

Determining Delivery Point Codes on Handwritten Addresses

Sargur N. Srihari, Edward Cohen, Venu Govindaraju and Ajay Shekhawat

CEDAR, 226 Bell Hall
State University of New York at Buffalo
Buffalo, NY 14260.

phone: 716-636-3191

fax: 716-636-3966

e-mail: govind@cs.buffalo.edu

Abstract

A method for determining the delivery point codes (DPCs) for handwritten addresses is described. Determining the DPC requires locating and recognizing address components (*e.g.*, ZIP Code, street number, P.O. box number) and using multiple information sources to assign a five, nine or eleven digit barcode (*i.e.*, the DPC) to an address. Our method uses diverse pattern recognition and image processing algorithms (thresholding, underline removal, separation of lines, location and recognition of address components, and Postal directory access) to determine the DPC.

Three separate encoding techniques are used, each accessing a separate USPS database. There are three methods of determining the DPC: (i) the ZIP code recognized happens to be a unique ZIP Code. In such cases, the 5-digit ZIP Code itself is the DPC; (ii) the address has a P.O. Box number. In such cases, the ZIP Code and the P.O. Box number are sufficient to determine the 4-digit add-on to make up the DPC; (iii) the address has a street line. In such cases, the ZIP Code and the street number confine the possible street names to a small set using a USPS directory. The word recognition algorithm is invoked to find which of the street names in the set matches the street name image in the address. The selected street name has a 4 digit add-on which can be obtained from the Postal directory. The last two digits of the street number are appended to the ZIP+add-on to generate an 11-digit DPC.

1 Introduction

Determining the delivery location for mail pieces based on handwritten addresses is a problem that trained humans can normally solve. As a problem in machine reading and interpretation, it presents many challenges like word recognition, digit recognition, combining high level and low level information, control structure to satisfy mutual constraints, *etc.*

The United States Postal Service sorts mail by utilizing an encoding of each mail piece by a string of digits, the first five of which are referred to as the ZIP Code, and together with the next four as the ZIP+4. The first three digits encode the state, city and postal sectional center facility (SCF), the next two the post office, and the final digits encode destination information such as block face, post office box number, *etc.* The destination encoding can be a five-digit ZIP Code or nine-digit ZIP Code or an eleven-digit delivery point code (ZIP+4 plus an additional two digits corresponding, for example, to the last two digits of the street address).

Examples of address blocks taken from live mail (Figure 1) illustrate some of the challenges posed by unconstrained handwriting and wide variations in address structure. Figure 1a is a very poorly handwritten address block in which the ZIP Code is located in a different line from the city name and the state name. Figure 1b is an example of an address block with guide lines and guide text which must be treated as noise by the address interpreter. Figure 1c shows the apartment number with ambiguous line position. Figure 1d illustrates one of the most challenging tasks in address interpretation. When components from adjacent lines "spill over" so as to touch each other, assigning components to their respective address lines is difficult. Forcing separation of address lines causes components from different lines to interfere in the recognition process.

A complete system for unconstrained address interpretation needs several components. At the pre-processing level, the written matter from the handwritten address block image must be distinguished from the background, and irrelevant streaks such as underlines must be removed¹. The text line locations in the address must be determined, the text lines separated into words, and candidates for the ZIP Code, street number, P.O. Box number, and street name must be determined. The ZIP Code and street number/P.O. Box number are read. The Postal directory is queried with the street number and the ZIP Code to obtain a lexicon of street names. The word recognition algorithm ranks the lexicon based on several features it extracts from the street name image.

The control structure of the system is shown in Figure 2. There are three methods of determining the DPC as described below:

¹ Approximately 5% of the time, the words in a handwritten address are written on machine printed guide lines that are provided to assist the writer. Often the handwritten text intersects these horizontal lines. The text and the lines are then fused during thresholding and the result is an image that is not conducive to connected component analysis for locating the ZIP Code.

1. The ZIP code recognized happens to be a unique ZIP Code. In such cases, the 5-digit ZIP Code itself is the DPC.
2. The address has a P.O. Box number. In such cases, the 5-digit ZIP Code and the P.O. Box number are sufficient to determine the 4-digit add-on with a Postal directory lookup. The 9-digit number becomes the DPC (see example in Figure 3).
3. The address has a street line (street number + street name). In such cases, the ZIP Code and the street number confine the possible street names to a small set with a Postal directory lookup. The word recognition algorithm is invoked to find which of the street names in the set matches the street name image in the address. The selected street name has a 4 digit add-on which can be obtained from the Postal directory. The last two digits of the street number are appended to the ZIP+add-on (9-digits) to generate a 11-digit DPC (see example in Figure 4).

