

Figure 3: Power plots for the Bell Labs Defect Model. Both the reference and probe distributions had the same sensitivity (0.10) and threshold (0.15). The reference distribution had $blur = 1.1$. Notice that the power function is minimum for the probe distribution of $blur = 1.1$. The plots correspond to sample sizes 5 (crosses), 10, 25 (boxes). Note that the (inverted) peaks are sharpest for a sample size of 25 and broadest for a sample size of 5.

mical behavior of these systems
scribed by sets of coupled equa
ations for formal neurons with
estigation of essential features o
bility and adaptation depends ;
the mathematical theory to be
te and efficient analysis of dyna
stract theoretical research in w
jects adopted are frequently as
onical form, the neurodynamic
due to various biological facts
ount of to a degree as large as

(a)

mical behavior of these systems
scribed by sets of coupled equa
ations for formal neurons with
estigation of essential features o
bility and adaptation depends ;
the mathematical theory to be
te and efficient analysis of dyna
stract theoretical research in w
jects adopted are frequently as
onical form, the neurodynamic
due to various biological facts
ount of to a degree as large as

(b)

mical behavior of these systems
scribed by sets of coupled equa
ations for formal neurons with
estigation of essential features
bility and adaptation depends
the mathematical theory to b
te and efficient analysis of dyn
stract theoretical research in
jects adopted are frequently a
onical form, the neurodynam
due to various biological facts
ount of to a degree as large as

(d)

mical behavior of these systems
scribed by sets of coupled equa
ations for formal neurons with
estigation of essential features o
bility and adaptation depends ;
the mathematical theory to be
te and efficient analysis of dyna
stract theoretical research in w
jects adopted are frequently as
onical form, the neurodynamic
due to various biological facts
ount of to a degree as large as

(c)

mical behavior of these systems
scribed by sets of coupled equa
ations for formal neurons with
estigation of essential features o
bility and adaptation depends ;
the mathematical theory to be
te and efficient analysis of dyna
stract theoretical research in w
jects adopted are frequently as
onical form, the neurodynamic
due to various biological facts
ount of to a degree as large as

(e)

Figure 2: Morphological Document Degradation Model. (a) Subimage of the noise free document. (b) Reference degraded document generated with $\alpha = \beta = 1.5$. (c) Probe sample accepted, $\alpha = \beta = 1.7$. (d) Probe sample rejected, $\alpha = \beta = 0.9$. (e) Probe sample rejected, $\alpha = \beta = 2.0$. Sample size used was 60.




Figure 1: Power plots for the Morphological Document Degradation Model. The reference distribution had $\alpha = \beta = 1.5$. Notice that the power function has a minimum near $\alpha = \beta = 1.5$. The power function corresponding to sample size of 60 (boxes), is sharper; that corresponding to a sample size of 10 (crosses) is broader.

- [KH95] T. Kanungo and R. M. Haralick. Morphological degradation parameter estimation. In *SPIE Proceedings*, San Jose, CA, February 1995.
- [KHB⁺94] T. Kanungo, R. M. Haralick, H. S. Baird, W. Stuetzle, and D. Madigan. Document degradation models: Parameter estimation and model validation. In *Proc. of Int. Workshop on Machine Vision Applications*, Kawasaki, Japan, December 1994.
- [KHP93] T. Kanungo, R. M. Haralick, and I. Phillips. Global and local document degradation models. In *Proc. of Second International Conference on Document Analysis and Recognition*, pages 730–734, Tsukuba, Japan, October 1993.
- [KHP94] T. Kanungo, R. M. Haralick, and I. Phillips. Non-linear local and global document degradation models. *Int. Journal of Imaging Systems and Technology*, 5(4), 1994.
- [LLT94] Y. Li, D. Lopresti, and A. Tomkins. Validation of document defect models for optical character recognition. In *Proc. of Third Annual Symposium on Document Analysis and Information Retrieval*, pages 137–150, Las Vegas, Nevada, April 1994.
- [MS88] M. Maltz and J. Szczepanik. Mtf analysis of xerographic development and transfer. *Journal of Imaging Science*, 32(1):11–15, 1988.
- [Nag94] G. Nagy. Validation of ocr data sets. In *Proc. of Third Annual Symposium on Document Analysis and Information Retrieval*, pages 127–135, Las Vegas, Nevada, April 1994.

In figure 3 we give one of the power functions that resulted from this experimental protocol.

