterns as derivative images are prepared, but typical processes use blurring filters, which do not discriminate between screen rulings, and the results can vary dramatically. Further, this blurring process degrades resolution, which is already compromised in the scaling effort. Note the comparison of 100 dpi scaled images—the one on the left was resampled without using a blur filter; the one in the middle was created using the standard blur filter; and the one on the right was created using the halftone utility.

DERIVATIVE IMAGES



Resampled from original halftone image	Resampled from a blurred halftone image	Resampled from HPU processed image
---	--	---------------------------------------

Picture Elements:

The utility is well suited to the creation of derivative images. By descreening first, moiré patterns are not created during scaling. Rather than simply blurring, the descreen process attempts to filter out the dominant frequency of the halftone screen. This is done by cascading a low-pass filter with a high-frequency emphasis filter, lessening the blurring effect. For most of the halftones, the descreening algorithm produced images that can be sub-sampled at any frequency without moiré patterns. For others, including those placed at angles other than 45 degrees, only about 95% success can be claimed, since faint patterns still appear at a few frequencies, but they are much less pronounced than in the original images, and they do not occur at most frequencies.

2.10.3 Printing

Project staff printed the original image files and the processed image files to determine the effect on print quality. Black and white printers print grayscale images by using a halftone dithering pattern to simulate the gray. The combined influence of the original halftone pattern and the printer's halftone pattern will increase the likelihood of moiré patterns. Removing (or greatly reducing) the screen pattern of the original halftone virtually eliminates this problem. The converted grayscale file is still at the mercy of the printer's resolution and halftoning algorithm, but the additional challenges posed by interpreting the original halftone screen have been nearly eliminated. Results from the HP Mv4 laser printer, which imparts a 106 lpi, produced prominent moiré on the unprocessed halftones screened at 120 and 133, and noticeable moiré on the 150; the processed images printed cleanly, with little to no moiré. Prints were also created on an HP 4500 Color Laser Jet, which imparts a 150 lpi. Similar results were produced, although the moiré patterning was less prominent on the unprocessed images and the processed images came out beautifully.

2.10.4 Other Observations

Cornell:

Although the utility was designed specifically for halftones, it was tested on an engraving, an illustration type that is subject to the similar pattern interference issues as halftones when printed or scaled. The halftone utility detected the engraving and treated it, but the end results were predictably disappointing, in part because the engraving really has no single, constant frequency. Additionally, the utility is designed to smooth out halftone dot patterns, enabling one to simulate the effect of viewing detail from the original illustration, rather than merely a grid of dots. In the case of engravings, the close up view reveals the essential attributes of the process. Because the halftone utility softened the edges of lines and hatchmarkings, the process resulted in an obvious degrading of the detail view.

Picture Elements:

Proper behavior on engravings is not to be expected. The approach used for halftone detection and frequency measurement would likely never yield an accurate result when applied on an engraving. One stated assumption for the halftone utility is that the page image presented to it must be known to contain a halftone. Future versions could be modified to reject the processing of pages which seem to contain no halftones.

Cornell:

The halftone conversion utility is also constrained to rectangular shapes. When confronted with irregularly shaped halftones, the utility processes a rectangular area determined to be part of the halftone region.

If that region contains text, the text is treated as halftone information, which results in a blurring of the characters.

Picture Elements:

The use of rectangular regions for halftone processing represents a simplified approach that is reasonably well justified for older materials. Typographic innovations in the twentieth century have introduced more non-rectangular halftone regions, but rectangular regions are still much more prevalent. The high-frequency emphasis pass in the algorithm helps sharpen up any text that intrudes into the rectangular area.

2.10.5 Compound Documents

While one point of view contends that the halftoning is an essential part of the illustrated page artifact, another holds that the original illustration or photograph that preceded the halftone is of more direct interest. Putting aside these philosophical questions, the practical problems inherent in digitizing and digitally manipulating halftones argue strongly for application of the processing utility to scanned halftones. This leads, however, to a new technical problem—how best to re-aggregate this distinct grayscale image with the balance of the content from the enclosing page.

A possible, but disappointing, approach is to take the descreened halftone image and “rescreen” it for some target output device and then merge this back with the rest of the page which is then entirely bitonal. This

still has moiré problems and does not allow any given printer controller to decide how to optimize the re-halftoning for its print engine. Grayscale data gives that flexibility.