Intermediate Steps

Figures 3 & 4 show some intermediate steps of the Handwritten Address Interpretation system. Figure 3b is a thresholded image obtained from the grey-level image of Figure 3a. The word classifications for each word of the bottom line are considered (c). The first word is classified as *city* because it has text (7-characters) followed by a comma. The second word is classified as "PO" because it is a short word with a component recognized as zero at the end. The third word is classified as *Box* as it is also a short word with 3-characters.

The selected word segmentation results in a bottom-line syntax of *city-3-characters-5-digits* (c) corresponding to the *city-state-ZIP*. Based on this match, the system selects the 5-digits number as the ZIP Code candidate. The ZIP Code candidate is segmented (d) and the isolated digits recognized with confidences shown in (e). The recognized ZIP Code 43613 is checked in the Postal directory and found to be valid. The middle line of the address is parsed to "P.O."-*Box*-4-digits (f) which matches the "P.O."-*Box*-*Box-number* syntax. This same line could be matched to 2-digits-2-characters-4-characters for a street line (*street-number-street-name*), but the P.O. Box line has higher confidence. The Box number is recognized as 5639 (g) with high confidence, and the 4-digit add-on 0639 is determined using a lookup in the Postal directory. The ZIP Code and the 4-digit add-on make up the 9-digit DPC (h).

Figure 4 is an example of the 3rd method of determining DPC. Each text line is separated and word hypotheses within each line are created. After word classification, the top parsing hypothesis for the *city-state-ZIP* is found in the second line from the bottom. One 3-word hypothesis (c) for this line locates *5-characters-state-abbreviation-5-digits* as a syntax match. This match is incorrect, since the first digit of the ZIP Code is included in the state name. This mistake occurs because the ZIP Code contains a connected component (consisting of two large touching digits, 0 and 8) that is estimated to contain three digits (making the

last word a 5-digit word). In addition, the middle word has two medium-sized connected components which provides a good match to the *state-abbreviation* classification. The ZIP Code for this parse is segmented (d) and the digit recognition results are shown in (e) which are not all of high confidence, therefore the ZIP Code candidate is rejected.

The second parsing hypothesis chosen for the *city-state-ZIP* is the correct one. This parse of the same line into three words gives a syntax match of *5-characters-state-abbreviation-6-digits* (*6-digits* match to the ZIP Code is allowed because digit string length estimation was required, and the system knows that estimation errors may occur). The ZIP Code is segmented (g) and all digits are recognized with high confidence (h). The street address line (i) is correctly located by matching the second line from the top to *3-digits-5-characters*. The street number is correctly recognized as 318 with a high confidence (j). Using a Postal directory lookup, the street number, and the 5-digit ZIP Code a dictionary is created (k). This dictionary contains all valid street names in the 01805 ZIP Code whose street number is 318. Their associated 4-digit add-ons (k) are known. The dictionary and the street name image are passed on to the word recognition algorithms which rank the dictionary (l). The confidence levels of the ranked dictionary indicate that *E. Main St.* is the correct street name. A ZIP+4 Code is created by adding the 4-digit add-on to the 5 digit ZIP Code (m). The last 2 digits of the street number (18) are then added to the ZIP+4 Code to create the final DPC (n).

2 Line Separation & Word Classification

After thresholding and noise removal, the first step is separating the text lines in the address block (Figure 2). Connected components are placed into text lines based on positional information. If the connected components cannot be placed into single lines due to insufficient information and/or due to components from adjacent lines touching, they are assigned to a set of possible text lines.

Since non uniform spacing often makes exact word separation impossible, a number of word separation hypotheses are generated for each text line (Figure 5,6). Each grouping is associated with a confidence. The words are classified into one of the following categories: *text*, *city*, *state-abbreviation*, *digits* (along with an estimate of the number of digits in the string), *digits-dash*, *9-digit ZIP+4 Codes*, *barcodes*, *P.O. Box* and *noise*. Next each connected component is identified (based on size and shape alone) as a *dash*, a *comma*, a disconnected *5-hat*² or none of these. In addition, each connected component is identified as a digit or a non-digit. There may be multiple identifications, *e.g.*, a comma and the digit 1 may have identical shapes, but these multiple identifications are reduced after the positional

²Handwritten 5's are often written in two strokes. The vertical bar and the the bottom curve are written in one stroke and the horizontal bar is written as a separate stroke. When the horizontal stroke is not connected to the other stroke, it is considered to be a disconnected 5-hat.

information is considered.

