

- Gold B & Morgan N (1999). *Speech and audio signal processing*. Chichester/New York: Wiley.
- Hermansky H & Morgan N (1990). 'RASTA processing of speech.' *IEEE Transactions on Speech and Audio Processing* 2(4), 578–589.
- Holmes J N & Holmes W J (2001). *Speech synthesis and recognition* (2nd edn.). London: Taylor & Francis.
- Huang X, Acero A & Hon H-W (2001). *Spoken language processing: a guide to theory, algorithm and system development*. Upper Saddle River, NJ: Prentice Hall.
- Jelinek F (1998). *Statistical methods for speech recognition*. Cambridge, MA: MIT Press.
- Jurafsky D & Martin J (2000). *Speech and language processing*. Upper Saddle River, NJ: Prentice Hall.

Handwriting Recognition, Automatic

S Srihari, State University of New York at Buffalo, Amherst, NY, USA

© 2006 Elsevier Ltd. All rights reserved.

Handwriting is a means of recording personal information as an aid to memory, as well as a means of communication with others. A pen together with paper or a small note-pad is often more convenient than a computer keyboard. Automatic recognition of handwriting has significance in reading handwritten notes in a personal digital assistant, postal addresses on envelopes, amounts in bank checks, forms, field notes, historical records, etc. The ultimate handwriting computer will have to process handwriting in an unconstrained environment, deal with many writing styles and languages, work with arbitrary user-defined alphabets and understand any handwritten message by any writer.

Fundamental characteristics of handwriting are that it consists of artificial graphical marks on a surface; its purpose is to communicate something; and this purpose is achieved by virtue of the mark's conventional relation to language. Each script has a set of icons, which are known as characters or letters, that have certain basic shapes. There are rules for combining letters to represent shapes of higher-level linguistic units. For example, there are rules for combining the shapes of individual letters so as to form cursively written words in the Latin alphabet.

Several types of analysis, recognition and interpretation can be associated with handwriting. Handwriting recognition is to transform written language in its spatial form of graphical marks into symbolic representation. For languages based on the Latin alphabet, the basic units of this symbolic representation are characters. Due to nonseparability of cursively written characters, recognition of words may be performed holistically; however, the resulting word is represented in terms of its characters. Handwriting interpretation is to determine the meaning or implication of a body of handwriting, e.g., determining the destination bar code of a handwritten address.

Handwriting recognition and interpretation are processes whose objectives are to filter out the variations so as to determine the message—which are aided by knowledge of the subject domain, e.g., in the case of the physician's prescription, where a pharmacist uses knowledge of drugs. Handwriting identification is to determine the writer of a handwritten document from among a set of known writers. Signature verification is to determine whether or not a questioned signature matches the genuine signature samples of a given person. In contrast to recognition and interpretation, identification and verification are processes that determine the special nature of the writing of a specific writer. All of the above processes depend upon features extracted from the writing, some of which are illustrated in [Figure 1](#).

Handwriting Input

Handwriting data is converted to digital form either by scanning the writing on paper or by writing with a special pen on an electronic surface such as a digitizer combined with a liquid crystal display. The two approaches are distinguished as static or off-line and dynamic or on-line handwriting, respectively. In the dynamic case, the two-dimensional coordinates of successive points of the writing as a function of time are stored in order, i.e., the order of strokes made by the writer is readily available. In the static case, only the completed writing is available as an image. The dynamic case deals with a spatiotemporal representation of the input, whereas the static case involves analysis of the spatioluminance of an image. The raw data storage requirements are widely different. The data requirements for an average cursively written word are: in the dynamic case, a few hundred bytes, typically sampled at 100 samples per second, and in the static case a few hundred kilobytes, typically sampled at 300 dots per inch. Paper documents, which are an inherently analog medium, can be converted into digital form by a process of scanning and digitization. This process yields a digital image. For instance, a typical 8.5 × 11-in page is scanned at a

resolution of 300 dots per inch to create a gray-scale image of 8.4 Mb.

Recognition rates are reported for either character recognition given a fixed alphabet or word recognition given a fixed set of words (or lexicon). Recognition rates are much higher for the dynamic case, e.g., static word recognition rates with lexicon sizes of 10 100 and 1000 were 95%, 85%, and 78%, respectively, for the top choice, whereas the dynamic word recognition rate for a lexicon of size 21 000 was 80%.