A variety of emerging standards are up to the task of linking a grayscale subregion to an enclosing bitonal page, including ITU-T Recommendation T.44 for color internet facsimile, and the TIFF-FX standard for internet color fax (RFC 2301). We consider yet another one here: Adobe's Portable Document Format, PDF.

PDF permits multiple pieces of image content of varying types to be placed accurately onto an enclosing page. This allows the descreened halftone to remain a grayscale image (optionally at a lower resolution and JPEG compressed—often good choices) while the textual portions of the page are thresholded to bitonal and compressed using Group 4 (ITU-T Recommendation T.6) compression. Substantial space savings are achieved in this way.

Picture Elements has hardware that implements a thresholding operation using the Multiple Scale Thresholding algorithm of the VST-1000 integrated circuit. By this means, the original raw 600 dpi page is thresholded, producing a high-quality bitonal image. The rectangular region(s) found to contain halftones are then “whited-out” or overwritten as all zero values. This bitonal image is then ITU Group 4 compressed.

Using another utility Picture Elements has created, the original page then can be recomposed by laying the grayscale of the descreened halftone region on top of the bitonal text and white background and storing the result as a page in a PDF file.

Some example compound PDF files are found at: <http://www.picturel.com/halftone>.

2.10.6 Desirable Enhancements to the Utility

As in any software project, the designers kept a wish list of ways in which the halftone processing utility could be improved. Since it is offered as public domain source code under the BookTools Project (<http://www.picturel.com/booktools>), others may undertake these enhancements and contribute the resulting improvements back to the community. These include:

- Measure screen angle, modify processing accordingly.
- Handle color halftones, where each of the process inks' screens occurs at a different angle.
- Reject pages which appear to not contain any halftones.
- Modify the screen frequency measurement approach to allow it to work at lower resolutions (300, 400dpi), perhaps using frequency domain methods.
- Minimize boundary effects by using a slightly wider cropping window during processing, then cropping back to proper size.
- Increase processing speeds; this is research code, written for correctness and clarity, not for production throughput.
- Improve filter design to reduce residual diagonal texture still further.

3.0 Conclusions

This study has produced a number of important results. The means for characterizing the key features of book illustrations as they pertain to digital imaging have been developed, and guidelines for assessing conversion requirements recommended. This is especially critical for publications from the mid-19th century to the mid-20th centuries, which were printed on paper that has become brittle. These volumes must be copied to preserve their informational content, and by defining quality requirements for electronic conversion, digital imaging can become an attractive alternative to conventional means of reformatting, such as microfilming and photocopying.

The basic groundwork for preparing an automated method for detecting and processing different illustration types has been prepared, and an example utility for processing halftones developed and tested. The halftone processing utility in particular will be a most welcome addition in the preservation tool kit. One of the major difficulties encountered by institutions converting text based material has been in the capture of halftone information. With its introduction in the late 1880s, halftone printing revolutionized the way illustrations were created in mass publications. Within 20 years, it virtually replaced the used of wood engraving in relief printing, and resulted in an increase in the graphical content of many books and journals.[24]

Library of Congress reports have noted the special problem of printed halftone illustrations, which are prone to distortion and moiré in both the capture and presentation stages. The Library has identified four suggested means for treating halftones, all of which present their own problems.[25] Obviously the ability to automate their treatment in a manner to ensure good capture that is free of distortion would be of tremendous benefit to cultural repositories that are converting late 19th century and early 20th century materials.

Since the halftone utility addresses a vertical slice of the more general problem of distinguishing and appropriately processing a wide range of illustrations, it will likely not perform properly when presented with other illustration types. Nonetheless, this work prepares the ground for characterizing and processing other graphic illustration types. Beyond the scope of this present project, the intent is to later develop additional utilities for processing the remaining illustration types.

This project also facilitates a shift in thinking about how to create the highest possible image quality for a given collection. This new capture architecture has the appropriate raw grayscale or color data collected from any scanner whose document handling capabilities suit the peculiarities of a particular item, such as a bound volume, a 35mm slide, or a 40 inch wide architectural drawing. The scanner choice can be made on the basis of its physical suitability and the quality of its raw image data, without regard to any special processing needs associated with the source document itself. All special processing and manipulation of raw data from these various sources is then performed in an off-line, largely scanner-independent manner by a centralized server we might call a post-processing server. In this way we are not constrained by the variable and inconsistent processing offered within the many different scanners which are needed to overcome the physical peculiarities of each item in a collection. This work will be particularly important in developing the means for capturing bound volumes without having to resort to disbinding or to the creation of photo-intermediates.

