

From Pixels to Paragraphs: the Use of Contextual Models in Text Recognition

Sargur N. Srihari

Center of Excellence for Document Analysis and Recognition
State University of New York at Buffalo
Buffalo, NY 14228-2576

Abstract

Contextual models are necessary for text recognition when there are document image imperfections, (e.g., multi-generation copies/faxes, poor contrast), or the document is inherently complex, (e.g., unusual fonts, non-standard layout, spelling errors, non-standard abbreviations). A hierarchy of models that characterize alternatives corresponding to characters, words, phrases, sentences, and paragraphs can be defined. A domain where explicit models are available up to the paragraph level is that of postal addresses. At each level of the contextual hierarchy, the models and their use in text recognition algorithms are described.

storage. The same document requires only one-eighth (or 33,750 bytes), if represented as a binary (black and white) image. Assuming that a typical address has only 75 characters, the recognition of the characters yields an ASCII representation of 75 bytes. Since there are about 20 words in an address, if we assume that all words in an address arise from a 250,000 word lexicon, then each word requires 18 bits, or the address requires 45 bytes. The entire address is a paragraph of several sentences, one of which has the syntax: city-state-ZIP Code, another has the syntax: street number street name-suite number, and a third is a phrase representing the firm name. Since there are 120 million de-

1. Introduction

The task of OCR (optical character recognition) or text recognition is to convert scanned and digitized machine-printed text into its intended symbolic form. This is a data reduction process where text in one form of representation is reduced to another more parsimonious form.

Several levels of representation and storage requirements for a text image are shown in Table I. Recognition is the process of converting the text from an image level representation to any of the symbolic levels. The levels correspond to either the signal (image) level or the symbolic (ASCII, linguistic, and textual structure) levels. The linguistic level is subdivided into: word, phrase, and sentence levels. Higher level units corresponding to paragraphs and beyond (sections, chapters, stories, etc.) are document structure levels involving spatial structure rules. The storage requirements for a document decrease as we move up the representation-level hierarchy. The rate of decrease is dependent on document semantics. We consider a document class where the semantics are fairly well defined, viz., typed US addresses (See Figure 1).

A 1" x 3" address label when scanned and digitized at 300 pixels per inch gray-scale requires 270,000 bytes of

Table I: Representation levels and storage requirements for US address image.

Text Rep. Level		Storage Requirement
Textual Structure	Document	
	Paragraph	27 bits (120 million delivery points)
Linguistic	Sentence	88 bits (4 sentences)
	Phrase	139 bits (7 phrases)
	Word	45 bytes (20 Words, 250,000 word lex)
ASCII	Character (8 bits/char)	75 bytes (75 chars/address)
Image	Binary	33,750 bytes
	Gray (300 dpi, 8 bits gray)	270,000 bytes

livery points in the United States, the address can be represented by just 27 bits (or 3.4 bytes). Thus there is a 80,000-to-one reduction in data content between the gray-scale representation and its meaning.

2. Models

A model is essentially knowledge about pattern variations. The models capture knowledge of image degradation processes or they can be about intrinsic variations in the text itself.

Models and recognition algorithms pertaining to different levels of representation are shown in Figure 2. Models at each linguistic level can be represented in terms of models at the next lower level, as well as in terms of spatial models: pixels (bit maps) and features. A bit map of a prototype character in a given font is an image level model. A stroke level representation of a given character is a feature level model. A word model is a representation of a word in terms of its characters or its characters' features or images. The possible words are modeled by explicit lexicons, n-grams, or statistics of letter combinations.

A phrase is a group of words that form a meaningful unit. In addresses, a multi-word street name is a phrase. A sentence is several phrases that complete a concept. A street address consisting of the house number and a street name is a sentence.

A paragraph in the English language is a cohesive unit of text. Concepts and relationships expressed in a paragraph are relevant to each other thus enabling us to make plausible inferences about the underlying meaning. An address is a paragraph. One sentence is the name associated with an individual. The name can be a first name, middle name, and initial. Another sentence identifies the firm — which is a business or organization. A third sentence gives the street address, which consists of a street number and street name. The underlying meaning of an address is the identification of an individual and the mail drop-off.