When the positional and classification information are combined, a more accurate interpretation of the text line structure in the address is achieved. At this time, connected components can be moved to different lines, *e.g.*, a comma may be moved from the line to which it is spatially close to the line above. Connected components that span more than one text line may be placed into a single text line or they may be split into two or more components and placed in separate text lines.

3 ZIP Code Recognition

The basic ZIP Code recognition procedure performs the task of locating the ZIP Code within the address, followed by segmenting and recognizing the ZIP Code digits. There are various characteristics that make the task non trivial: the ZIP Code is often on the last line (90% of the cases), but not always; underlines make segmentation of lines and digits difficult; poorly formed addresses without a ZIP Code may be incorrectly identified as having a ZIP Code; the segmentation of a ZIP Code into digits is difficult when digits touch, have connecting ligatures or have overhangs between digits. The digit recognizer needs to be able to handle a variety of writing styles and the artifacts that arise due to segmentation.

Based on statistical analysis of information on mail pieces [5], we have determined the likely syntax constructions to be found in handwritten addresses. Although the most likely syntax of the address has the bottom line containing the city, state, and ZIP Code, many other address syntaxes are possible. For instance, the bottom line may contain the ZIP Code and the second last line may contain the city and state name. Matches between the syntax, word groupings and word classifications are determined and the most consistent matches are examined in order to find likely ZIP Code and state name candidates.

Words that match the syntax of a ZIP Code are selected as ZIP Code candidates. Depending on the number of consistent word classifications, zero or more ZIP Code candidates are selected. The system processes up to three candidates to determine the ZIP Code of the address.

Semantic information is applied to determine the recognition reliability. If all digits are identified with sufficient confidence, the ZIP Code candidate is assigned a confidence value based on recognition confidence of the individual digits and the amount of segmentation required. If a candidate has nine digits, the confidence is increased, because a candidate with nine recognized digits is very likely to be the correct ZIP Code candidate.

Currently the ZIP Code recognition system has a performance of 76.4% correct with 1.2% error rate [6].

Digit String Segmentation

Given the image of a string of digits, the goal of the segmentation algorithm is to partition the image into regions, each containing an isolated digit. A recognition aided iterative method is used.

The segmenter can be invoked in three different modes: (i) estimate the number of digits (street number / P.O. Box number) (ii) force either 5 or 9 digits in the digit string (ZIP Code) and (iii) force a given digit string length. The number of digits is initially estimated from the aspect ratio of the digit string and successive estimates are obtained by a linear regression model. Digits that are recognized with high confidence are removed after each iteration. The effective contribution of the removed digits to the density of the digit string is recorded. This information is fed to a least squares linear model [4]. By setting the density to zero (all digits are removed), the least squares equation can estimate the number of digits in the string. Connected digit components are split into required number of digits. The segmenter has a correct segmentation rate of 93.33% when the input is specified as a 5 or 9 digit ZIP Code, and is 83.03% correct when the number of digits has to be estimated.

Digit Recognition

Handwritten digits have a wide shape variability – from neat printing to broken and touching characters. Digit recognition is performed using a suite of independent and uncorrelated digit recognition algorithms whose decisions are merged using a combination scheme. Four digit recognizers are employed [2]: polynomial discriminant function, a mixed approach classifier, a stroke based recognizer, and a structural contour-based chain-code classifier whose individual performances were 93.6%, 89.4%, 86% and 84.1% respectively on the same test set. The neural network decision combiner achieves 96% accuracy.

4 Street Number and P.O. Box Recognition

The goal is to locate street lines and P.O. Box lines in an address block. The input to the system is a list of connected components in the address block, line number assignment for each component, coordinates and size of each component and the result of digit recognition for each component. Each text line is segmented into clusters of connected components where each cluster can be considered a word. The components in each word are then matched against a set of syntactic descriptions of street lines and P.O. Box lines, and confidence values are assigned to each match.

Words are formed by ordering components of a text line from left to right and selecting word breaks between components. Word breaks are hypothesized based on horizontal spaces

between connected components, a shift in the mean vertical position of connected components, or the locations of punctuation (Figure 5). Confidences of word gaps are based on their size, presence of punctuation, *etc.*

Given a ZIP code and a street number, the following statistical information is available. Under a ZIP Code, the valid street lengths and their distribution is known. Moreover, it is also known as to which digits can occur at various positions of a street number of specified length. For example, ZIP Code 14008 allows only 4-digit street numbers, and in ZIP Code 14069 all 3-digit street numbers begin with '1', all 4-digit street numbers begin with '9', and all 5-digit street numbers begin with '10'. *A priori* statistical information is used to recover from errors in Street number/P.O. Box number recognition.