Dynamic recognition systems are available in personal digital assistants (PDAs) and tablet PCs. The performance of PDAs is acceptable for processing hand-printed symbols, and, when combined with keyboard entry, a powerful method for data entry has been created. Static recognition systems have made a significant economic impact for constrained domains such as postal addresses and courtesy amounts on bank checks.

Dynamic Handwriting Recognition

Dynamic recognition deals with the automatic processing of a message as it is written, using a digitizer or a stylus that captures information about the pen tip, generally its position, velocity or acceleration as a function of time. The sampled data is preprocessed

for noise reduction, trace normalization and segmentation into units. Noise arises due to digitizer quantization, the digitizing process, erratic hand or finger movements, inaccuracies of the pen-up/pen-down indicator, etc. Noise reduction deals with data smoothing, signal filtering, and breaking connections. Normalization accounts for baseline drift, writing slant and script size.

Segmentation is at the message and word levels. Message segmentation processes are line detection, word segmentation, separating nontextual gesture commands, etc. Words are segmented into characters and sub-character strokes. Tentative segmentation is corrected during classification. Segmentation can be avoided by holistic recognition, particularly for small vocabularies. Some methods combine holistic recognizers with segmentation-based algorithms.

A pen-based computer needs to process a handwritten message as it is produced. The steps, ranging from various shape classification processes to ultimate shape recognition, will have to cope with the variability of message production. Some sources of variability are geometric, allographic and neurobiomechanical. Geometric variations refer to position, size, baseline orientation and slant. Allographic variations deal with the various models that are associated with a single character by different populations of writers. Neurophysiological and

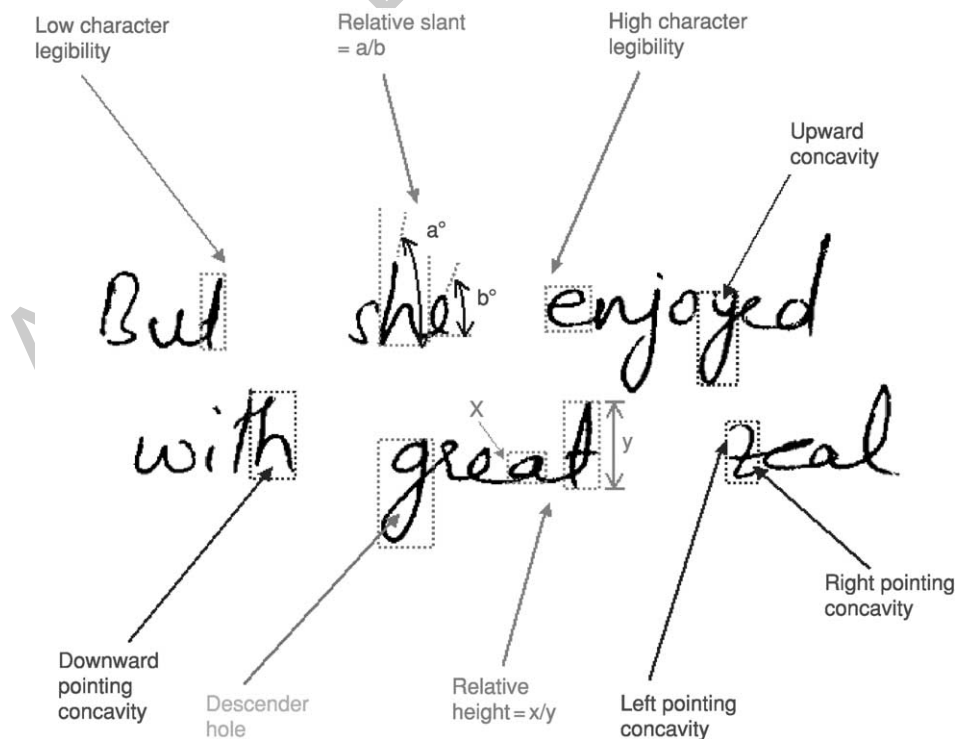


Figure 1 Examples of handwriting features.

biomechanical factors affect both the execution of an action plan or the production of individual strokes.

Recognition methods are either structural or statistical. Structural methods use rules to describe character shape. The rules are embellished by probabilities to capture variations of characters and words. Statistical methods represent shape by a feature list and classes by probability distributions in multidimensional feature space. Statistical methods can be grouped into generative and discriminative.