References

1. Anne R. Kenney, and Stephen Chapman, *Digital Imaging for Libraries and Archives*, Cornell University Library, Ithaca, New York, 1996.
2. Picture Elements, Inc., *Guidelines for Electronic Preservation of Visual Materials*, Library of Congress Preservation Directorate, Washington, D.C., 1995.
3. Vito J. Brenni, *Book Illustration and Decoration*, A Guide to Research, Greenwood Press, Westport, Connecticut, 1980.
4. Bamber Gascoigne, *How to Identify Prints*, Thames and Hudson, New York, New York, 1986.
5. James M. Reilly, *Care and Identification of 19th-Century Photographic Prints*, Eastman Kodak, Rochester, New York, 1986.
6. Geoffrey Wakeman, *Victorian Book Illustration*, Newton Abbot, Norwich, Great Britain, 1973.
7. William Shakespeare, *Shakspeare's Songs and Sonnets*, S. Low, London, England, 1875(?), p. 26.
8. *Harper's Magazine* 99 (1899), New York, p. 55.
9. Samuel Rogers, *Italy; a Poem*, E. Moxon, London, England, 1854, p. 79.
10. William S. Russell, *Pilgrim Memorials, and Guide to Plymouth*, Crosby and Ainsworth, Boston, Massachusetts, 1866, p. 80
11. Philip G. Hamerton, *Chapters on Animals*, Seeley, Jackson, and Halliday, London, England, 1883, p. 144.
12. Ada Earland, *Ruskin and His Circle*, Hutchinson & Co., London, England, 1910, frontispiece.
13. Benvenuto Cellini, *The Life of Benvenuto Cellini*, John C. Nimmo, London, England, 1896, frontispiece.
14. David T. Valentine, *History of the City of New York*, Putnam, New York, New York, 1853, p. 183.
15. Hester Lynch Piozzi, *Dr. Johnson's Mrs. Thrale: Autobiography, Letters, and Literary Remains of Mrs. Piozzi*, Foulis, Edinburgh, Scotland, 1910, p. 89.
16. Brian A. Wandell, *Foundations of Vision*, Sinauer Associates, Inc., Sunderland, Massachusetts, 1995.
17. Michael Ester, *Digital Image Collections: Issues and Practice*, The Commission on Preservation & Access, Washington, DC, 1996, p. 6.
18. L.H. Sharpe, II and Basil Manns, "Document Understanding Using Layout Styles of Title Page Images," in Donald P. D'Amato, Wolf-Ekkehard Blanz, Byron E. Dom, Sargur N. Srihari, editors, *Machine Vision Applications in Character Recognition and Industrial Inspection*, Proc. SPIE 1661 (1992), pp 50-60.
19. H.S. Baird, H. Bunke, and K. Yamamoto, editors, *Structured Document Image Analysis*, Springer-Verlag, Berlin, 1992.
20. Lawrence O'Gorman and Rangachar Kasturi, *Document Image Analysis*, IEEE Computer Society Press, Los Alamitos, CA, 1995.

21. Robert Ulichney, *Digital Halftoning*, MIT Press, Cambridge, MA, 1987.
22. David P. Becker, *Drawings for Book Illustration: The Hofer Collection*, Harvard University Department of Printing and Graphic Arts and the Houghton Library, Stinehour Press, 1980.
23. Louis H. Sharpe, D. Michael Ott, Carl Fleischhauer, *Library of Congress Manuscript Digitization Demonstration Project, Final Report*, October 1998. <http://lcweb2.loc.gov/ammem/pictel>.
24. David Clayton Phillips, *Art for Industry's Sake: Halftone Technology, Mass Photography and the Social Transformation of American Print Culture, 1880-1920*, dissertation, May 1996, <http://pantheon.yale.edu/~davidp/halftone/index.html>.
25. Carl Fleischhauer, *Digital Formats for Content Reproductions*, July 13, 1998, <http://memory.loc.gov/ammem/formats.html>.

