

Implementing Word Retrieval in Handwritten Documents using a Small Dataset

Y. Liang, R.M. Guest, M.C. Fairhurst*

School of Engineering and Digital Arts, University of Kent

**Corresponding author: m.c.fairhurst@kent.ac.uk*

Abstract

A novel approach to the problem of keyword retrieval in cursive handwritten documents is introduced in this work. Two issues are addressed: small dataset size and uneven sample distribution across the character set. The proposed strategies utilise graphemes (fragments of a handwritten word) to implement a recognition model which is subsequently used to form the feature model for the query word.

1 Introduction

The requirement for automated handwriting recognition has long been established across many application domains. While automatic handwriting recognition is still a very challenging task, the past decade has seen a proliferation of applications for this technology. The work described here aims to address difficulties often encountered in the context of a specific application: keyword retrieval in handwritten document images. The application is generally defined by the following characteristics:

- 1) The stored handwritten documents are in an image format.
- 2) The query word is expressed as a set of ASCII characters.
- 3) Samples of the query word are not necessarily available to the system before the retrieval.
- 4) The result of the query is a list of images segmented from the documents representing a potential match of the query word.

The second and third characteristics further distinguish keyword retrieval from a similar problem: word spotting [1-6]. In word spotting approaches, a model is created for each query word, and is trained with samples of the exact word. Consequently, in this approach, the query word must have been “seen” (i.e. instances of the query word must be provided for the training process) by the system before it is retrieved, and hence the system is not able to retrieve words that have not been seen by the system. These words are referred to as out-of-vocabulary or OOV words. In keyword retrieval approaches [7-12], the query word is represented by models of the individual characters. Specific instances of the query word are not required

either during the training process or upon a query request. Therefore, systems developed using this type of approaches are able to search for OOV words.

Keyword retrieval approaches are generally preferable to the word spotting methods due to the enhanced flexibility in application as discussed above. However, in addition to the difficulties commonly found in all handwriting analysis, keyword retrieval approaches introduce further possible performance related issues: character segmentation, human effort in providing the training data, and the likelihood of an uneven sample size. The aim of this study is to address these issues in a novel approach to implementing the keyword retrieval application.

2 Issues and proposed strategies

Human intervention is typically used to provide suitable training data for word retrieval systems. This task usually entails line segmentation, word segmentation, transcription, and labelling (associating the segmentations with their precise transcription). Character segmentation and character labelling were implemented manually in [7], providing a moderate and uniform sized training dataset - 32 samples - for each character. The character models in [7] are established using a joint-boosting classifier and the probability of detecting a query word is evaluated by a HMM-based method. In other work [10-12], segmentation and labelling was carried out on a line-by-line basis. A method based on recurrent neural network is developed to infer the probability relation between each character and the set of features extracted from a vertical one-pixel-wide window slide across the text line. The character recognition models described in [8, 9] are established using a publicly available English character database. However, the recognition rate achieved shows a significant decline from those in the previous two approaches.

In terms of system usability, reducing human effort in the process of providing suitable training data is an important goal of our study. Hence, the first strategy in our work is that an automatic character segmentation process will be devised to provide the training dataset based on a small number of pages from each writer.

Automatic character segmentation in itself is a

challenging process in handwriting analysis [13], often resulting in over- or under- segmentation. A commonly adopted method is to consider the segmentations as preliminary outcomes, termed as “graphemes”, which are subsequently subjected to further analysis based on linguistic context [13-16].

Due to the nature of language, the number of samples extracted from a piece of text will typically vary for each character in the alphabet. A significant imbalance in the number of training samples across classes is not favourable in solving pattern recognition problems [17, 18], neither is the potential for small sample sizes.