Generative methods are Bayesian methods, principal component analysis, cluster analysis, etc. Discriminative methods are exemplified by linear discriminant functions and support-vector machines. Artificial neural networks are implicitly generative. For instance, Kohonen self-organized feature maps are used to detect shape prototypes in a training set of characters – which is analogous to *k*-means or hierarchical clustering.

The hidden Markov model (HMM) is a doubly stochastic process that can be based on either discrete or continuous symbols. The former converts the feature vector into a discrete symbol using vector quantization. Symbol probabilities are for stroke shapes in a sliding window. The continuous approach uses variances and covariances of features to estimate the probability of the occurrence of an observed feature vector, which is usually assumed to be Gaussian. The goal of the HMM algorithm is to find the probability that a specific class is the most likely to occur given a sequence of observations. The essence of this approach is to determine the *a posteriori* probability for a class, given an observed sequence where the jump from one state to another is described by a Markov process. HMMs can be incorporated into a stochastic language model.

Cursive script recognizers do not yet take into account contextual anticipatory phenomena: for example, once handwriting is well learned, the neuromuscular effectors involved in that task normally act concurrently to speed up the execution. This generally leads to co-articulation and context effects. The sequence of strokes is not produced in a purely serial manner, i.e., one after the other, but parallel articulatory activity does occur, and there is important overlap between successive strokes or graphemes. The production of an allograph is thus affected by the surrounding allographs: it depends both on the preceding and following units. Many methods take into account the effect of the previous stroke over the actual stroke being processed but often neglect the simultaneous effect of the forthcoming stroke.

Personal adaptation can improve accuracy and usability of dynamic systems. A basic user-dependent system comes with a set of recognizable allographs

for each character, but it allows the user's definition of his own set of symbols or gestures to take user preferences into account. This allows taking into account cultural determinants, handwriting learning systems as well as personal styles and evolution of handwriting habits over a long period of time.

Signature Verification

In signature verification, a test, or questioned, signature is compared with one or a few reference specimens collected from an enrolled user. It requires the extraction of writer-specific information from the signature signal, irrespective of its handwritten content.

Signature verification is based on the assumption of handwriting being singular, exclusive and personal. Signature verification presents a double challenge. The first is to verify that what has been signed corresponds to the unique characteristics of an individual, ignoring the content of the writing. Rejection of an authentic signature is a false negative. The second and more demanding task is avoiding the acceptance of a forgery as authentic, which is referred to as a false positive.

Tolerance levels for signature verification is smaller than that for handwriting recognition. Some applications may require an error of 1 in 100,000 for false negative and even less for false positive. A trade-off exists between the errors. Signature verification systems work with an error margin of about 2–5% between the two errors.

Evaluation of signature verification algorithms raises several difficulties. Theoretically, it is not possible to measure the false-positive rate, since it is not possible to define a good forger and prove existence/nonexistence of a forger. Practically, several methods have been proposed. The simplest ones rely on the use of random forgeries, i.e., picking a random true signature of a person and considering it as a forgery of the signature of another person. Many studies incorporate unskilled forgeries, and, in some rare cases, highly skilled forgeries. Definitions of random, skilled and unskilled imitations vary.

Principal verification methods are probabilistic classifiers, time warping or dynamic matching, signal correlation, neural networks, hidden Markov models, Euclidian or other distance measures, hierarchical approach combining a few methods, and Baum-Welch training. At the analysis level, the main approaches have focused on spectral analysis, cosine transforms, direction encoding, velocity, timing and shape features sets, shape features force, pressure and angles functions.

Dynamic signature verification occupies a specific niche among identification systems. On the one hand,

they differ from systems based on the possession of something (key, card, etc.) or the knowledge of something (passwords, personal information, etc.) because they rely on a specific well-learned gesture. On the other hand, they also differ from systems based on the biometric properties of an individual (fingerprints, voice print, retinal prints, etc.) because the signature is still the most socially and legally accepted means for identification. Its unique, self-initiated motoric act provides an active means to simultaneously authenticate both a transaction and a transactioner.

Static Handwriting Recognition

The central tasks in off-line handwriting recognition are character recognition and word recognition. A necessary preliminary step to recognizing written language is the spatial issue of locating and registering the appropriate text when complex, two-dimensional spatial layouts are employed, a task referred to as document analysis.

