<http://www.diva-portal.org>

Postprint

This is the accepted version of a paper presented at *22nd International Conference on Pattern Recognition (ICPR), 24-28 Aug, 2014, Stockholm, Sweden.*

Citation for the original published paper:

Dahllöf, M. (2014)

Scribe attribution for early medieval handwriting by means of letter extraction and classification and a voting procedure for larger pieces.

In: *22nd International Conference on Pattern Recognition (ICPR)* (pp. 1910-1915).

International Conference on Pattern Recognition

<http://dx.doi.org/10.1109/ICPR.2014.334>

N.B. When citing this work, cite the original published paper.

© 2015 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Permanent link to this version:

<http://urn.kb.se/resolve?urn=urn:nbn:se:uu:diva-241360>

Scribe attribution for early medieval handwriting by means of letter extraction and classification and a voting procedure for larger pieces (ACCEPTED VERSION)

Mats Dahllöf

Dept of Linguistics and Philology and Dept of Information Technology, Division of Visual Information and Interaction
Uppsala University, Sweden
E-mail: mats.dahllof@lingfil.uu.se

Abstract—The present study investigates a method for the attribution of scribal hands, inspired by traditional palaeography in being based on comparison of letter shapes. Data was drawn from early medieval Caroline minuscule manuscripts. The generation of a prediction for a page image involves writing identification, letter segmentation, letter classification, and, finally, hand prediction, by voting, from a letter proposal list. Letters and sequences of connected letters are identified by means of connected component labeling and split into letter-size pieces. The hand (and character) prediction makes use of a dataset containing instances of the letters *b*, *d*, *p*, and *q*, cut out from manuscript pages whose scribal origin is known. Letters are represented by features capturing the distribution of foreground and of background enclosed by foreground. Cosine similarity is used for nearest neighbor classification. The hand behind a page is finally predicted by means of a voting procedure taking the highest scoring letter-level hits as its input. This hand prediction method was evaluated on pages from five different hands and reached an accuracy above 99% for four of them and 84% for a fifth significantly more difficult one. The hand behind single top-listed letters was correctly predicted in 82% of the cases.

I. INTRODUCTION

An important issue in palaeography is to classify styles and to attribute pieces of handwriting to scribes. This paper describes a method for doing this by extracting plausible letter candidates from manuscript pages and to predict the scribal hand behind these by means of a nearest neighbor classification procedure using a dataset of cut-out letter instances from known hands. Page-level scribe predictions are generated from the letter-level ones, by a voting procedure. The same letter-level classifier also produces character predictions, and its performance in that regard will also be discussed. The system was developed for the early medieval style Caroline minuscule, and evaluated on five works, each representing one scribal hand, containing about 1000 manuscript pages.

II. PREVIOUS STUDIES

The issue of predicting who physically produced a piece of handwriting (the “scribe”, “hand”, or “writer”) has received considerable attention. In the recently established field of “digital palaeography”, where historical material has been at the centre of interest, there is a hope that the use of computational methods will be a way of supporting and developing the methods for scribal attribution in traditional palaeography. When

it comes to modern handwriting, an important motivation for computational analysis is provided by forensic science, where the purpose is to establish facts about the origin of documents with a level of certainty appropriate for conclusions put forth as legal evidence. As can be expected, computational systems for hand prediction make use of machine learning. A feature extraction mechanism and a classification framework can be seen as the two central components, both of which allow a variety of design choices.

Following Jain and Doerman [6], we can distinguish between literate and illiterate methods for feature generation, where the former analyze stretches of writing as sequences of letters, and the latter treat them as pen tip traces, without ascribing any conventional linguistic structure to them. The virtue of the illiterate approaches is that they do not rely on transcriptions of the manuscript images, and that they, at least in principle, can be applied independently of writing system and style. Examples of features of this kind are probability distributions of character fragment contours (based on a codebook generated from training data), as suggested by Schomaker and his coworkers [9]. The idea is developed by Bulacu and Schomaker [2], with character fragments represented by normalized bitmaps. These “sub or supra-allographic fragments” are called “graphemes”, and do not in general directly correspond to characters (as “graphemes” do in the terminology of linguists). Other examples of “illiterate” features are distributions of stroke fragments (Tang, et al. [11]), and joint probability distributions of the orientations of any two hinged edge fragments (Bulacu and Schomaker [2]). Brink et al. [1] propose features based on the relation between the local width and direction of ink traces in a probability distribution. The approach of Jain and Doerman [6] relies on a simple character segmentation scheme and clustering of proposed segments to derive a “pseudo-alphabet” of contour gradient descriptors. The proposal involves a distance measure over these pseudo-alphabets.