Motivated by the strategy of automatic character segmentation, a novel approach based on the analysis of graphemes, termed a *grapheme spectrum*, is proposed in this work to address the issue of small and imbalanced sample size. The grapheme spectrum approach to character modelling uses the same underlying principle in a technique known as bag-of-features (BOF) [19] in that a word image is decomposed into a number of areas (i.e. graphemes in our work) each of which is represented by a set of features. This technique has the potential be used directly in word spotting. In addition to a detailed implementation of the BOF technique, however, our approach addresses the added obstacles to keyword retrieval..

Two strategies are adopted in the proposed approach: a) the graphemes correspond to short strokes in handwriting that are always smaller than or equal to a character. The proposed approach differs from reported grapheme-based character segmentation methods [20] in that the aim is to decompose a character into recognisable portions that can be extracted reliably and repeatedly, which forms the basis of the following grapheme spectrum recognition method. b) The recognition models are trained to recognise graphemes instead of characters. A benefit of this approach is that by replacing characters with graphemes as the classes in the recognition problem, the sample size of each class is effectively boosted. Graphemes are shared by more than one character: for example, a loop can be observed in a number of characters, including, for example, a, b, e, g, o, p, and q. Therefore, a second benefit of this approach is that the imbalanced sample size is not as big an issue with graphemes as it is with characters.

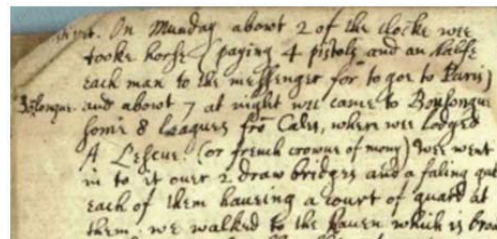
3 Datasets

For the purpose of this study, three manuscripts of diverse writing style and age are analysed: Bargrave’s travel diary (1645) [21], George Washington’s documents [5] (1755), and a modern handwriting sample donated by a local writer at the time of this

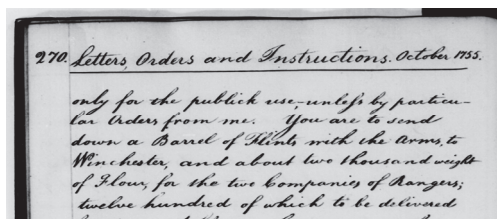
study. A fragment of a page from each document is shown in Figure 1.

3.1 Ground truth data for training


For training, word images are extracted manually from three pages of each manuscript. Transcription is provided for each word image.



a) Diary - John Bargrave's travel diary



b) GW - George Washington's documents



c) Modern – Handwriting sample provided by a local writer

Figure 1 - Manuscripts

3.2 Image pre-processing for testing data

For testing data, a fully automated process is devised to acquire the word segmentations. After a binarisation process using the technique proposed in [22], only black-and-white information is retained in the manuscript images. Preliminary line segmentation is obtained by analysing the horizontal projection of the pixel values. Automatic skew correction is performed on a line-by-line basis by estimating the regression angle of all pixels on the text line. The text lines are not explicitly segmented into words at the pre-processing stage. Instead, in response to the query of a keyword, the region(s) within the text line that most likely contain(s) an instance of the keyword is (are) extracted in real-time. The likelihood is assessed by the proposed grapheme spectrum method which will be described in the following section.

4 Grapheme spectrum approach

This section describes the steps taken to construct the grapheme spectrum for each character, including: grapheme segmentation, grapheme recognition and forming the grapheme spectrum, followed by the method to test the hypothesis that a sub-image within a text line is an instance of the query word.

4.1 Grapheme segmentation

Most approaches found in similar grapheme segmentation studies (see, for example [20]) aim to obtain fragments approximating individual characters, which may result in parts of a character or combinations of two to three characters. The approach adopted in this work, however, aims to segment individual characters into meaningful portions that closely represent natural handwritten strokes, e.g. horizontal/vertical strikes, diagonal strikes, loops, concave/convex strokes.