Preprocessing operations performed prior to recognition are thresholding, or converting a gray-scale image into a binary black-white image; noise removal, extraction of foreground textual matter by removing, say, textured background, salt-and-pepper noise and interfering strokes; line segmentation, the separation of individual lines of text; word segmentation, the isolation of textual words; and character segmentation, the isolation of individual characters, typically those that are written discretely rather than cursorily.

Thresholding is to extract the foreground (ink) from the background (paper). The histogram of gray-scale values of a document image typically consists of two peaks, a high peak corresponding to the white background and a smaller peak corresponding to the foreground. So the task of determining the threshold gray-scale value (above which the gray-scale value is assigned to white and below which it is assigned to black) is one of determining an optimal value in the valley between the two peaks. The gray-scale histogram defines the optimal threshold as one that maximizes the between-class variance, where the distributions of the foreground and background points are regarded as two classes. Each value of the threshold is tried, and one that maximizes the criterion is chosen. Improvements to this basic idea, such as handling textured backgrounds, measure attributes of the resulting foreground objects to conform to standard document types.

Stroke connectivity has to be preserved during noise removal. Digital capture of images can introduce noise from scanning devices and transmission media. Smoothing operations are used to eliminate

artifacts introduced during image capture. Adaptive stroke filling with a neighborhood operator emphasizes stroke connectivity and prevents aggressive overfilling.

Interference of strokes from neighboring text lines is a problem that is often encountered. One approach is to follow strokes in thinned images to segment the interfering strokes from the signal. Thinning algorithms hold the promise of converting static handwriting to dynamic data, although artifacts such as spurs limit their use.

Segmentation into lines is more difficult than in the case of printed text. If the writing is on ruled lines, line segmentation is accomplished by examining the horizontal histogram profile at a small range of skew angles. When lines of text undulate up and down, ascenders and descenders frequently intersect characters of neighboring lines. One method is based on the notion that people write on an imaginary line that forms the core upon which each word of the line resides. This imaginary baseline is approximated by the local minima points from each component. A clustering technique is used to group minima of components to identify text lines.

Line separation is followed by a procedure to separate text into words. The focus is on identifying physical gaps assuming that gaps between words are larger than gaps between characters (Figure 2). However, exceptions are commonplace because of flourishes in writing styles with leading and trailing ligatures. Another method incorporates cues that humans use and does not rely solely on the one-dimensional distance between components. Writing style in terms of spacing is captured by characterizing the variation of spacing between adjacent characters as a function of the corresponding characters themselves. The notion of expecting greater space between characters with leading and trailing ligatures is incorporated into the segmentation scheme.

Isolation of words in a textual line is usually followed by recognizing the words themselves. Most recognition methods call for segmentation of the word into its constituent characters. Segmentation points are determined using features such as ligatures and concavities. Gaps between character segments




Figure 2 Word separation based on gaps.

(a character segment can be a character or a part of character) and heights of character segments are used in the algorithm.

Character Recognition

The basic problem is to assign the image of a handwritten character to its symbolic class. We discuss here some of the issues in the recognition of English orthography in its handwritten form. The typical classes are the upper and lower case characters, the ten digits and special symbols such as the period, exclamation mark, brackets, dollar and pound signs, etc. Shape features are extracted so as to assign the observed character to the appropriate class. Support-vector machines and artificial neural networks have emerged as fast methods for implementing classifiers. Algorithms based on nearest-neighbor methods have higher accuracy but are slower.

Recognition of a character from a single machine-printed font family on a well-printed paper document can be done very accurately. Difficulties arise when handwritten characters are to be handled. Some examples of segmented handwritten characters are shown in Figure 3. In the difficult cases, it becomes necessary to use models to constrain the choices at the character and word levels. Such models are essential in handwriting recognition because of the wide variability of handprinting and cursive script.

Word Recognition

A word recognition algorithm attempts to associate the word image to choices in a lexicon. Typically, a ranking is produced. This is done either by the analytic approach of recognizing the individual characters or by the holistic approach of dealing with the entire word image. The latter approach is useful in the case of touching printed characters and handwriting. A higher level of performance is observed



Figure 3 Examples of handwritten characters segmented from images.

by combining the results of both approaches. There exist several different approaches to word recognition using a limited vocabulary.