5.3 Protocol for Real vs. Real Experiment

In this section we outline the experimental protocol that was used while validating the real-degraded characters against real-real degraded characters.

Real data generated on the computer using Times Roman 8pt font. Text in various European languages was printed using the Adobe Times-Roman typeface, at 8 point, on a Canon laser printer and then scanned at 400 pixels per inch using a Canon scanner. Lower-case ‘e’s were extracted semi-automatically by OCR (thus some possess artifacts resulting from resegmentation). From among these, 3000 were selected by two of the authors, working independently to avoid misclassifications.

The validation parameters used were same as that used for Bell Labs Defect Model. Before selecting the two populations, we randomly shuffled the real data in order to obscure any systematic per-page dissimilarities (due, e.g., to skew scale variations). The validation procedure accepted the null hypothesis that the two samples were from the same underlying population. Repeated trials gave a reject rate close to 0.05, the significance level designed into the test.

6 Summary

We have proposed a statistical methodology for validating document degradation models such as [Bai90, Bai92, KHP93, KHP94]. To illustrate the general applicability of the methodology, we applied it to two different degradation models. Extensive experiments have been conducted on both synthetic and real images. Currently we

are working on applying this methodology to test synthetic against real data for two models. The power function of the validation procedure can be used to estimate model parameters. We conjecture that competing validation procedures can be judged by the discriminations made possible by their power functions.

Acknowledgement: Authors would like to thank Werner Stuetzle and David Madigan for discussions on permutation tests and Tin Kam Ho and David Ittner for comments on the validation procedure.

References

- [Arn90] S. F. Arnold. *Mathematical Statistics*. Prentice-Hall, New Jersey, 1990.
- [Bai90] H. Baird. Document image defect models. In *Proc. of IAPR Workshop on Syntactic and Structural Pattern Recognition*, pages 38–46, Murray Hill, NJ, June 1990.
- [Bai92] H. S. Baird. Document image defect models. In *Structured Document Image Analysis*. Springer-Verlag, New York, 1992.
- [Bai93] H. Baird. Calibration of document image defect models. In *Proc. of Second Annual Symposium on Document Analysis and Information Retrieval*, pages 1–16, Las Vegas, Nevada, April 1993.
- [HS92] R.M. Haralick and L.G. Shapiro. *Robot and Computer Vision (vols. 1 and 2)*. Addison-Wesley Publishing Co., Inc., Reading, Massachusetts, 1992.

figure 1. The power function corresponding to sample size 10 is the widest, and the power function corresponding to sample size 60 is the narrowest. Note all the three power functions give a misdetection (reject) rate close to $\epsilon = 0.05$ when the probe distribution has a parameter value close to that of the reference distribution ($\alpha = \beta = 1.5$). Furthermore, when the $\alpha(= \beta)$ are far from 1.5, the misdetection rate is close to 1, which implies that the validation procedure can distinguish the two samples with high probability. An image generated with $\alpha = \beta = 1.7$ that the validation procedure accepted with a probability close to 0.9, is shown in figure 2 (c). Two images of documents generated with parameter values $\alpha = \beta = 2.0$ and $\alpha = \beta = 0.9$, which the validation procedure could reject the null hypothesis easily, are shown in figure 2 (d) and figure 2 (e), respectively.

5.2 Protocol for Bell Labs Image Defect Model

The parameters of the image defect model that were varied are: the sensitivity (Θ_s), the blurring (Θ_b), and threshold (Θ_t). Furthermore, the spatial quantization error (x and y offsets) were allowed to vary randomly uniformly within $[0, 1]$ (pixels). The rest of the model parameters for the reference and probe distributions were fixed at the following values:

font	=	Times Roman
character symbol	=	the letter ‘e’
resolution	=	400
skew (rotation)	=	0
jitter	=	0
x -scale	=	1
y -scale	=	1 .

For the reference distribution, we generated degraded populations with parameter

values:

$$\begin{aligned}\mu_b &= \{0, 0.3, 0.7, 1.1, 1.5\} \\ \sigma_b &= \{0\} \\ \mu_s &= \{0, 0.05, 0.10\} \\ \sigma_s &= \{0\} \\ \mu_t &= \{0.15, 0.25, 0.35\} \\ \sigma_t &= \{0\}\end{aligned}$$

The probe distribution parameters were varied over the following range of values.

$$\begin{aligned}\mu'_b &= \mu_b + \{0, \pm 0.05, \pm 0.10, \pm 0.30, \pm 0.4\} \\ \sigma'_b &= \{0\} \\ \mu'_s &= \mu_s + \{0, \pm 0.05\} \\ \sigma'_s &= \{0\} \\ \mu'_t &= \mu_t + \{0, \pm 0.05\} \\ \sigma'_t &= \{0\}\end{aligned}$$

The range of model parameter values for the reference distribution were chosen, somewhat subjectively, to cover a range of defects observed in reality (the same range of parameter values were used in the creation the BLidm0 images in the University of Washington English Document Database I).