Active digit segmentation is incorporated into the street number location system. When a potential digit string is examined and it is found that most of the components are digits, although some of the digits may be touching, a digit segmenter is consulted. The segmenter attempts to split the component into digits, and outputs its confidence of the components being good digits. This information is used to refine the street number (P.O. Box) hypothesis. The current performance of the street number/ P.O. Box number location system is 67% accept and 16% error.

5 Word Recognition

The street address line is located by searching for a line in the address that contains digits followed by text. Everything to the right of the street number is assumed to be the street name. Based on statistical analysis of address images, we know that the street address line must be located beneath the top address text line and above the text line containing the city name. If a street line address is found (75% of letter mail have street addresses), the system tries to recognize the street number. If the street number is recognized with a high confidence, a dictionary of possible street names is generated using the National Carrier Walk Sequence (NCWS) directory. The dictionary and the street name image are passed on to the handwritten word recognition algorithms.

Several approaches to word recognition have been explored. One method is based on segmentation and recognition of individual cursive characters in the word [3]. The algorithm follows a *Hypothesis Generate and Reduce (HGR)* paradigm where the lexicon is continually reduced based on the word identity hypothesis at different stages of the algorithm. The segmentation process is actually a search that generates a series of segmentations of the word. Potential characters that result from segmentation are passed to a character recognizer. The system keeps track of the rank of each character as segmentation proceeds. Depending on the desired depth-of-search, the top n characters and their respective segmentations are chosen and the segmentation process continues from that point. This results in an n - way branch at each character candidate segmentation point in the word. The final result is a series of

strings which represent the characters detected from each segmentation sequence of the word. Since many segmentations result in words that contain improbable sequences of characters, most strings are rejected. The remaining strings are passed to a dictionary matching which generates a ranking of the dictionary.

Another approach first divides the word into segments and then recognizes the segment string [1]. The recognition method is based on the *Hidden Markov Model (HMM)*. The observed segments are referred to as symbols and the characters giving rise to the symbols are the hidden states. The Viterbi algorithm and a dictionary are used in the final word recognition.

Ranked outputs of several uncorrelated word recognition algorithms can be combined so as to improve overall performance. The two algorithms mentioned above (HGR & HMM) were considered for combination. Both were tested on images of handwritten city and street names involving a dictionary of words. The size of the dictionary was varied during the experiment. The HGR method resulted in performances (top choice) of 90%, 76%, and 63% with dictionary sizes of 10, 100, and 1000 respectively. On the same dictionaries, the HMM method had performances of 91%, 68%, and 45%. The combination performance was 92%, 87%, and 77%.

6 Performance Evaluation

The system performance was tested on a set of 973 images³. The performance at various error rates (by different thresholding) is shown in Table 1 and Figure 7. We have identified

	Accept	Error
No thresholds	58%	36%
Threshold ZIP Code and Street number	42%	33%
Threshold Word Recognition	30%	14%

Table 1: DPC performance at various thresholds

the major causes of failures (errors and rejects). Rejects are caused by failure to determine a ZIP Code (20%), failure to locate a street number or P.O. Box number (21%), no dictionary found for the street name (5%)⁴ incorrect dictionary generated because of incorrect ZIP Code/ street number recognition (37%), and rejection by word recognition algorithm (1%).

³all addressed to locations inside Buffalo

⁴In these cases, the system correctly recognized the ZIP Code and the street number, but search of the Postal directories revealed no valid entries. We have not been able to determine if these errors are due to incomplete directories or patron errors.

Errors were caused by incorrect recognition of street name (9%), incorrect recognition of street number or ZIP Code (11%), and incorrect dictionary created from a correct street number and ZIP Code (2%). The test of 973 images contained only 7 (1%) P.O. Box numbers (the general mail stream is expected to contain 22%), so P.O. Box number recognition does not figure prominently in the analysis.

7 Summary

Determining the DPCs is a challenging problem. Our system can correctly determine the DPC 37% (Accept rate \times Correct rate) of the time. All the major components of the system are in place. Each of the components is undergoing critical evaluation and refinements. Research is actively underway to develop better techniques of dividing text lines into words using improved spatial measurements and better punctuation detection. The goal is to achieve a 50% DPC rate with less than 1% error rate.