Analytic word recognition based on determining presegmentation points followed by determining an optimal path through a state transition diagram is shown in Figure 4. In the holistic approach, the word image is represented by a fixed size vector, for example by imposing a fixed grid on the word image (Figure 5). Other holistic features are the upper and lower profiles of word images, which are represented as a series of vectors describing the global contour of the word image and bypass the segmentation phase. Word recognition involves preprocessing, a possible segmentation phase that could be avoided if global word features are used, recognition, and postprocessing. Performance of the methods is a function of the size of the lexicon.

When dealing with large lexicon sizes, performance can be improved by using a dynamic lexicon. For example, the lexicon is represented as a tree and the results of recognizing the first few characters are used to eliminate possible paths in the tree. In another method, word images are oversegmented such that after segmentation no adjacent characters remain touching. Instead of passing on combinations of segments to a generic recognizer, a lexicon is brought into play early in the process. A combination of adjacent segments is compared to only those character choices that are possible at the position in the word being considered. The approach can be viewed as a process of accounting for all the segments generated by a given lexicon entry. Lexicon entries are ordered according to the goodness of match. Dynamic programming is used to string the potential character candidates into word candidates, combine heuristics to disqualify certain groups of primitive segments from being evaluated if they are too complex to represent a single character, and take into account compatibility between consecutive character candidates.

Application of Static Handwriting Recognition

A noteworthy success of static handwriting recognition has been in reading postal addresses. The task of interpreting handwritten addresses is one of assigning a mail-piece image to a delivery address. An address for the purpose of physical mail delivery involves determining the country, state, city, post office, street, primary number (which could be a street number or a post office box), secondary number (such as an apartment or suite number) and finally the firm name or personal name.

The recognition task is considerably aided by knowledge of the postal domain. The task is one of

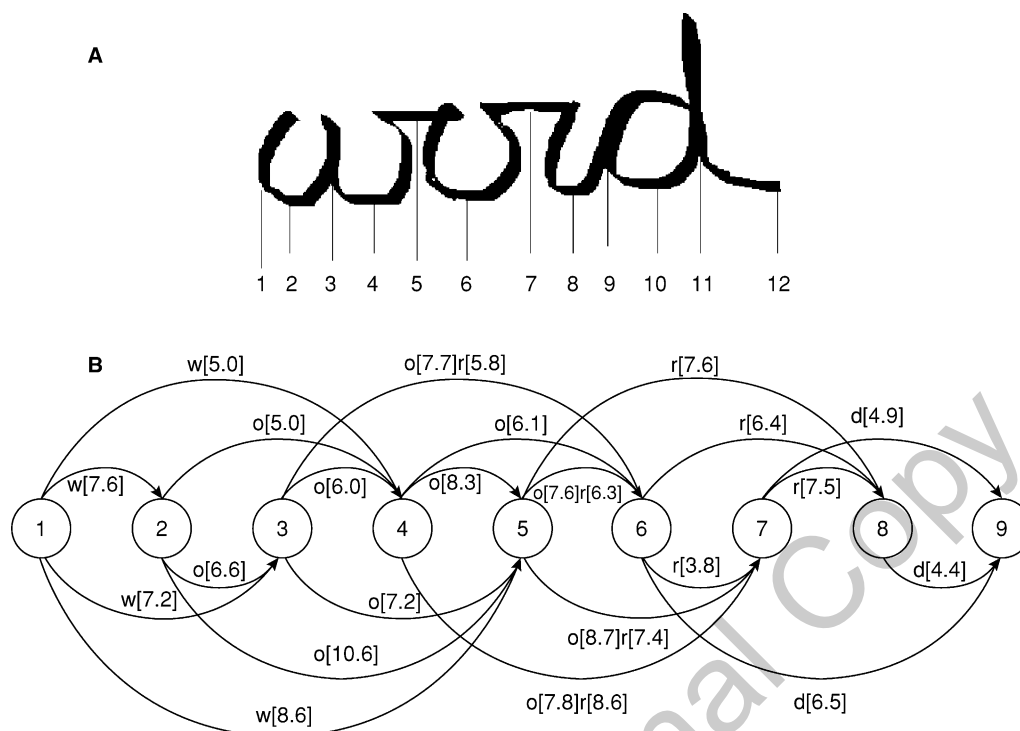


Figure 4 Analytic word recognition: (A) word with presegmentation points shown and (B) corresponding state transition diagram.