Appendix 1. Advisory Committee Members

Library of Congress Affiliation

Mark Dimunation, Chief, Rare Book and Special Collections Division
Carl Fleischhauer, Digital Conversion Technical Coordinator, National Digital Library
Basil Manns, Senior Physical Scientist
Irene Schubert, Chief, Preservation Reformatting Division

Cornell University Affiliation

John Dean, Director, Department of Preservation & Conservation
Nancy Green, Curator of Prints and Drawings, Johnson Museum of Art
Michele Hamill, Photographic Conservator, Department of Preservation & Conservation
Gregory Page, Associate Professor, Art Department
Tatyana Petukhova, Paper Conservator, Department of Preservation & Conservation
John Reps, Professor Emeritus, Department of City & Regional Planning

Appendix 2. Halftone Utility User's Manual

The attached document is included with the source code distribution of the halftone processing utility software at <http://www.picturel.com/halftone> as the file `peiHalfTone.pdf`.

Detection and De-Screening of Halftone Illustrations
Algorithm Description and Simulation User's Manual
Version 1.250 of 12/31/98

Rick Crowhurst

Picture Elements, Inc.

1. Scope

This document describes PEI's halftone illustration detection and de-screening algorithm and gives instructions for using the simulation of the algorithm provided in the software package **pei1.250**.

2. Introduction

When a high-resolution gray-scale image containing halftone illustrations is reduced in resolution by conventional methods, interference between the frequency of the halftone "dots" and the frequency of sub-sampling typically causes the resulting image to have "Moire" patterns that make it visually unacceptable. The algorithm described here detects halftone regions in a 600 dpi image, measures the dot frequency of each, and uses these measurements to "filter out" the dots, producing an image that can be sub-sampled without acquiring Moire patterns.

The software package **pei1.250** simulates the algorithm; it takes gray-scale images as inputs and produces a de-screened gray-scale output image for each halftone illustration detected.

3. Algorithm Description

For purposes of this description, the algorithm is divided into three phases: the detection phase, the frequency measurement phase, and the de-screening phase. The input is a gray-scale image with a resolution of approximately 600 dpi, and the output is a de-screened gray-scale image of the same resolution for each halftone illustration found in the input.

3.1. The Detection and Extraction Phase

The detection phase is represented in Figure 1. This figure, as well as subsequent ones, represents the algorithm in terms of the simulation, so that each rectangle represents a simulation step that reads one or more input files, represented as incoming lines, and produces one or more output files, represented as outgoing lines. Doubled lines are used to indicate that there may be more than one file for each image, and the file suffix assigned to each file is shown.

In order to detect halftone illustrations, the input image is passed through a high-pass filter that responds to the halftone dot pattern, as the figure shows, and the maximum response of this filter in each small neighborhood is used as a measure of "halftone-ness" for that neighborhood. The measure image is then reduced in resolution by a factor of 16, so that each pixel represents a 0.027" X 0.027" region of the original image; next, the resulting reduced measure image is examined using connected component analysis (object tracking) in order to find all connected regions with width and height of at least 0.427" of pixels that exceed a threshold. Then, connected regions whose bounding rectangles overlap are combined, so that a list of disjoint rectangles, each representing a halftone illustration, is produced. Finally, the rectangle list is used to extract the halftone illustrations from the original image.