3. Character level

A character model is necessarily in terms of image level primitives. The models (implicit or explicit) correspond to

each character, e.g., 77 models for upper/lower case characters (52), numerals (10), and special symbols (15) are necessary for postal addresses. The models are either: (i) bit maps (templates) for each character in each font (an image level model), (ii) or character features that are font-independent (a feature level model). The bit maps or feature models can be explicitly stored or distributed as weights in a neural network. Recognizers based on bit maps (template matching) are not as robust as those based on feature sets, but they are fast and perform well on standard fonts. Some useful features are pixel pairs, contour features, stroke features, end points, etc.

3.1 Pixel models, OCR systems

Commercial OCR systems typically employ template matching. They can reliably recognize cleanly typed machine-print. The areas where performance needs improvement are in handling addresses with smear, broken characters, touching characters, poor contrast, textured backgrounds, etc.

The performance of three OCR software packages for the Macintosh computer (Calera Wordscan Plus 1.0, Caere Omnipage Professional 3.0, and Xerox Imaging Systems AccuText 3.0) were compared on the address image of Figure 1, which was scanned on a desktop scanner at 300 ppi binary.

The results with the address image for each of the OCR software is shown in Figure 3 (a-c). The result of Figure 3 (d) (combination Pb and H) refers to research software. On this image, research system performance at the character, word, phrase, sentence, and paragraph level are 98%, 80%, 71%, 50%, and 0% respectively.

3.2 Feature models

Recent surveys of feature-based methods are [1] and [2]. The use of gray-scale imagery in recognition, use of *cognitively justified features such as those derived from Gabor filters and artificial network classifiers* provide a new impetus for research in this area. Four methods, Ga, Gr, H, and P, were explored from the viewpoint of recognizing segmented gray-scale data.



Line 4: [faded]
Line 3: [faded] INC.
Line 2: 2400 MEIDINGER TOWER
Line 1: LOUISVILLE, KY 40202

Figure 1. Image of US Address used as a paragraph example throughout this paper.