Figure 5 Holistic word recognition: a word is represented in an 8×4 grid and a feature vector is extracted.

interpretation rather than recognition since the goal is to assign the address to its correct destination irrespective of incomplete or contradictory information present in the writing. Systems currently in use by the United States Postal Service and other countries' postal services are able to correctly sort about 90% of handwritten addresses.

A gradation of class-discrimination problems is encountered: handwriting versus machine-print with two classes, handwritten numeral recognition with ten classes (Figure 6), alphabet recognition with 26 classes (Figure 7) and touching-digit pair recognition with 100 classes (Figure 8). Word recognition with a lexicon is a problem where the number of classes is dynamically determined by contextual constraints. Another problem encountered is similar to the problem of object recognition in computer vision: determining the destination address in a cluttered background.

Signature Verification

In static signature verification, a questioned signature image—as scanned and extracted from a bill, a check or document—is compared with known signature references. Unlike the dynamic case, there is no temporal information directly available, and the verification process relies on features extracted from the luminance of the trace only.

Although the extraction of a signature from a document background is already a very difficult problem in itself, particularly for checks, most of the studies published to date assume that an almost perfect extraction has been done. From a practical point of view, most researchers agree now that a solution will rely on the extraction from a signature of pseudodynamic features reflecting, for example, some specific characteristics used by a forensic document examiner as well as the automatic recovery of the stroke sequence in the signature image.

Handwriting Identification

Handwriting identification deals with comparing questioned writing with known writing exemplars and determining whether the questioned documents and exemplars were written by the same or different authors. Two issues of concern are the variability of handwriting within individuals, which are individual

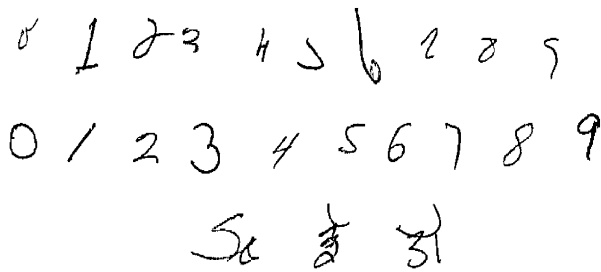


Figure 6 Ten-class discrimination (digits). Three sets of images are shown. The top row consists of difficult-to-read numerals. The middle row consists of fairly standard ones. The bottom row has non-digits, supplied to a digit recognizer, which must be rejected.

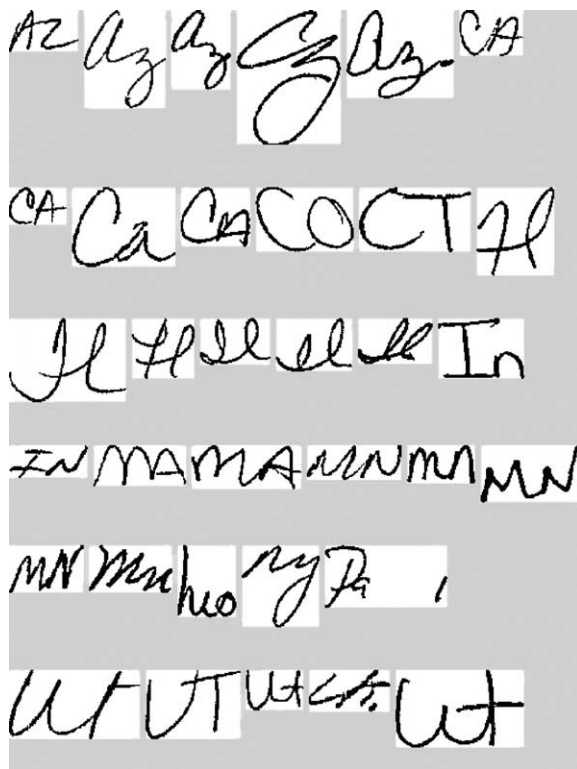


Figure 7 State abbreviations recognition with 66 classes. The valid abbreviations are AA, AE, AK, AL, AP, AR, AS, AZ, CA, CM, CO, CT, CZ, DC, DE, FL, FM, GA, GU, HI, IA, ID, IL, IN, KS, KY, LA, MA, MD, ME, MH, MI, MN, MO, MP, MS, MT, NC, ND, NE, NH, NJ, NM, NV, NY, OH, OK, OR, PA, PR, PW, RI, SC, SD, SQ, TN, TX, TT, UT, VA, VI, VT, WA, WI, WV, WY. The figure shows image snippets containing two-letter abbreviations collected from actual mail.

characteristics, and between individuals, which are class characteristics. The extraction of distinctive individual traits is what is relied on to determine the author of the questioned document.