The validation procedure itself has few parameters that were varied while others were fixed. The inferences derived from the tests are function of these parameters and so have to be specified.

1. Number of repetitions: $T = 100$.
2. Sizes of samples selected from bootstrap pool: $N = M = \{5, 10, 25\}$.
3. Number of permutations, $K = 500$.
4. Significance level of the test, $\epsilon = 0.05$.
5. The character-to-character distance, $\delta(x, y)$, used was the Hamming distance.
6. The set-to-set distance, $\rho(X, Y)$, used was the mean nearest-neighbor distance.

to provide enough information so that anyone can replicate our experiments; and second, to design experiments that will experimentally test the theoretical formulations developed in our previous sections.

There are three types of experiments possible:

Synthetic vs. Synthetic: One reference sample is synthetically created using the image defect model, with a fixed model parameter value. Then many probe samples are generated, again using the model, but with varying parameters. The validation procedure can be run on the probe and reference samples, and the power function generated. This experiment is in part a sanity check for the methodology: if it does not work on synthetic data, there's little point trying it on real data. Also, the parameter estimation methodology can be studied in this way since the true parameter Θ is known and the variance of the the estimated parameter $\hat{\Theta}$ can be calculated.

Real vs. Synthetic: In this case the reference sample constitutes of real degraded characters and the probe sample is generated by varying the model parameter Θ . The validation procedure is run on the probe and reference distribution pairs, and a power function is generated. This experiment tests whether or not the synthetic characters are actually close to the real characters.

Real vs. Real: This could test for systematic dissimilarities between two image populations (e.g. rotations, fonts, etc.). Note that this use is independent of any degradation models.

In this paper we report results on two types of experiments, synthetic vs. syn-

thetic, and real vs. real. Work on the synthetic vs. real experiment is in progress.

5.1 Protocol for Morphological Degradation Model

The following protocol used for creating the reference sample X and the probe sample Y . The reference distribution parameter Θ_r was fixed with the following parameter component values: $\eta_f = \eta_b = 0$, $\alpha_0 = \beta_0 = 1$, $\alpha = \beta = 1.5$, and the structuring element size $k = 5$. The probe distribution parameter Θ_p was varied by varying $\alpha (= \beta)$. Other probe distribution parameter components, $\eta_f, \eta_b, \alpha_0, \beta_0, k$ were same as that in the reference model parameter Θ_r . In all cases the noise free document was the same (a Latex document page formatted in IEEE Transaction style) and the same set of 340 character 'e' (Computer Modern Roman 10 point font) were extracted from the page, for creating the reference population X and the probe population Y .

The validation procedure parameters used were as follows:

1. Sizes of samples, X , and Y : $N = M = \{10, 20, 60\}$.
2. Number of permutations: $K = 1000$.
3. Significance level of the test: $\epsilon = 0.05$.
4. Number of repetitions, T , for computing the power function: $T = 100$.
5. The character-to-character distance, $\delta(x, y)$, used was the Hamming distance.
6. The set-to-set distance, $\rho(X, Y)$, used was the mean nearest-neighbor distance.

The noise-free document is shown in figure 2 (a). The reference degraded document generated with model parameter Θ_r is shown in figure 2 (b). The power function for the sample sizes 10, 20, 60 are shown in

fact could be used to compare our validation procedure with other validation procedures (e.g. [Nag94, LLT94]): the validation procedure with a better power function is better. See [Arn90] for an introduction to power functions.