The segmentation method can be described as follows:

- 1) Extract the skeleton of the word
- 2) The word will be divided into graphemes by the following pixels on the skeleton of the word:
 - a) local minima
 - b) local maxima
 - c) branch point (a pixel that has more than two neighbouring pixels in the 8-connected neighbourhood)
- 3) Preserve loops by connecting graphemes that comprise a loop

4.2 Grapheme recogniser

Three observations can be made from the obtained graphemes:

- 1) The graphemes are always smaller than or equal to a single character.
- 2) The same grapheme can be observed in a large number of characters. In addition to the example given in Section 2, many characters contain a vertical strike, including b, d, h, p, and q.
- 3) Using the method described in 4.1, the same set of graphemes can be repeatedly extracted from most instances of the same character.

These properties are exploited in this work to address the issues regarding the small and imbalanced sample size across all characters. In comparison with a dataset consisting of character samples that can be extracted from the same piece of text, the first and second properties result in a relatively large and uniformly sized dataset across all graphemes. The repeatability of the segmentation method allows the

implementation of grapheme recognisers, which are subsequently used in character modelling.

Because the graphemes are not labelled, an unsupervised learning algorithm has been chosen to implement the grapheme recogniser. From the candidate unsupervised learning algorithms, the self-organising map (SOM) [23] is chosen, because it offers the advantage of learning the topological structure of the data as well as the class identities, and it has been successfully employed in the analysis of handwriting styles [16]. A grid layout topology of the SOM is adopted in this study. A k-fold validation experiment on keyword matching using segmented word images from the three datasets is devised to determine the optimal size of the map, with the result being 9-by-9. As an input to the SOM, graphemes are expressed by the x-y coordinates. The outcome of the training is termed a map-of-graphemes (MOG).

4.3 Character segmentation and grapheme spectrum

The output of the grapheme recogniser is utilised to form the character models. Before we continue to construct the character model, an automated process must be devised to associate the graphemes with the character from which they are most likely extracted. Because the words in the training dataset are labelled, the character segmentation process makes use of the contextual information provided in the word labels, i.e. orders of the characters, the presence of ascenders, descenders, and/or capital letters.

A character model is initially a collection of the graphemes extracted from all instances of this character. Each grapheme is expressed by the topological position of its winning node in the MOG. The model keeps a count for each neuron in the MOG. Therefore, the model is a vector, of which the value of each element represents the frequency of the corresponding neuron being assigned to the graphemes extracted from the instances of this character. By dividing the frequency by the total number of instances of the character, the value is translated to the probability of observing this character if such a grapheme is detected. The grapheme spectrum is expressed in Eq. 1, where n_i is the number of times the i -th neuron is the winning node of a grapheme of the represented character, and S is the total number of instances of the character in the training set.

$$p(i) = \frac{n_i}{S}, \quad 1 \leq i \leq N \quad \text{Eq. 1}$$

Each character model, as illustrated in Figure 2, is a frequency spectrum, hence the designation *grapheme spectrum*.

Note that the errors of the character segmentation

are carried over to the grapheme spectrum. Based on the assumption that most graphemes are assigned to the correct character, the errors can be identified by the low frequency entries in the spectrum. Therefore, by setting the frequency entries that are smaller than the t -th percentile of the entire spectrum to zero (t is determined empirically), most of the errors resulting from character segmentation are filtered.




Figure 2 – Grapheme spectrum for ‘a’

4.4 Keyword retrieval hypothesis evaluation

For each query word, a template is formed by referring to the models of the characters comprising the word. The word model is expressed in Eq. 2, where k corresponds to the k -th character in the word, and M is the total number of characters comprising the word. Definitions of other symbols are as defined in Eq. 1

$$p(k, i) = \frac{n_{ik}}{S_k}, \quad 1 \leq k \leq M, \quad 1 \leq i \leq N \quad \text{Eq. 2}$$

Retrieval is essentially a process of evaluating the hypothesis that a word image is an instance of the query word. Thus, the image must also be expressed in terms of graphemes, as illustrated in Eq. 3, where L_j is the label of the winning neuron for the j -th grapheme.

$$q(j) = L_j, \quad 1 \leq j \leq G, \quad 1 \leq L \leq N \quad \text{Eq. 3}$$

The distance between two graphemes is assessed by their topological positions on the MOG [23]. Using variable d to denote the distance function, the distance between the j -th grapheme in the test word image and the i -th entry in the grapheme spectrum of the corresponding character is written as $d_{i,q(j)}$.