References

1. M. Chen, A. Kundu and J. Zhou. Off-line handwritten word recognition using a single contextual Hidden Markov Model. *Proc. of the Computer Vision and Pattern Recognition Conf.*, Champaign, IL, pp. 669-672, 1992.
2. E. Cohen, J.J. Hull and S.N. Srihari. Understanding handwritten text in a structured environment: determining ZIP codes from addresses. *Int. Journal of Pattern Recognition and Artificial Intelligence*, vol. 5, no. 1&2, pp. 221-264, 1991.
3. J.T. Favata and S.N. Srihari. Off-line recognition of handwritten cursive words. *Proc. of the SPIE symposium on electronic imaging science and technology*, San Jose, CA, 1992.
4. R. Fenrich. Segmentation of automatically located handwritten words. *Proc. Int. workshop in handwriting recognition*, Bonas, FRANCE, pp. 33-44, 1991.
5. J.J. Hull, D.S. Lee and S.N. Srihari. Characteristics of handwritten mail addresses: A statistical study for developing an automatic ZIP Code recognition system. *TR 88-06*, Dept. of Computer Science, SUNY Buffalo, 1988.
6. S.N. Srihari, E. Cohen, J.J. Hull and L. Kuan. A system to locate and recognize ZIP Codes in handwritten addresses. *Int. Journal of Research and Engineering Postal Applications 1*, pp. 37-45, 1989.

Dominican Convent

[REDACTED] Dent St

Buffalo, N.Y.

14211

(a)

QUANTITY	DESCRIPTION
[REDACTED]	[REDACTED]
[REDACTED]	ELMWOOD AVE.
[REDACTED]	Buffalo, N.Y. 14222

(b)

[REDACTED]
[REDACTED] Graystone Apts.
apt [REDACTED] [REDACTED] Johnson Park
Buffalo N.Y. 14201

(c)

91
[REDACTED]
[REDACTED] Elliott Square
Buff, N.Y. 14203

(d)

Figure 1. Examples of handwritten address blocks extracted from actual mail. The names of the addressees have been blackened out intentionally.