Distance functions: Various set distance functions $\rho(X, Y)$ can be used for computing the distance between the sets of characters X and Y . We used the following:

$$\rho(X, Y) = (\rho(Y|X) + \rho(X|Y))/(N + M)$$

where,

$$\rho(Y|X) = \sum_{x \in X} \left(\min_{y \in Y} \delta(x, y) \right)$$

$$\rho(X|Y) = \sum_{y \in Y} \left(\min_{x \in X} \delta(x, y) \right)$$

$$\delta(x, y) = \text{HammingDistance}(x, y).$$

The character-to-character distance, $\delta(x, y)$, used was Hamming distance, which was computed by counting the number of pixels where the characters x and y differed after their centroids had been registered. A variety of other character distances such as Hausdorf distance could have been used. Similarly, other set distance functions, $\rho(X, Y)$, could have been used, e.g., squared difference of the set means, the Hausdorf distance, etc. Of course, the choice of distances $\delta(x, y)$ and $\rho(X, Y)$ determines the effectiveness of the validation procedure; we conjecture that the best choice is the one whose power function allows the finest discriminations.

4 Estimation of the Degradation Model Parameters

Given a degraded document we would like to estimate the parameters of the document degradation model, $\hat{\Theta} = (\hat{\alpha}_0, \hat{\alpha}, \hat{\eta}_f, \hat{\beta}_0, \hat{\beta}, \hat{\eta}_b, \hat{k})^t$, that could be used to create degraded documents which are “similar” in the sense discussed above.

We will use the following procedure to estimate the parameter vector $\hat{\Theta}$.

1. Given a fixed sample X of size N .
2. Generate a sample Y of size M and with model parameter Θ
3. Check if the validation procedure accepts the null hypothesis that X and Y come from the same underlying population.
4. Repeat step 2 and 3 K times and estimate the reject rate.
5. Change the parameter Θ of the sample Y and repeat steps 2 through 4, to get a reject rate function.
6. Find the parameter value Θ_0 where the reject rate function is minimum.
7. Θ_0 is the best estimate.

There is a subtle difference between the power function and the reject rate function generated in the estimation procedure. In the validation procedure, the power function is generated by creating new samples of X and Y , in each step. But during estimation, we have only one given sample of X , fixed in all the experiments, while multiple samples of Y are generated. Application of this method for estimating the structuring element size k is discussed in [KH95].

5 Experimental Protocol and Results

In this section we outline the protocol we use to conduct the experiments. Here we give all the sample sizes used, the number of trials that were run at different stages, the exact model parameter values used for the synthetically degraded characters, etc. The purpose of this section is twofold: first,

ate degraded instances $y_1, \dots, y_M \in B$ from the ideal pattern ω , such that the distribution of y_1, \dots, y_M is close to that of x_1, \dots, x_N .

3 Model Validation

In this section we describe a method that can be used to statistically validate the degradation model. Suppose we are given two sequences of degraded characters: “real” images $X = \{x_1, x_2, \dots, x_N\}$, and “synthetic” images $Y = \{y_1, y_2, \dots, y_M\}$. We ask whether or not the distribution of x_i ’s is the same as that of y_i ’s. Specifically, we statistically test the null hypothesis that the distributions are the same. (We will see that, in the design of such a test, it does not matter where the x_i ’s and the y_i ’s came from: either could be real or synthetic, both real, etc.) We now describe a procedure to perform this test.

1. Given (i) real data $X = \{x_1, x_2, \dots, x_N\}$, (ii) synthetic data $Y = \{y_1, y_2, \dots, y_M\}$, (iii) a distance metric on sets, $\rho(X, Y)$, where X, Y are sets of characters. (iv) size of test ϵ , (usually 0.05).
2. Create a new sample $Z = \{x_1, \dots, x_N, y_1, \dots, y_M\}$. Thus Z has $N + M$ elements labeled $z_i, i = 1, \dots, N + M$.
3. Randomly partition the set Z into two sets as follows. Randomly select N elements z_{i_1}, \dots, z_{i_N} as the first set X' , and the rest as the second set, Y' .
4. Compute $d_i = \rho(X', Y')$.
5. Now repeatedly permute the elements of Z , create new partitions X' and Y' and compute d_i . Let us say we make K repetitions.

6. Empirically compute a distribution of d_i ’s as follows $P(d \geq v) = \#\{k | d_k \geq v\} / K$
7. Compute $d_0 = \rho(X, Y)$.
8. Compute the P-value: $p_0 = P(d \geq d_0)$.
9. Reject the hypothesis that the two samples come from the same population if $p_0 < \epsilon$.

Power Functions: If the above procedure is repeated T times, each time with true null hypothesis, the procedure will reject the true null hypothesis, on the average, $\epsilon \cdot T$ number of times. That is, the misdetection rate will be ϵ . In fact, one can generate the *power function*, which is the misdetection rate as a function of the parameter Θ , of the testing procedure as follows.