Regardless of the spatial position of the graphemes in the test word, the hypothesis can be evaluated character-wise based on two factors: a) the topological distance between the individual graphemes in the test word image and the non-zero entries in the grapheme spectrum of the corresponding character, b) the frequency values in the grapheme spectrum of the entry corresponding to the smallest topological distance to the individual graphemes in the test word image. These criteria are expressed in Eq. 4 for the K -th character in the query word. The definitions of p , i

and N can be found in Eq. 2, whereas q , j , and G are defined in Eq. 3, and d is the distance function as described above. The \max function in Eq. 4 expresses a maximisation process, which assigns an individual grapheme in the test word to the entry in the grapheme spectrum of the assumed character that maximise the outcome.

$$\gamma = \sum_{1 \leq j \leq G} \max_i \left(\frac{p(K, i)}{d_{i,q(j)}} \right), \quad 1 \leq i \leq N \quad \text{Eq. 4}$$

However, using Eq. 4, it is possible to find that a grapheme at the left hand side of a word image is considered to be part of a character at the end of the query word. In order to include spatial position of the graphemes into the equation, a process called *hypothetical character segmentation* is introduced here – the graphemes in the test word image are segmented into characters based on contextual information in the query word. The result is written as in Eq. 5, where M is the total number of characters contained in the query word, and the use of k in combination with j denotes that the j -th grapheme is associated with the k -th character.

$$r(k, j) = \begin{cases} 1, & \text{if the } j\text{-th grapheme is associated with the } \\ & k\text{-th character} \\ 0, & \text{otherwise} \end{cases} \quad 1 \leq j \leq G, \quad 1 \leq k \leq M \quad \text{Eq. 5}$$

Combining Eq. 4 and Eq. 5, the evaluation of the hypothesis is updated to Eq. 6.

$$\gamma = \sum_{1 \leq k \leq M} \left(\sum_{1 \leq j \leq G} \max_i \left(\frac{p(k, i) \times r(k, j)}{d_{i,q(j)}} \right) \right), \quad 1 \leq i \leq N \quad \text{Eq. 6}$$

5 Experiments and Results

5.1 Experimental configuration

From each document, word images are extracted from three pages to form the training dataset. The testing dataset contains one to two page images from each document. The chosen keywords appear on the testing pages, but are not in the training dataset.

As discussed in Section 2, the algorithm devised in this work is intended to perform under the constraints of a small training dataset and uneven sample size. The configurations of the experiments, therefore, aim to assess the performance of the algorithm under these constraints and the potential for improvement when the constraints are relaxed. Therefore, the keywords within each document are divided into two groups based on the smallest sample size for the characters contained in

the word, and ten samples is considered here as the separation between small and moderate sample sizes, because this results in relatively even separation between the two groups of the keywords. A brief summary of the experimental configurations can be found in Table 1.

Table 1 – Experimental configurations

Smallest sample size		>10	<=10	>10	<=10	>10	<=10
Document		Diary		GW		Modern	
Training data	p	3		3		3	
	u	163		226		298	
	w	267		372		437	
Testing data	p	1	1	1	1	2	2
	o	40	21	25	22	80	53

p: number of pages, u: number of unique words, w: number of word samples, o: number of unique OOV keywords

5.2 Assessment metric


The performance is evaluated using two common metrics in information retrieval – precision and mean average precision (MAP). Both metrics result in a value ranging from 0 to 1 with a higher value representing better performance. Definition of these two metrics can be found in the literature relating to information retrieval [24].