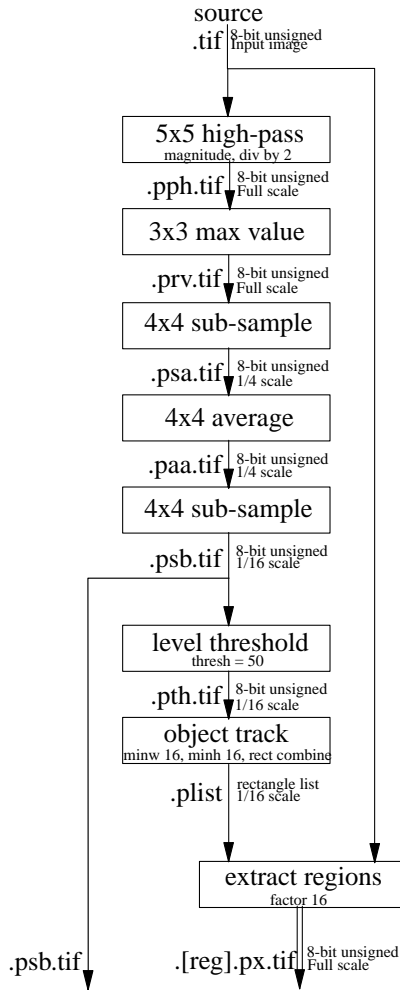


Figure 1
Detection and Extraction of Halftone Illustrations

3.2. The Frequency Measurement Phase

Figure 2 shows the measurement of halftone frequencies that is performed in the second phase. Two one-dimensional frequency measurements are made for each illustration: one is the frequency of dots encountered while scanning the illustration in a left diagonal ("\") orientation, and the other is similarly measured in a right diagonal ("/") orientation. Note that the left and right frequencies could be different due to variations in producing and scanning the original halftone illustration.

In order to determine these frequencies, 8 64 X 64-pixel "samples" are selected from regions of each illustration where the halftone response is especially high. These regions are extracted from the illustration and subjected to filters that "smear" the sample along one diagonal and "sharpen" it along the other, so that the dot frequency can then be determined by counting the zero-crossings encountered while scanning along the "sharpened" diagonal. Statistical methods are used to reject dubious zero-crossings and to validate the measurements. The results from all 8 samples are then combined, producing a left frequency and a right frequency for each illustration.

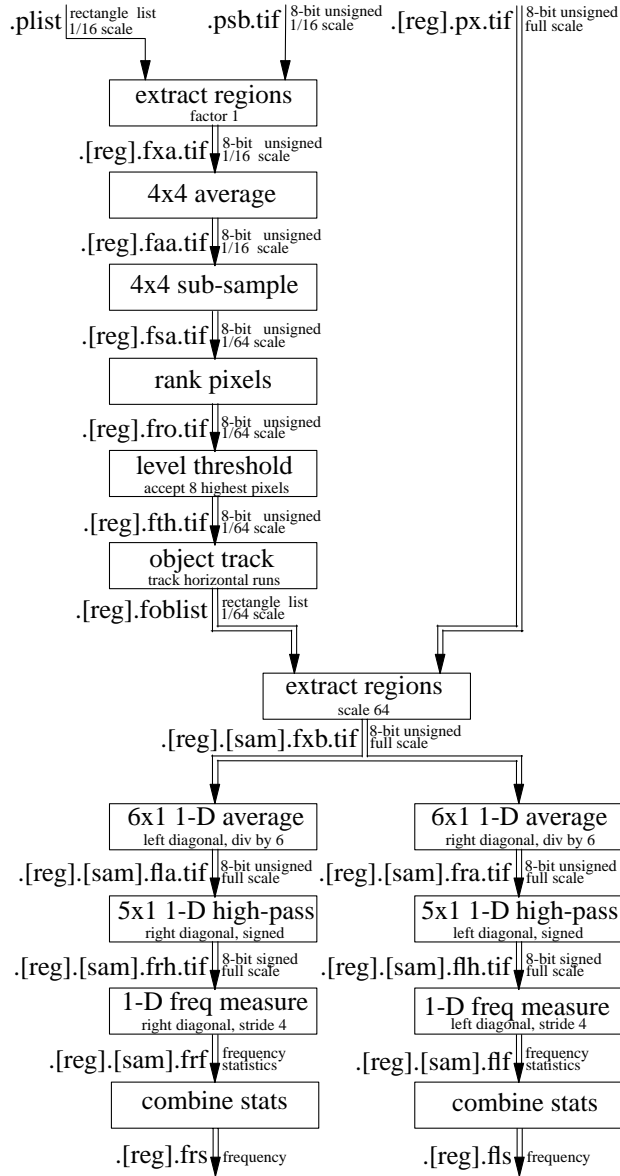


Figure 2
Measurement of Halftone Frequencies

3.3. The De-Screening Phase

Figure 3 shows how the dot frequencies are used to de-screen the halftone illustrations. First, each illustration is filtered by a right-diagonal low-pass filter determined by the right-diagonal dot frequency, then by a left-diagonal low-pass filter determined by the left-diagonal dot frequency; the result is a low-passed image with almost all of the halftone dot component removed. Finally, a pair of diagonal high-pass filters are used to sharpen the image.