First, generate the reference sample, X , with the model parameter Θ_r . Now, generate the probe sample, Y , with model parameter Θ_p and compute the reject rate $\gamma(\Theta_p)$, by repeating the validation procedure T times, and computing the fraction of times the hypothesis was rejected. Next, keep Θ_r fixed and keep varying Θ_p and for each value of Θ_p record the reject rate. The plot of the reject rate $\gamma(\Theta_p)$ versus the parameter value Θ_p is the power function. This function should have a minimum at $\Theta_p = \Theta_r$, and should increase on either side and go upto 1 when Θ_p is far from Θ_r . The sensitivity, i.e, the width of the notch, is a function of the sample sizes N and M and the various metrics used. When the sample size is small, the notch is broader and when the sample size is large, the notch is sharper. If we compute the power functions γ_1 and γ_2 for two different validation procedures, for same sample sizes, and $\gamma_1(\Theta) > \gamma_2(\Theta)$ for all values of Θ , then γ_1 is a better validation procedure. Note, by design $\gamma_1(\Theta = \Theta_r) = \gamma_2(\Theta = \Theta_r) = \epsilon$. This

- **jitt**: jitter, the distribution of per-pixel discrepancies of the pixel sensor centers from an ideal square grid (output pixel);
- **sens**: pixel sensor sensitivity, the distribution of per-pixel additive noise (the standard error of a normal distribution with zero mean, in units of intensity);
- **thrs**: the binarization threshold (in units of intensity, where 0.0 represents white and 1.0 black).

The input to the model is an essential perfect bilevel image, derived from artwork purchased from typeface manufacturers, and described at a spatial sampling rate much higher than **resn**. When the model is simulated, the parameters take effect in this order: the input image is rotated, scaled, and translated; then the output resolution and per-pixel jitter determine the locations of the centers of the output pixel sensors; for each pixel sensor the blurring kernel is applied, giving an analog intensity value; per-pixel sensitivity noise is added; finally, each pixel's intensity is thresholded. The output image is bilevel, at spatial sampling rate **resn**.

2 Statistical Problem Definition

In this section we formulate degradation model parameter estimation and model validation as statistical problems. Although degradation of the document is over the entire page, the degradation process itself is local. That is, degradation in one region does not influence the degradation process in another sufficiently far region. More precisely, the degradation at a pixel is influenced only by pixels within a disk of diameter k , which is the size of the disk structuring element used in the morphological

closing process. Thus, one way to characterize the degradation process is to study the degradation of local patterns. Since the most common patterns that occur on a document page are characters, we will statistically characterize the degradation of individual characters on the page and use this characterization to estimate the parameters of a degradation model that produces similar degradations.

Assume that a scanned character is represented by 30×30 matrix with 0 or 1 entries. This matrix can be represented as 1000×1 vector ($30 \times 30 \approx 1000$). Let, B be the space of $D = 1000$ dimensional binary vectors, that is, $B = \{0, 1\}^D$. Now, let $x_1, x_2, \dots, x_N \in B$ be independent and identically distributed D -dimensional vectors representing instances of degraded characters produced from the same class ω . That is, each of these x_i 's were produced from the same ideal pattern ω (say the ideal character 'e') and the same degradation parameters Θ . Now, in our case D is large, typically on the order of 1000. Thus, the number of possible values x_i can take up is 2^{1000} , which is approximately equal to 10^{300} , a dauntingly large number. Available sample sizes, N , are typically on the order of 1000. Thus, samples x_i occupy the space B extremely sparsely.

Two problems we need to address are:

Model Validation: Suppose we are given a set of *real* degraded instances $x_1, \dots, x_N \in B$ of the pattern ω and the another set of *synthetic* degraded instances $y_1, \dots, y_M \in B$ of the pattern ω . Test the null hypothesis that the distribution of y_1, \dots, y_M is same as that of x_1, \dots, x_N , to a specified significance level ϵ .

Parameter Estimation: Suppose we are given a set of degraded instances $x_1, \dots, x_N \in B$ of the pattern ω . Estimate the degradation model parameter $\hat{\Theta}$, which can be used to gener-

1.1 A Morphological Document Degradation Model

The model accounts for (i) the pixel inversion (from foreground to background and vice-versa) that occurs independently at each pixel due to light intensity fluctuations, sensitivity of the sensors, and the thresholding level, and (ii) the blurring that occurs due to the point-spread function of the scanner optical system.