Information about these two classes of variability is gathered based on the features for characterizing




Figure 8 Hundred-class discrimination (digit pairs). Touching digit pairs can be hard to separate, and it may be preferable to treat them as a single unit consisting of 100 classes. The figure shows touching-digit image pairs extracted from ZIP Codes.

handwriting. Some of the elements of comparison are alignment (reference lines), angles, arrangement (margins, spacing), connecting strokes (ligatures and hiatuses), curves, form (round, angular or eyed), line quality (smooth, jerky), movement, pen lifts, pick-up strokes (leading ligatures), proportion, retrace, skill, slant, spacing, spelling, straight lines and terminal strokes.

Several of these features are readily computable based on handwriting recognition techniques. For instance, handwriting recognition procedures routinely compute baseline angle and slant so that a correction can be applied prior to recognition. The result of applying these procedures is then used to cluster different samples of handwriting in a multidimensional feature space. The writership of the questioned document is then established from its proximity to the exemplars. Most handwriting identification experts today almost entirely rely on manually intensive techniques. Tools have recently become available for automated handwriting verification and identification (Figure 9).

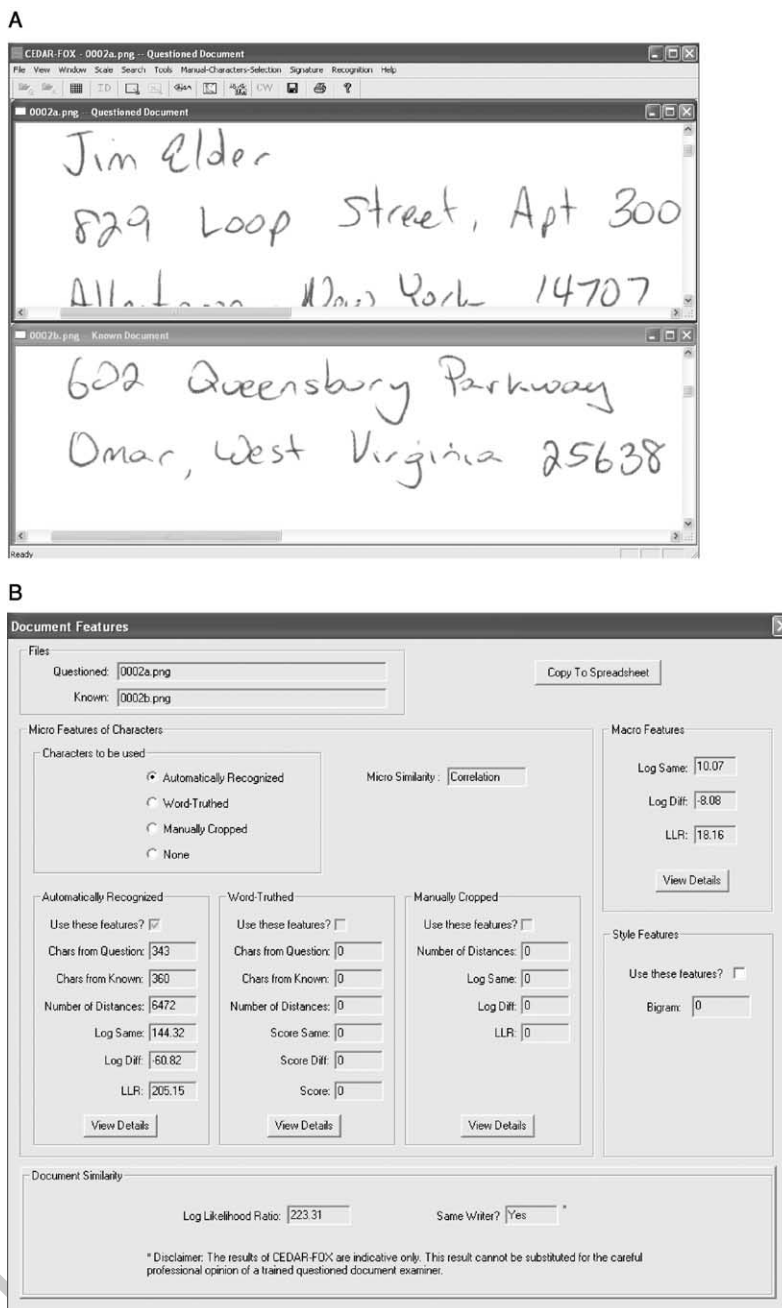


Figure 9 Screen shots of a software system for writer verification. The total log-likelihood ratio of the two documents (A) being written by the same/different writer are accumulated by the scores provided by several handwriting features (B).