5.3 Performance and discussion


The performance in the six experiments described in Table 1 is assessed by MAP and precision at rank one as shown in Figure . The best performance is a MAP of 57% achieved with the Modern manuscript shown in Figure 3 a) when the smallest training sample size is greater than ten for all characters, corresponding to the precision at rank one of 53% in Figure 3 b).

A review of relevant work reported in the literature, in particular in terms of the ability to search for OOV words, is given in Table 2. In comparison with other studies, a considerably smaller training dataset is adopted in this present study. With the exception of the sixth experiment, the performance achieved in our study is superior to that achieved in [7-9], while at the same time the proportion of testing data adopted in our work is greater than that in [7] and similar to that in [8, 9]. The work described in [10] addressed the keyword retrieval problem using a handwriting recognition approach. The GW20 adopted in [10] is a small database comparable to the number of pages from the GW manuscripts adopted in this study. When trained and tested on the GW20 database using a four-fold cross validation, the system achieved an average precision of 86% on the chosen lexicon words, although the method is capable of spotting OOV words. Therefore, it is difficult to compare the

performance achieved in this study with that reported in [10], with respect to the ability of retrieving OOV words.



a) Mean average precision



b) Precision at rank one

Figure 3 – Keyword retrieval performance

Table 2 – Comparison of reported works

Ref.	Dataset	Claimed performance
[7]	GW20 ^a	Accuracy: 84% for lexicon words 32% for OOV words
[8, 9]	1125 PCR forms ^b	Precision at rank one <30%
[10]	1,539 pages from the IAM [25]	Average precision 59-77%
	GW20 ^a	Average precision 67-86% for lexicon words

a) 20 pages taken from George Washington's manuscripts [7]

b) New York State Pre-hospital Care Report (PCR) forms

In addition to the headline performance, the most important aspect of this work is to demonstrate the potential of improving the performance when the constraints are relaxed, i.e. when the number of samples available for training purpose increases. It can be seen from Figure 3 that the performance improved within each manuscript with respect to the configuration outlined in Table 1. Another observation that can be made is that the performance within an older manuscript is poorer. Instead of associating the

performance with the age of manuscript, the actual writing style and layout are considered to be the cause. Regardless of the poorer performance in the Diary, the potential of improving the performance by increasing the number of training samples for each character is encouraging.

6 Conclusion

In summary, we describe in this paper a novel approach to the keyword retrieval problem in cursive handwritten documents. The goal of this study is explicitly to retrieve OOV words, while at the same time addressing two prominent issues: small training dataset sizes and non-uniform sample distributions for the characters. The method introduced in this paper has achieved very encouraging results, which also show advantages over other comparable methods with respect to the particular context of application.

It is also worth noting that, unlike most similar work reported in the literature, automated pre-processing procedures (including skew correction, line segmentation, and implicit word segmentation at the testing phase) can be applied to the manuscript page images in the testing dataset, and hence no human intervention is required once the system is trained. However, the errors produced by the automated segmentation are carried over to the retrieval stage. Therefore, the retrieval performance can possibly be improved in the future by enhancing the pre-processing stage. While a dataset with limited size is used in pattern recognition studies to evaluate the performance expectancy, the performance does not always scale linearly. It is our intention to investigate this scalability issue in our future work,

Acknowledgement: The authors gratefully acknowledge the support of the EU INTERREG IVA France (Channel) – England Programme and the Canterbury Cathedral Archives in the production of this work.