We model the probability of a pixel changing from its ideal value as a function of the distance of that pixel from the boundary of a character. Let d be the distance (four connected or eight connected) of a foreground or background pixel from the boundary of the character and let Θ be the parameters of the model. Let $P(1|d, \Theta, f)$ and $P(0|d, \Theta, f)$ be the probability of a foreground pixel at a distance d from the background to remain as 1 and to change to a 0, respectively. Similarly, let $P(1|d, \Theta, b)$ and $P(0|d, \Theta, b)$ be the probability of a background pixel at a distance d changing to a 1 and remaining a 0, respectively. The foreground and background 4-neighbor distance can be computed using any distance transform algorithm (see [HS92]). The random perturbation process then proceeds to change pixel values in a pixel by pixel independent manner. The following forms for the background and foreground conditional probabilities were used in the model.

$$P(1|d, \Theta, b) = P(1|\alpha_0, \alpha, \eta_b) \quad (1)$$

$$= \alpha_0 e^{-\alpha d^2} + \eta_b \quad (2)$$

$$P(0|d, \Theta, f) = P(0|\beta_0, \beta, \eta_f) \quad (3)$$

$$= \beta_0 e^{-\beta d^2} + \eta_f \quad (4)$$

Here α_0 and β_0 are the initial values for the exponentials; α and β control the decay speed of the exponentials; η_f and η_b are the uniform probability of a foreground and background pixels flipping, respectively. The independent pixel degradation is followed by a morphological closing

operation with a disk of diameter k to account for the correlation introduced by the optical point spread function preceding the thresholding operation which produces the noisy image. Since the closing operation is a nonlinear, it is difficult to model the probability of pixels flipping after the closing operation.

The degradation model parameter vector Θ is a vector of seven parameters, $\Theta = (\alpha_0, \alpha, \eta_b, \beta_0, \beta, \eta_f, k)^t$, where the last entry k is the size of the disk used in the morphological closing operation.

Software for simulating noisy documents using the above degradation model is available from University of Washington English Document Database I. Examples of synthetically degraded document images are shown in figure 2.

1.2 Bell Labs Image Defect Model

The image defect model is described in detail in [Bai90, Bai92]; what follows is a brief summary.

The parameters of the model (and their units) are:

- **size**: the nominal text size of the output (in units of points);
- **resn**: the spatial sampling rate (output pixels/inch);
- **skew**: rotation (degrees);
- **xscl, yscl**: multiplicative scaling factors (horizontally and vertically);
- **xoff**: horizontal translation offset (output pixels);
- **yoff**: vertical translation offset (output pixel);
- **blur**: defocusing, modeled as a Gaussian point-spread function (psf) centered at the pixel sensor center (the standard error of the psf kernel in units of output pixels);

Validation and Estimation of Document Degradation Models

Tapas Kanungo[†], Henry S. Baird[‡], and Robert M. Haralick[†]

[†] Department of Electrical Engineering, FT-10

University of Washington

Seattle, WA 98195 USA

{tapas,haralick}@ee.washington.edu

[‡] AT&T Bell Laboratories

600 Mountain Avenue, Room 2C-322

Murray Hill, NJ 07974 USA

hsb@research.att.com

Abstract

Printing, photocopying and scanning processes degrade the image quality of a document. Baird ([Bai90, Bai92]) and Kanungo, Haralick and Phillips ([KHP94, KHP93]), have proposed document degradation models that model the local distortions introduced during the printing and scanning process. In this paper we propose a statistical methodology that can be used to validate such models. Essentially, we show how to test whether two samples of degraded documents are from the same population or not. Although we demonstrate the methodology on synthetic documents, degraded document populations could in general be (i) both real, (ii) both synthetic (artificially generated), or (iii) one synthetic and one real. This hypothesis testing methodology is independent of the degradation model and can be used to validate any document degradation model. We show how to apply the methodology to two different degradation models. Furthermore, we construct the probability of reject function and use it to the estimate

the parameters for the document degradation model from a degraded document page image.

1 Two Document Degradation Models

In this section we discuss two document degradation models that model the local degradations that occur when documents are printed, scanned and digitized. The first model we describe is the Morphological Document Degradation Model (Kanungo, Haralick and Phillips, [KHP93, KHP94]). The second model is the Bell Labs Image Defect Model (Baird, [Bai90, Bai92, Bai93]). These models have been described in detail in the cited literature, but we briefly describe them here for completeness. There are other models such as [MS88], but we will not be discussing them in this paper.