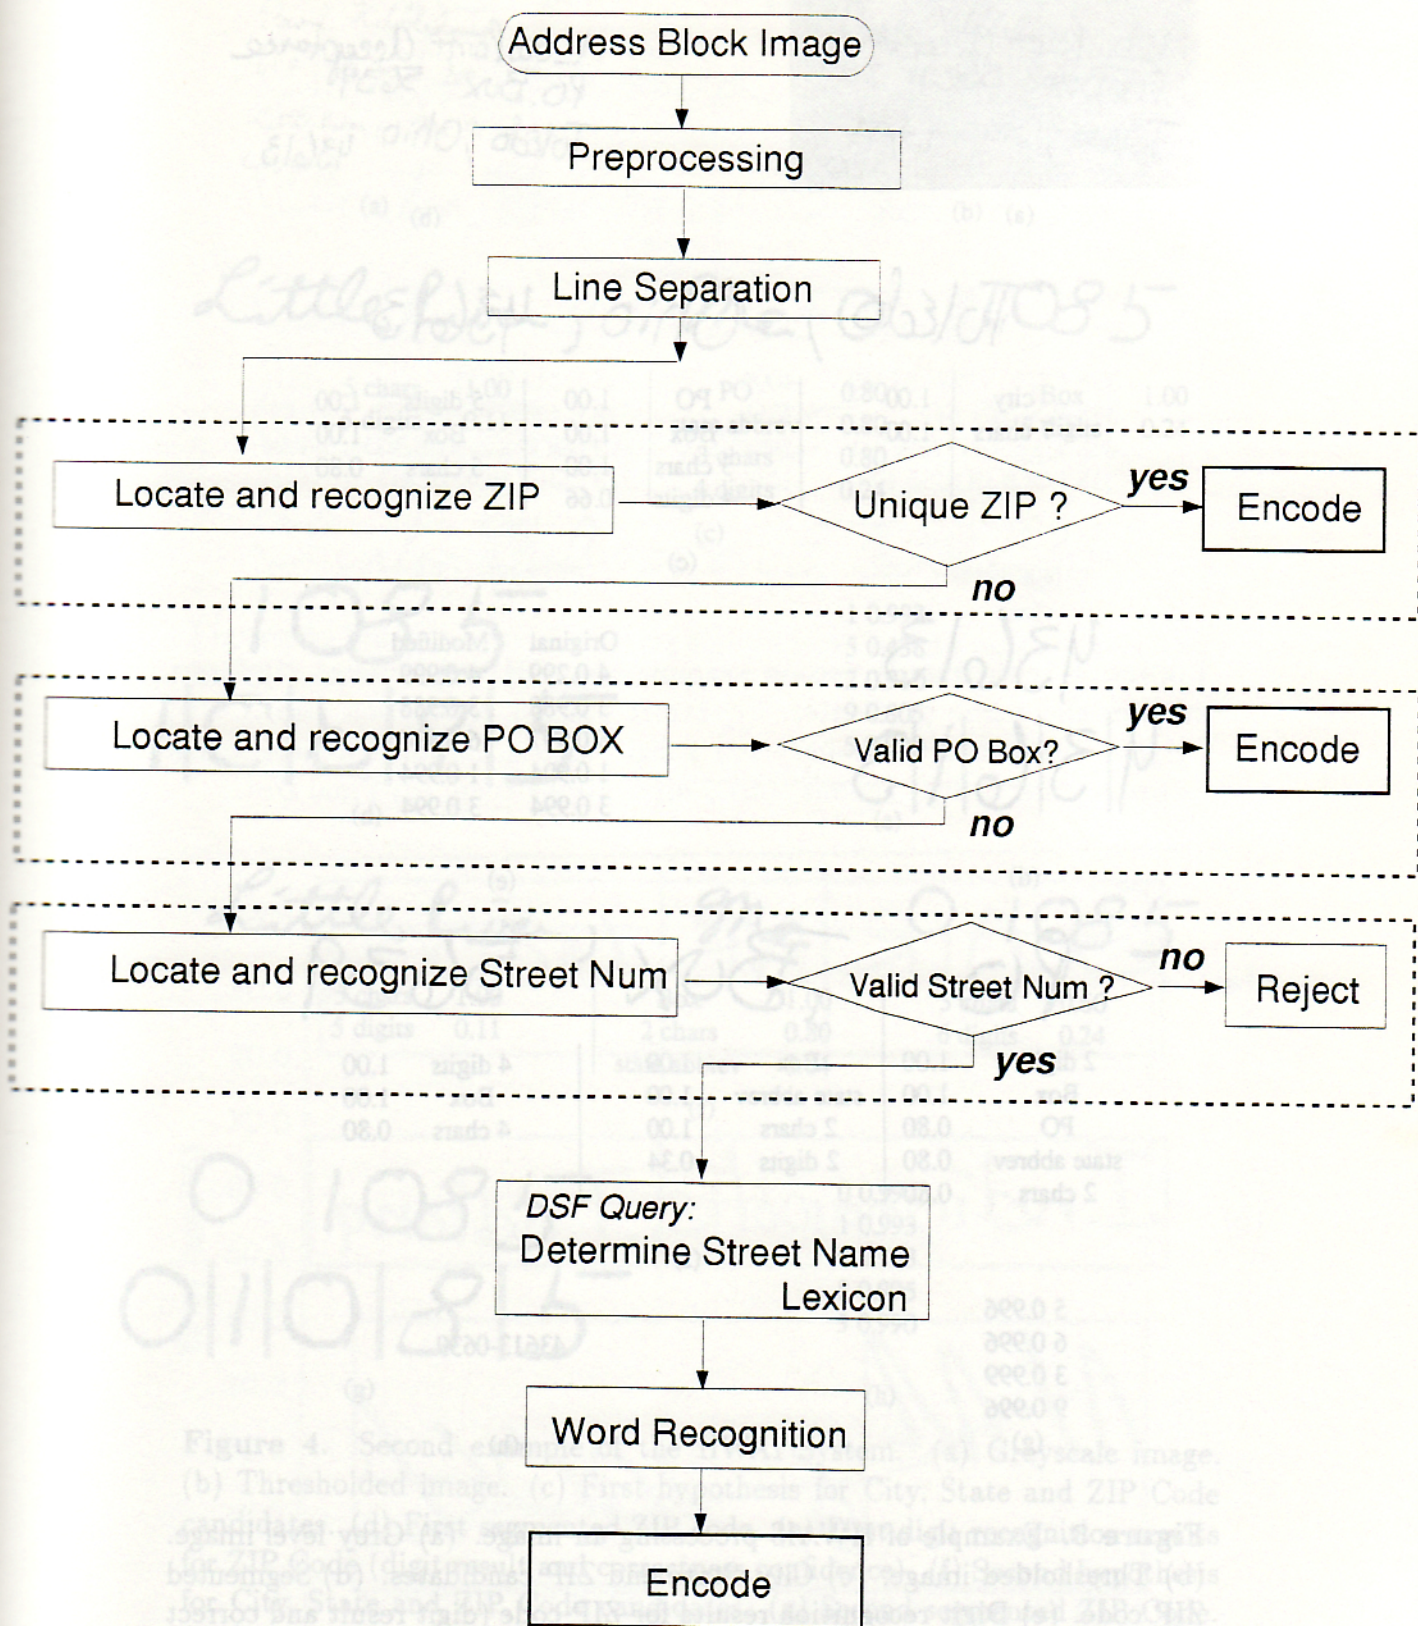
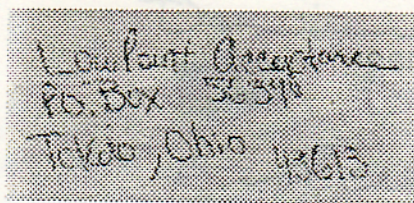
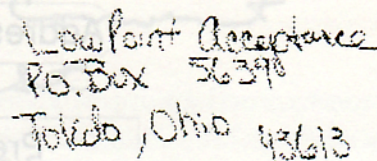