References:

- [1] N. R. Howe, *et al.*, "Boosted decision trees for word recognition in handwritten document retrieval," in *ACM SIGIR*, New York, USA, 2005, pp. 377-383.
- [2] T. van der Zant, *et al.*, "Handwritten-Word Spotting Using Biologically Inspired Features," *IEEE TPAMI*, vol. 30, pp. 1945-1957, 2008.
- [3] T. M. Rath, *et al.*, "A Statistical Approach to Retrieving Historical Manuscript Images without Recognition," Center for Intelligent Information Retrieval, University of Massachusetts 2003.
- [4] Y. Leydier, *et al.*, "Text search for medieval manuscript images," *PR*, vol. 40, pp. 3552-3567, 2007.
- [5] T. M. Rath, *et al.*, "A Search Engine for Historical Manuscript Images," presented at the ACM SIGIR, Sheffield, United Kingdom 2004
- [6] M. Rusinol, *et al.*, "Browsing Heterogeneous Document Collections by a Segmentation-free Word Spotting Method," 2011, pp. 63-67.
- [7] N. R. Howe, *et al.*, "Finding words in alphabet soup: Inference on freeform character recognition for historical scripts," *PR*, vol. 42, pp. 3338-3347, 2009.
- [8] H. Cao, *et al.*, "A probabilistic method for keyword retrieval in handwritten document images," *PR*, vol. 42, pp. 3374-3382, 2009.
- [9] H. Cao, *et al.*, "Unconstrained handwritten document retrieval," *IJDAR*, vol. 14, pp. 1-13, 2010.
- [10] V. Frinken, *et al.*, "A Novel Word Spotting Method Based on Recurrent Neural Networks," *IEEE TPAMI*, vol. 1, pp. 1-14, 2011.
- [11] A. Graves, *et al.*, "Connectionist temporal classification: Labelling unsegmented sequence data with recurrent neural networks," in *ICML*, 2006, pp. 369-376.
- [12] A. Graves, *et al.*, "A novel connectionist system for unconstrained handwriting recognition," *IEEE TPAMI*, vol. 31, pp. 855-868, 2008.
- [13] R. Casey and E. Lecolinet, "A survey of methods and strategies in character segmentation," *IEEE TPAMI*, vol. 18, pp. 690-706, 2002.
- [14] A. El-Yacoubi, *et al.*, "An HMM-based approach for off-line unconstrained handwritten word modeling and recognition," *IEEE TPAMI*, vol. 21, pp. 752-760, 2002.
- [15] K. M. Sayre, "Machine recognition of handwritten words: A project report," *PR*, vol. 5, pp. 213-228, 1973.
- [16] L. Schomaker, *et al.*, "Using codebooks of fragmented connected-component contours in forensic and historic writer identification," *Pattern Recognition Letters*, vol. 28, 2007
- [17] M. A. Mazurowski, *et al.*, "Training neural network classifiers for medical decision making: The effects of imbalanced datasets on classification performance," *Neural Networks*, vol. 21, pp. 427-436, 2008.
- [18] S. J. Raudys and A. K. Jain, "Small Sample Size Effects in Statistical Pattern Recognition: Recommendations for Practitioners," *IEEE TPAMI*, vol. 13, pp. 252-264, 1991.
- [19] E. Nowak, *et al.*, "Sampling strategies for bag-of-features image classification," *Computer Vision-ECCV 2006*, pp. 490-503, 2006.
- [20] T. Saba, *et al.*, "Methods and strategies on off-line cursive touched characters segmentation: a directional review," *JAIR*, pp. 1-20, 2011.
- [21] (2009), *The Bargrae Collection*. Available: <http://canterbury-cathedral.org/assets/content/bargrae/index.html>
- [22] Q. Huang, *et al.*, "Thresholding technique with adaptive window selection for uneven lighting image," *Pattern recognition letters*, vol. 26, pp. 801-808, 2005.
- [23] T. Kohonen, *Self-organising maps*, 3rd ed. Berlin: Springer, 2001.
- [24] R. Baeza-Yates and B. Ribeiro-Neto, *Modern information retrieval*. ACM press New York., 1999.
- [25] H. Hiary and K. Ng, "A system for segmenting and extracting paper-based watermark designs," *International Journal on Digital Libraries*, vol. 6, pp. 351-361, 2007.