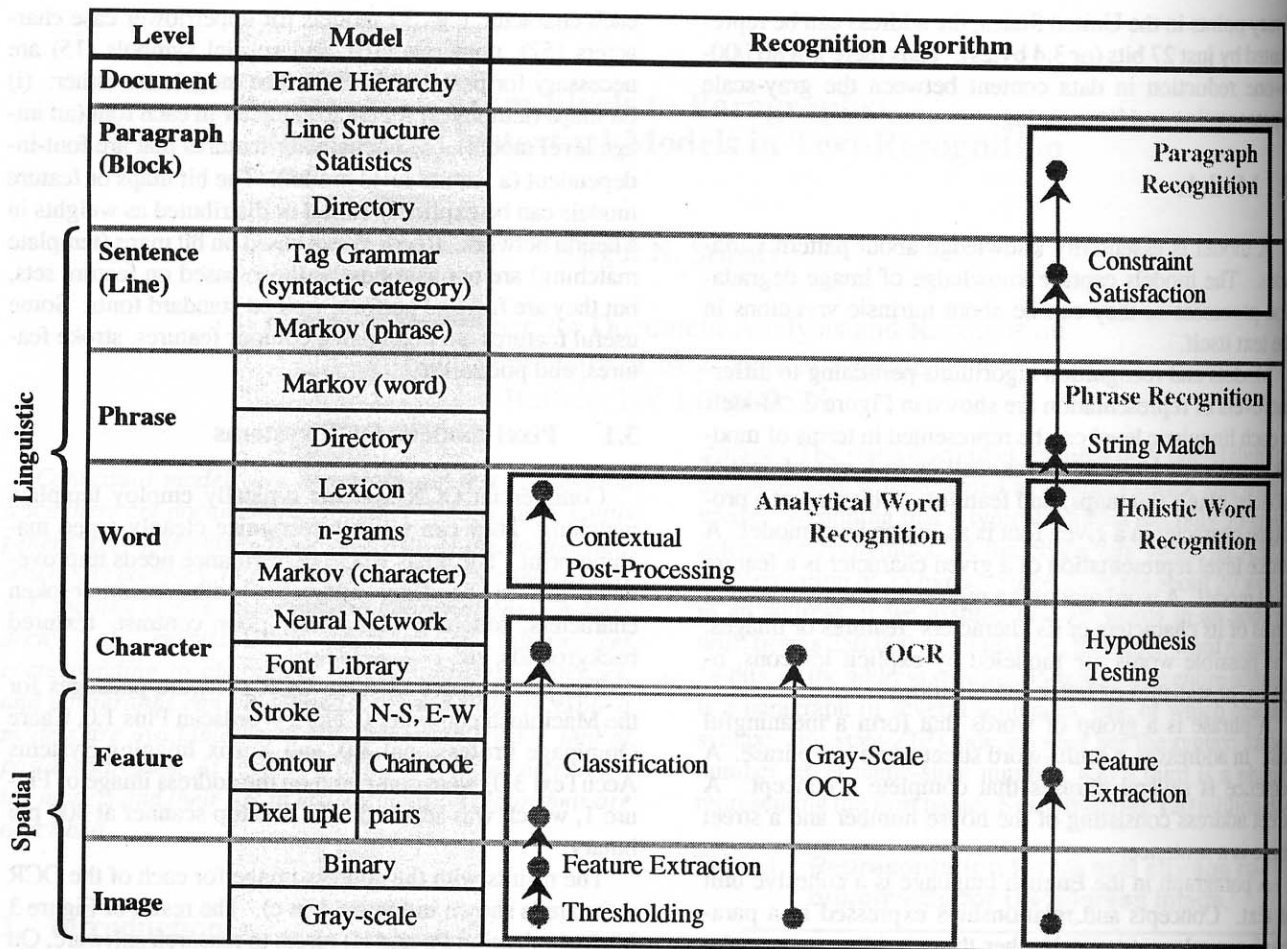


Figure 2. Hierarchy of levels, models, and algorithms in text recognition.

Recognizer P, based on polynomial discriminant functions [3], uses products of certain pixel pairs as features and a linear classifier (Table II). Recognizer Ga, based on Gabor functions, maximizes the joint resolution in the spatial frequency information domain. Recognizer Gr uses gradient magnitude and direction as features. Recognizer H uses features derived from vertical and horizontal histograms.

The four methods H, P, Ga, and Gr, were combined using logistic regression — which is equivalent to a neural

network with no hidden units and using a weight constraint. Results with folding similar characters together into equivalence classes are also shown.

4. Word level

Word models can be represented in terms of characters in three different ways: list of legal words (or lexicon), n-grams (legal letter combinations), and Markov models

Table II: Character recognition performance with no rejects

Training: 47,565 alphanumerics from postal images, Testing: 25,000 gray-scale, 300 ppi

Recognition Method	Features (No., Type)	Classifier	Alphanumeric Recognition
Pb (binary)	Pixel pairs (989, real)	Linear	89.74%
Pg (gray-scale)	Pixel pairs (989, real)	Linear	91.26%
Ga	Gabor Coefft (72, real)	Neural Network (40 hu, 100 hu)	94.06%
Gr	Gradient (192, binary)	Neural Network (101 hu, 127 hu)	92.07%
H	Histogram (72, real)	Neural Network	93.2%
Combination (Pg, Ga, Gr, H)	Decisions, confidences (186, real)	Logistic	95.12%
Combination (fold)	Decisions, confidences (186, real)	Logistic	98.76%

(a) AccuText	(b) OmniPage
JAMES TENNILL ADVEST j~j~• 2400 MEIDENGER TOWER LGUISVILLE, KY 40202	JAMES TENNILL ADVEST ~C> 2400 ME~INGER T~WER LGUISV~LLE, KY 402~2
(c) WordScan	(d) Combination Pb and H
JAMES TENNILL ADVEST INC. 2400 MEIDINGER TOWER LOUISVILLE, KY 40202	JAMES TENNILL ADVEST INC 2400 MEIOLNGER TOWER LCUISVILLE KY 40202

Figure 3. Performance of commercial OCR software and research software.

(first and higher order letter transitional probabilities). Three distinct approaches to utilizing word models in character recognition can be identified: character-based word recognition, segmentation-based word recognition and word shape recognition.

4.1 Analytic Model: Character-based

The word is modeled as consisting of characters and each character has a distinct shape. Recognition is a three-step approach. First the word image is segmented into character images. Second, the segmented characters are recognized by using a character recognition technique. Third, the resulting word is corrected by a model of the expected words. The third step is contextual post-processing (CPP). CPP techniques differ in the method of representing word models. Those that utilize lexicons use string matching techniques to measure a distance between the recognized word and each word in the lexicon.