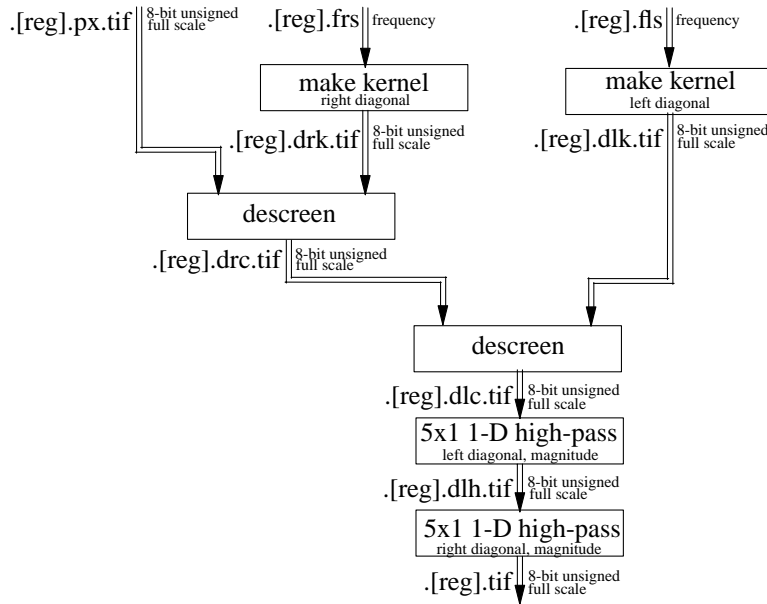


Figure 3
De-Screening of Halftone Illustrations

4. Simulation Description

The simulation is implemented on IBM, HP, Sun, and Linux Unix platforms. It consists of a number of C programs which are called by a C Shell script called **peiHalfTone**; it is distributed as the file **peiX.XXX.tar.gz**, where **X.XXX** is the version number, currently **1.250**.

This section gives Unix command sequences for performing various routine manipulations of the software package. In all cases, it is assumed that the directory **~/pei** is the package's home directory.

4.1. Simulation Requirements

The simulation uses the standard TIFF library **libtiff**, written by Sam Leffler at Silicon Graphics; it can be obtained from

```
ftp://ftp.sgi.com/graphics/tiff/...
```

In order to support input images in TIFF JPEG format, the simulation also requires the standard JPEG library provided by the Independent JPEG Group and available at

```
ftp://ftp.uu.net/graphics/jpeg/jpegsrc.v6.tar.gz
```

If these packages cannot be installed in their canonical places, their libraries and header files can also be put in **~/pei/import/tiff** and **~/pei/import/jpeg**, respectively.

4.2. Unpacking the Simulation Software

To unpack the software package from the distribution file **peiX.XXX.tar.gz** to the directory **~/pei**, copy the distribution file to that directory and execute:

```
cd ~/pei
gunzip peiX.XXX.tar.gz
tar xf peiX.XXX.tar
```

This creates the directory **~/pei/src**.

4.3. Generating the Simulation Executables

After the package has been unpacked, the executables can be generated by the following lines if Sun or Linux (Gnu) **make** is installed:

```
cd ~/pei/src
make
```

If Sun or Linux **make** is not available, the **Makefile** won't work, so a make script is provided on IBM and HP platforms; it is invoked as follows:

```
cd ~/pei/src
make_all
```

4.4. Installing the Simulation

After the executables have been generated, the simulation can be installed by executing the following commands as superuser:

```
cd ~/pei/src
make install
```

or, if Sun or Linux **make** is not available,

```
cd ~/pei/src
make_install
```

This has the effect of copying the executables, including the C Shell script **peiHalfTone**, to the directory **/usr/local/bin**.

4.5. Running the Simulation

After the simulation has been installed, it can be invoked on a list of input images by a command of the form

```
peiHalfTone name1.tif name2.tif ...
```

Each entry in the argument list should be the name of a gray-scale TIFF file and should have a **.tif** suffix. If a file name does not have the **.tif** suffix, the simulation assumes that it's not a TIFF file, so it attempts to convert it to one using Image Alchemy (if Image Alchemy is not present, or some other format conversion utility is preferred, the simulation can be changed accordingly). If the conversion fails, a message is printed and the program continues with the next name.

For each input file *name* that is a TIFF file, or is convertible to one, and for each halftone illustration found in *name.tif*, **peiHalfTone** produces an output file in TIFF LZW format named *name.XX.tif*, where *XX* is a unique positive integer that identifies the illustration.