Language Analysis and Processing: Language Models

Whatever the approach for recognition, dynamic or static, language models are essential in recovering strings of words after they have been passed through a noisy channel, such as handwriting or print degradation. The most important model for written language recognition is the lexicon of words. String matching algorithms between candidate words and

a lexicon are used to rank the lexicon, often using a variant of the Levenshtein distance metric that incorporates various edition costs into the ranking process. String matching methods are often improved by incorporating dictionary statistics in the training data. Lexical subsets, in turn, are determined by linguistic constraints, e.g., in recognizing running text, the lexicon for each word is constrained by the syntax, semantics and pragmatics of the sentence.




Figure 10 Handwritten sentence recognition. The path through top word choices is determined using part-of-speech tags.

The performance of a language model is evaluated in terms of the text perplexity, which measures the average number of successor words that can be predicted for each word in a text. The performance of a recognition system can thus be improved by incorporating statistical information at the word-sequence level. The performance improvement derives from selection of lower-rank words from word recognition output when the surrounding context indicates such selection makes the entire sentence more probable. Lexical techniques, such as collocational analysis, can be used to modify word neighborhoods generated by a word recognizer. Modification includes reranking, deleting or proposing new word candidates. Collocations are word patterns that occur frequently in language; intuitively, if word A is present, there is a high probability that word B also is present.

Methods to apply syntactic knowledge include N-gram word models, N-gram class (e.g., part-of-speech) models, context-free grammars and stochastic context-free grammars. N-Gram word models seek to determine the string of words that most probably gives rise to the set of output words that has been digitized or scanned. The problem with this approach is the difficulty of reliably estimating the parameters as the number of words grows in the vocabulary. A few alternatives to avoid this problem have been proposed: smoothing back-off models and maximum entropy methods. N-Gram class models map words into syntactic or semantic classes. In the first case, also referred to as the part-of-speech approach, for each sentence that has to be analyzed, a lattice of word/tag assignments is created to represent all possible sentences for the set of possible word candidates. The problem is to determine the best path through the lattice. The semantic approach relies mainly on machine-readable dictionaries and electronic corpora; it uses word definition overlaps between competing word candidates to select the correct interpretation. Other approaches that involve collocations are co-occurrence relations and the use of semantic codes that are available in some dictionaries.

So far, these approaches have been limited to proof of concept, and no large-scale experiments have been

reported to demonstrate the effectiveness of semantic information in resolving ambiguities, although real-life analysis of human behavior suggests that this is very often the only way to proceed. An example of a handwritten sentence together with recognition choices produced by a word recognizer and grammatically determined correct paths is shown in [Figure 10](#). An increase in the top-choice word-recognition rate, for example, from 80% to 95%, is possible with the use of language models.

Conclusion

Handwriting recognition has come of age from being a sleepy research area to general acceptance. In the dynamic or on-line case, handwriting recognition in personal digital assistants and tablet PCs is impressive. Dynamic signature verification systems have been marketed over the last few years. Most of the static or off-line successes have come in constrained domains, such as postal addresses, bank checks and census forms. The analysis of documents with complex layouts, recognition of degraded printed text and the recognition of running handwriting continue to remain largely in the research arena. While machine performance is beginning to become acceptable in both handwriting recognition and handwriting identification, major research challenges exist in word and line separation, segmentation of words into characters, recognition of words when lexicons are large, and the use of language models in aiding preprocessing and recognition.

Bibliography

- Plamondon R & Srihari S N (2000). 'On-line and off-line handwriting recognition: a comprehensive survey.' *IEEE Transactions on Pattern Analysis and Machine Intelligence* 22(1), 63–84.
- Srihari S N (1992). 'High performance reading machines.' *Proceedings of IEEE* 80(7), 1120–1132.
- Srihari S N, Cha S, Arora H & Lee S-J (2002). 'Individuality of Handwriting.' *Journal of Forensic Sciences* 47(4), 856–872.