The analytic method is appropriate when a reliable segmentation can be obtained and the extracted characters are not deformed by size normalization and other procedures that prepare a character for recognition. Certain ambiguities in the decisions on character classes can be resolved by CPP (Figure 4). After the recognition stage, the shape information has all been converted into class labels. The

class decision suggests only that the character is in a neighborhood of a standard shape of the decided class. Subtle information about how much the character deviates from the standard shape is forever lost. In some cases, this loss could be so severe that the true word cannot be recovered.

4.2 Holistic Model: Segmentation-based

Each word in a lexicon is modeled as a shape where each character shape is assumed to be segmentable. Recognition involves a two-step approach. In the first step, the word image is segmented into individual characters. In the second, recognition, step, features are extracted for each character image. Classification into a word is done by examining the entire (compound) feature set.

Word level knowledge is brought into play during the recognition step. Letter transitional probabilities and class-conditional probabilities of character feature vectors can be simultaneously utilized during recognition. Lexical information can also be brought to bear on the recognition phase (Dictionary-Viterbi algorithm).

The method is suitable for word images where characters can be reliably segmented but difficult to recognize in isolation, e.g., images so broken that most shape features are lost.

Line 4:	JAMES L 5	TENNILL L 8	
Line 3:	ADVEST L 6, K ADVERT 1.0	INC L 3, K INC 0.0	
Line 2:	2400 N 4	MEIOLNGER L 9	TOWER L 5, K TOWER 0.0
Line 1:	LCUISVILLE L 10, K LOUISVILLE 1.	KY L 2, K KY 0.0	40202 N 5

Figure 4. Word recognition using analytic approach and CPP (L, N, and K represent alphabetic, numeric, and keyword fields with confidences).

Table III: Word recognition performance

Test set: 1340 300 ppi grayscale, Lexicon size: 2067
Font types: 66

Recognition method	Top 1 choice	Top 5 choices	Speed (secs/word)
Analytical	95.8%	97.6%	10.07
Holistic (seg)	85.4%	95.0%	23.64
Holistic (WS)	89.1%	93.1%	0.9

4.3 Holistic Model: Word-shape

The word is modeled as a shape where boundaries between character shapes is unimportant. Recognition is a one-step holistic approach that bypasses segmenting a word image into characters. Features are extracted from the unsegmented word image. The recognition problem is to associate the given word image with one or more words in a lexicon (hypothesis testing). Each word in the lexicon is represented by a feature model based on an expected set of fonts.

The advantage of word shape recognition is that some errors in character segmentation and premature decisions on character identities are avoided. It is suited for images that are difficult to segment into characters, or where characters are distorted when normalized.

4.4 Recognizer comparison and combination

Performance and speed of the different methods are shown in Table III. Improved overall performance can be obtained by combining results. In the logistic regression method of combination, weights are associated with classifiers to take into account the performance statistics of each individual classifier and correlations of classifiers [5].

5. Phrase level

When the size of the lexicon is 250,000 and each phrase is of average length 2, then there are 62.5 billion possible phrases. Of these, only a fraction are legal and it is these that need to be modeled. As with words, phrases/sentences can be modeled at the word level (analytic), with interword constraints, or modeled at the lower character level (holistic). The latter is also referred to as line-based as it involves concatenating the words and matching against legal records. We consider each approach.

When the number of phrases is not too large it is possible to deal with phrases as units. By considering the problem of identifying personal and firm names in postal addresses (lines 3 and 4) we were able to implement a whole-phrase approach and compare it with a word-based approach. The records retrieved from a database with each approach are shown in Table IV. The performance summaries for the approaches are given in Table V [8].

In the word-based approach, abbreviations can be recognized intelligently, e.g., First and 1st. The insertion and deletion of keywords is not critical. The system is more tolerant of parsing errors since only one correct firm is needed to index the database records. The disadvantages are that relative word positions are not taken advantage of. They are adversely affected by the presence of unimportant words, e.g., Company.

In the line-based approach, word segmentation errors do not have any effect on performance. Also, the relative word positions are taken advantage of. Its disadvantage is that performance depends heavily on the database containing all firm name variations to overcome: abbreviations, acronyms, aliases and keyword insertions/deletions.

Table IV: Result of matching recipient phrases against directory of recipient phrases.