Figure 2. Overview of the Handwritten Address Interpretation System.



(a)



(b)

Toledo, Ohio 43613

city	1.00
7 chars	1.00

PO	1.00
Box	1.00
3 chars	1.00
4 digits	0.66

5 digits	1.00
Box	1.00
5 chars	0.80

(c)

43613

4|3|6|1|3

(d)

Original	Modified
4 0.299	4 0.999
3 0.988	3 0.988
6 0.997	6 0.997
1 0.994	1 0.994
3 0.994	3 0.994

(e)

P.O. BOX 5639

2 digits	1.00
Box	1.00
PO	0.80
state abbrev	0.80
2 chars	0.80

Box	1.00
state abbrev	1.00
2 chars	1.00
2 digits	0.34

4 digits	1.00
Box	1.00
4 chars	0.80

(f)

5 0.996
6 0.996
3 0.999
9 0.996

(g)

43613-0639

(h)

Figure 3. Example of HWAIS processing an image. (a) Grey level image. (b) Thresholded image. (c) City, State and ZIP candidates. (d) Segmented ZIP code. (e) Digit recognition results for ZIP code (digit result and correct confidence) and adjusted digit recognition results. (f) "P.O.", "Box" and Box-number candidates. (g) Digit recognition results of P.O. Box number. (h) 9-digit ZIP+4 code (also the DPC).

David Kibblesome
 318 E Main Str
 Little River, Me 01085
 CASA

(a)

David Kibblesome
 318 E Main Str
 Little River, Me 01085
 CASA

(b)

Little River, Me 01085

5 chars 1.00
 5 digits 0.11

PO 0.80
 state abbrev 0.80
 3 chars 0.80
 4 digits 0.24

Box 1.00
 5 digits 0.24

(c)

1085
 ||C||L||R||5

(d)

1 0.993
 5 0.438
 2 0.910
 9 0.805
 5 0.990

(e)

Little River, Me 01085

5 chars 1.00
 5 digits 0.11

Box 1.00
 2 chars 0.80
 state abbrev 0.64

5 chars 0.60
 6 digits 0.24

(f)

01085
 0||1||0||8||5

(g)

0 0.992
 1 0.993
 0 0.993
 8 0.995
 5 0.990

(h)

Figure 4. Second example of the HWAI System. (a) Greyscale image. (b) Thresholded image. (c) First hypothesis for City, State and ZIP Code candidates. (d) First segmented ZIP code. (e) First digit recognition results for ZIP Code (digit result and correctness confidence). (f) Second hypothesis for City, State and ZIP Code candidates. (g) Second segmented ZIP Code. (h) Second digit recognition results for ZIP Code (digit result and correctness confidence).

318 E Main Str

3 digits 1.00
 Box 1.00
 3 chars 0.80

5 chars 0.67
 8 digits 0.04

(i)

3 0.986
 1 0.998
 8 0.999

(j)

318 CITY VIEW BLVD	-4006	318 E MAIN ST	-3311
318 E MAIN ST	-3311	318 E MOUNTAIN RD	-1853
318 E MOUNTAIN RD	-1853	318 CITY VIEW BLVD	-4006
318 LOCKHOUSE RD	-1212	318 LOCKHOUSE RD	-1212

(k)

(l)

01085-3311

01085-3311-18

(m)

(n)

Figure 4 (continued) Second example of the HWAI System. (i) Street number and name candidates. (j) Street number recognition results. (k) Original street name dictionary and 4-digit addons. (l) Ranked street name dictionary and 4-digit addons. (m) 9-digit ZIP+4 code. (n) Final DPC.