Recipient-line matching (holistic):						
	Name	Street	City	State	ZIP+4 Code	Confidence
1	ADVEST INC	2400 MEIDINGER TOWER	LOUISVILLE	KY	402023446	0.888889
2	ADVEST INC	2400 MEIDINGER TOWER	LOUISVILLE	KY	402023463	0.888889
3	JAMES TENNILL	2400 MEIDINGER TOWER	LOUISVILLE	KY	402023446	0.888889
Recipient-word matching (analytic):						
1	ADVEST INC AC	2400 MEIDINGER TOWER	LOUISVILLE	KY	402023446	1.000000
2	ADVEST INC	2400 MEIDINGER TOWER	LOUISVILLE	KY	402023446	1.000000
3	ADVEST INC	2400 MEIDINGER TOWER	LOUISVILLE	KY	402023463	1.000000
4	ADVEST JERRY DEVERIS	2400 MEIDINGER TOWER	LOUISVILLE	KY	402023446	1.000000
5	ADVEST	2400 MEIDINGER TOWER	LOUISVILLE	KY	402023463	1.000000
6	ADVEST	MEIDINGER TOWER	LOUISVILLE	KY	402023463	1.000000
7	JAMES A. GRAY	2500 MEIDINGER TOWER	LOUISVILLE	KY	402023446	0.976000
8	JAMES TENNILL	2400 MEIDINGER TOWER	LOUISVILLE	KY	402023446	1.000000
9	MR JAMES D TENNILL JR	2400 MEIDINGER TOWER	LOUISVILLE	KY	402023446	1.000000

Table V: Phrase recognition performance

Database size: 138,521 records;
 Test set: 2,000 images (using character rec)

Recog method	Encoded	Correct in Nbrhd	Errors	Speed (s/phrase)
Analytic	79.4%	87.2%	12.8%	4
Holistic	78.4%	88.2%	11.8%	14

6. Sentence level

Sentence level models are used to determine the correct syntactic category of each word/phrase in a sentence. Once the category is chosen, the appropriate lexicon can be chosen for recognition. Models can be purely syntactic, purely statistical or hybrid.

6.1 Syntactic parsing

A grammar captures the syntax of all legal input sentences and rejects others. The syntactic categories are as follows:

· cw = city word, fw = firm word, nn = recipient name, ow = other word, sn = street number, ss = street suffix, st = street word, sw = state word, and z5 = five digit ZIP Code.

A rule-based parser works under the premise that certain word types are likely to precede/follow one another, e.g., a street number is usually followed by either a pre-directional or a street name. Prior to parsing, keyword matching is done first. This is not always able to assign word types to all words, nor can it resolve conflicts of word types (i.e., NORTH can be a city word, street prefix, street suffix or a firm word). There are a set of approximately 100 rules of the form $s_1, s_2 \Rightarrow s_1, s_3$ which will assign and resolve word types. An example is:

street_number, street_prefix? \Rightarrow street_number, street_prefix

This rule states that a street number followed by a possible street prefix, yields a street number followed by a "known" street prefix. This would help resolve the problem above with the word NORTH. The rule engine applies relevant rules to these cases until a line has been parsed (Table VI). The problems with such a method are (i) the inability of the method to cover all possibilities, and (ii) the computational complexity of parsing.

Table VI: Result of rule-based parsing of address

Line 4:	unknown	0.000	ow	ow	
Line 3:	recipient	1.000	fw	fw	
Line 2:	street	0.667	sn	st	st
Line 1:	city-state-zip	0.833	cw	sw	z5

6.2 Statistical parsing

Statistical methods employ n-gram statistics reflecting the transition frequencies between syntactic categories represented by n-tuples of words, i.e., use hidden Markov methodology at the word level. Here the transitional probabilities between syntactic categories and the word-class conditional probabilities associated with the word images would be used.

Prior to parsing, keyword matching determines possible word types. The nearest keywords identified a lexicon and the penalties associated are shown in Table VII. Next, is to assign to each line a set of line configurations from the collected training set data, with some of the following constraints. The first two are number of lines and number of words on a line. The parser limits its search space to only line configurations found in similarly sized address blocks, and obviously to configurations with the same number of words. A third criterion is based on results from keyword matching. If for example the words "RURAL" and "BOX" were matched on a line, then configurations with word types of rural_word and rural_box are given increased confidence. The final constraint used is training set statistics which contain information as to the number of times that this configuration has previously been seen. The problem with this approach is that local maxima often cause the correct sentence choice to be overlooked.

7. Paragraph level

Recognition at the paragraph level implies association with a semantic model. The number of such models can be very large in any domain. In the case of an address paragraph, the number of semantic models can be enumerated. For instance there are 120 million delivery points in the United States which are represented by a nine digit ZIP+4 code together with a two-digit add-on.

7.1 Syntax: vertical parsing

A paragraph model at the syntactic level needs to capture the dependencies between sentences. In the case of an address, there are dependencies between address lines. For example, most often the bottom line (line 1) is the City-State-ZIP line, line 2 is a street address line, line 3 a personal name, etc. We refer to the assignment of line (sen-

Table VII: Result of statistical parsing of address

Line 4:	recipient	0.046	nn	nn	
Line 3:	recipient	0.548	fw	fw	
Line 2:	street	0.385	sn	st	ss
Line 1:	city-state-zip	0.749	cw	sw	z5

tence) type as vertical parsing and the assignment of phrase types within a line as horizontal parsing. Address parsing can be performed with a statistical or a rule-based approach.

Statistical parsing: Various combinations of horizontal line configurations are compared with the training set data. The combinations that are statistically more probable are given greater confidence, in the set of multiple parses returned. The advantage of the statistical parser is that it can sometimes parse address blocks even when poor image quality results in poor character recognition. The results of statistical parsing together with the confidences associated with the parses for the address of Figure 1 are shown in Table VII. The overall confidence of the parse was 0.432

Rule-based parsing: A state vector is used to signify what line types previously seen. For example, if a city-state-zip line has already been parsed, then lines above it are considered less likely to be of this type. This also aids in horizontal parsing, because rules for city, state and zip can usually then be ignored. Some of the advantages that the rule based parser has over the statistical parser are the ability to parse address blocks of a form it has not seen before, and the ease of adding rules without retraining the entire system. The results of rule-based parsing on the image of Figure 1 are given in Table VI. The overall confidence of the parse was 0.625.

7.2 Semantics: constraint satisfaction

Once the parse is completed, the task is to associate the structure with a semantic model. In the case of postal addresses, the different choices supplied by the parser can be used to look up a directory. Based on the city-state-zip sentence, the ordered choice of ZIP Codes is 40202, 40220, 40221, and 40222. The street sentence allows us to determine the ZIP+4 to be 40202-3446. This determination can be done efficiently using constraint satisfaction [7].

A constraint satisfaction problem consists of n variables, $X_1 \dots X_n$, each having an associated domain of possible values. Each variable can be assigned any value in its domain but *constraints* are imposed on the legal combinations of values that two or more variables can assume. A solution consists of a set of values, one for each variable, such that the constraints over all variables are satisfied.

The constraint satisfaction paradigm maps very easily to postal image interpretation. Each word in lines 1 and 2 of the address block image is considered to be a variable. The domains for each variable are provided by the recognition routines. Finally the constraints are imposed by the ZIP+4 database. The constraints between two or more pieces of

information from a mail piece are satisfied if there is at least one record in the ZIP+4 database that contains all the specified information. This process is carried out by a query system built around the ZIP+4 database which allows access to the data in a fast and efficient method. Figure 5 illustrates a graph of the various fields in the ZIP+4 database that can be used in queries. Every field is linked to every other field. Thus, every piece of information provides a constraint on every other piece of information. For example, we may be interested in the constraints between the ZIP Code ($z5$) and the first street word ($st1$). A query in this case might take the form of the ZIP Code being "75248" and the street word being "HIDDEN." If this were a legal combination of values, the query would return "yes." That is, these two values would satisfy the constraints imposed by the ZIP+4 database.

Conflict analysis: This section is closely tied to the constraint satisfaction solver in that it uses information about which variables have been repaired in selecting a variable to be removed or modified, also known as the culprit. This section uses a heuristic that says the most likely culprit is the variable that was repaired first. This will be the variable that, using the top recognition choices, had the most conflicts. Depending on the category of the selected culprit, various methods may be tried to improve or modify the recognition results. An example of this is when the culprit is a PO box number. Sometimes this is written as "#248." The "#" is often misrecognized as a "4" due to low resolution and poor image quality. Thus a recognition choice for this number could be "4248." A simple modification to the recognition results in this case would be to drop the choices for the first digit from the recognition results. The constraint satisfaction solver will be called again, but this time "248" will be used instead of "4248." If nothing can be done to modify the recognition results of the culprit, then the culprit will be removed from further processing.