(a) BLAIR NE 68009

(b) Indianapolis, Indiana

(c) NM 87004

Figure 5. Word breaks can be indicated by three types of features: (a) a horizontal gap between words, (b) a punctuation mark present between words, or (c) a vertical shift in the text line.

Clinton Iowa 5 2 736
 Clinton Iowa 5 2 736
 Clinton Iowa 5 2 736
 Clinton Iowa 5 2 736
 Clinton Iowa 5 2 736

Figure 6. Possible ways of dividing a given text line into words. The correct choice is second from the bottom.

Introduction

The United States Postal Service (USPS) is currently involved in an extensive attempt to automate the letter sorting process associated with mail delivery. To date, they have deployed approximately 400 multiline optical character readers (MLOCs) that read machine printed addresses and assign corresponding ZIP+4 codes. Of the letters currently fed to MLOCs, 15 percent are machine printed and practically all are rejected. In addition, the percentage of machine printed addresses with bar codes increases, the percentage of the mail stream sent to the MLOCs that is handwritten also increases. Thus, a small volume must be processed by costly semi-automatic or manual processes. Therefore, it is the objective of the USPS to develop techniques to encode nationally handwritten addresses to delivery stop at near-real-time speed.

The goal of USPS research in the area of handwritten address interpretation is to read 100 addresses per minute with the following performance criteria: delivery-stop error rate of at least 50 percent, a 5-digit error rate of at most 1.5 percent, and a letter-on error rate of at most 2 percent [1].

This work involves extensive experience in machine printed address interpretation having been involved in USPS research programs since 1986. These programs have involved handwritten ZIP code recognition, contextual analysis of machine printed ZIP codes, and advanced research in machine printed character recognition. In particular, the knowledge gained from the end-to-end context sensitive system is highly valuable in designing a similar system.

2.0. This paper discusses the context sensitive system in building an end-to-end system. Section 2 discusses the context sensitive system in building an end-to-end system.

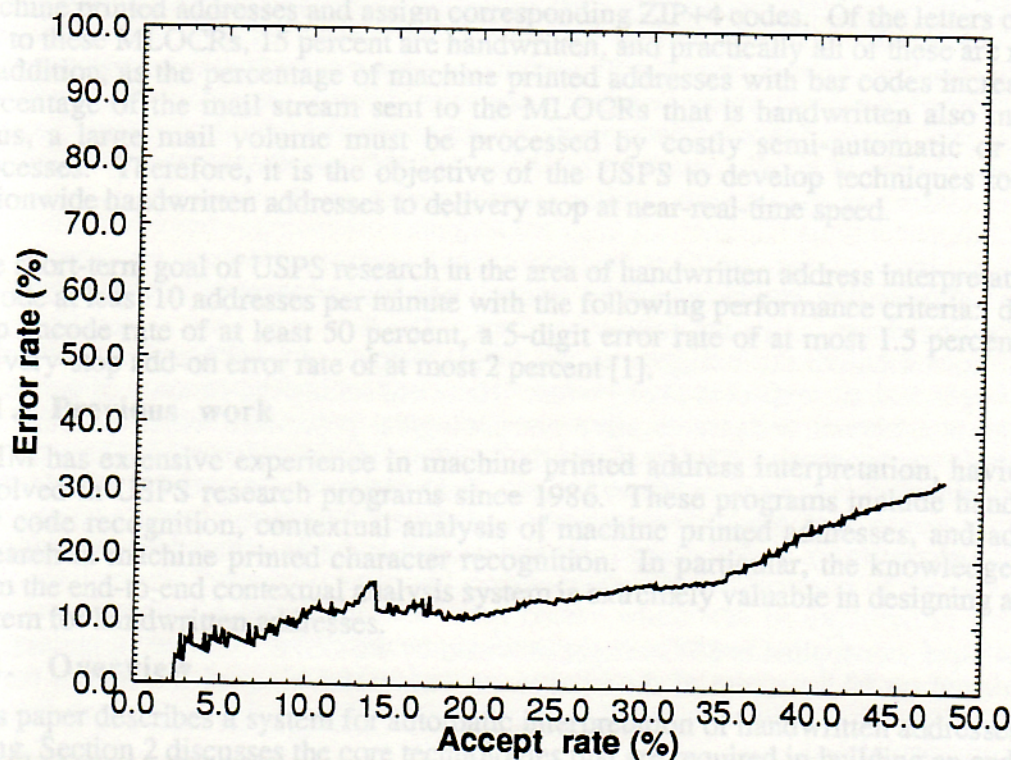


Figure 7. Performance of the Handwritten Address Interpretation System at various levels of threshold.