In some cases information should be removed not because of misrecognition but because it is written incorrectly. For example, if a person addressing an envelope accidentally writes the wrong ZIP Code, conflict analysis would want to identify the ZIP Code as being incorrect and remove it from further consideration. In this way our system is able to overcome errors on the address block itself, as well as correct the occasional recognition error.

Reanalysis: The reanalysis section is also closely tied to the constraint satisfaction solver. It uses the solution found by informed backtracking to retrieve all ZIP+4 records that match the address believed to be on the mail piece. If only one record is found, then the ZIP+4 Code for that record is assigned to that mail piece. Otherwise, reanalysis looks for

differences between the records. Unused information from the address block is used to distinguish between two or more records. This process continues until a single record is isolated or there is no more usable information on the address block.

The system was tested on 2000 address block images which had not been previously seen collected from real mail on a national scale. In the test, the system was able to encode 92% to nine-digits, with errors in about 9% of the nine-digit assignments.

7.3 Semantics: recipient matching

The recipient sentences (lines 3 and 4) can be used to arrive at an independent assignment of the ZIP+4 Code. The results of matching the recipient on a whole line basis and a word basis yields three and nine matching records in a firm directory as shown in Table IV. The firm directory consisted of records determined from a large learning set of actual addresses.

The results from lines 3 and 4 of the address paragraph provide independent confirmation of the interpretation provided by lines 1 and 2 using constraint satisfaction. When the two results are combined together, the encode error rate can be significantly reduced. For the 2,000 images described in Table V, the encode rate was 70% with a 0.98% error rate.

8. Summary and conclusion

We have discussed several hierarchical levels in modeling text. At each level, the model is typically based on elements at the next lower level, e.g., characters in terms of features, words in terms of characters, phrases/sentences in terms of words, and paragraphs in terms of sentences. Such models are referred to as analytic, since they are based on




Figure 5. Constraints from the ZIP+4 File

step-by-step analysis. Although such models are usually sufficient, in order to achieve system robustness it is necessary to use models that skip a level, e.g., words in terms of features/pixels and sentences in terms of characters. Such models are holistic in that decisions about entities are made in the context of decisions about other entities. Holistic methods are slower than analytic methods and do not necessarily perform better. However, when analytic and holistic methods are combined, overall system performance is higher than either alone. Recognition algorithms based on different models can be combined to achieve robustness.

Acknowledgments

I would like to thank Joseph Piazza who did the commercial OCR experiments. Thanks are also due to the contextual analysis group at CEDAR, notably Peter Cullen, Ron Curtis, Anthony Girardin, Jon Hull, Doug McKee, and Xin Zhao, who performed many of the experiments described.

References

- [1] S.N. Srihari and J.J. Hull, Character Recognition, *Encyclopedia of Artificial Intelligence*, Second Edition, S.C. Shapiro (editor), Wiley Interscience, New York, 1992, 138-150.
- [2] R. Fenrich, S. Lam and S.N. Srihari, Optical Character Readers, *Encyclopedia of Computer Science and Engineering*, Third Edition, A. Ralston (editor), Van Nostrand Reinhold, 1992, 993-1000.
- [3] J. Schurmann, A multifont word recognition system for postal address reading, *IEEE Transactions on Computers*, C-27(8), April 1978, 721-732.
- [4] S.N. Srihari, Recognition of handwritten and machine-printed text for postal address interpretation, *Pattern Recognition Letters*, 14, April 1993, 291-302.
- [5] T.K. Ho, J.J. Hull and S.N. Srihari, A computational model for recognition of multifont word images, *Machine Vision and Applications*, 1992, 5: 157-168.
- [6] P.W. Palumbo and S. N. Srihari, Text parsing using spatial information for recognizing addresses in mail pieces, *Proceedings Eighth International Conference on Pattern Recognition*, Paris, France, 1986, 1068-1070.
- [7] P.B. Cullen, J.J. Hull, and S.N. Srihari, A constraint satisfaction approach to the resolution of uncertainty in image interpretation, *Proceedings Eighth IEEE Conference on Artificial Intelligence for Applications*, Monterey, CA, 1992.
- [8] P.K. Wong, T. K. Ho, and S. N. Srihari, Firm name recognition for postal address interpretation, *Proceedings Fifth USPS Advanced Technology Conference*, Nov 30-Dec 2, 1992, 757-770.