The individual practices of different scribes are also reflected in the layout of the manuscript pages. This fact is used by De Stefano et al. [4] for hand prediction based on layout features, and the method is applied to a 12th century Caroline minuscule multi-hand manuscript.

These methods use different clustering procedures to generate codebooks and pseudo-alphabets of recurring writing ele-

© 2014 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Published article (2014 22nd International Conference on Pattern Recognition):
DOI:10.1109/ICPR.2014.334
<http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=6977046>

ments, whose distribution is used to characterize the individual features of handwriting. The task of deciding which hand to propose is in many systems left to nearest-neighbor classifier ([9], [2], [1], [6], [11]). A multi layer perceptron was used in [4].

According to an overview of nine approaches, whose author (Brink et al. [1]) however notes that “the numbers cannot be well compared because of differences in dataset material, required level of human interference, and number of writers”, hand identification systems achieve accuracy rates between 62% and 97% for modern handwriting, when several hundred hands are included in the data. The results reported by Brink et al. [1] for medieval manuscript data amount to lower accuracy rates for smaller sets of scribes, suggesting that hand prediction for medieval handwriting is a more difficult task. An explanation for this may be that medieval handwriting, in particular the items that have been preserved, typically represents the standardized craft of educated scribes, whereas modern handwriting is an informal and non-professional mode of expression, which allows various personal idiosyncrasies to shine through.

Discussing the advantages of using computational tools in palaeography, Stokes [10] sees the need for applying more quantitative methods in this field. Traditional palaeographers have in general relied on describing handwriting in qualitative terms. Their conclusions about the scribal origin of pieces of writing likewise tend to be based on qualitative evidence. The methods of such a discipline tend to be difficult to communicate, to reproduce, and to validate. By contrast, a computer-aided palaeography, making use of techniques inspired by those developed in forensic analysis of modern handwriting, will be better equipped to characterize historical writing in precise and quantitative terms, and to work with models which can be objectively validated. Another contribution to computer-aided palaeography is Ciula’s [3] “Java System for Palaeographic Inspections”, which allows researchers in palaeography to study letter shapes in quantitative detail.

III. SOURCE MATERIAL

A. Manuscript pages and letter images

The present study makes use of two kinds of data: First, there is a collection of manuscript images, from five codicological units, each of which represents a single scribe. Secondly, the generation of hand predictions relies on a dataset of letter images, cut out from the manuscript pages. Each letter image is labelled with a scribe and a character (letter type, grapheme) tag.

B. Codicological units

The Caroline (or Carolingian) minuscule is the most important book style of the early medieval period in Europe, well-known for being easier to read than the later Gothic styles for both visual and linguistic reasons. Strictly speaking, the term *minuscule* is used for both a kind of letter (corresponding to lowercase) and a writing style (in which minuscule letters are used in the body text). The scribes were generally skilled craftsmen, who produced a highly standardized and formal

TABLE I. THE FIVE CODICOLOGICAL UNITS (EACH REPRESENTING ONE “HAND”) PROVIDING DATA FOR THE PRESENT STUDY: TEXT EXAMPLE, SHELFMARK, DATE (CENTURY), NUMBER OF PAGES INCLUDED IN THE DATASET, AND THEIR DIVISION (INTO LETTER SOURCE PAGES, TUNING DATA, AND EVALUATION DATA).

Shelfmark, unit	Date	Pages	Division

Cod. Sang. 112.	9th c.	322	1+32+289

Cod. Sang. 186 (unit p. 3–146).	Early 9th c.	143	1+14+128

Cod. Sang. 557.	9th c.	270	2+27+241

Cod. Sang. 562 (unit p. 3–93).	Late 9th c.	91	2+10+79

Cod. Sang. 565 (unit p. 3–222).	10th or 11th c.	220	2+22+196
all five		1046	8+105+933

kind of handwriting. However, even if the Caroline minuscule is not a cursive style, many of the letters are connected.

Many Caroline codices and their metadata have been published by the *e-codices* website hosted by the University of Fribourg. Five well-preserved works, see table I, were selected for the present study, all written at the Abbey of St. Gall (Switzerland) and still belonging to the St. Gall Stiftsbibliothek.

High-quality digital reproductions of these books are found at <http://www.e-codices.unifr.ch/en/list/one/csg/0112>, etc. The highest resolution images provided by the website (“x-large”, “converted into JPEG files [...] and minimally processed to improve legibility”) have served as manuscript page data. (Pages from Cod. Sang. 562 are also included in the dataset compiled by Fischer et al. [5].) Their use is regulated by a Creative Commons License. According to palaeographers’ verdicts, quoted at the website, each codicological unit represents a single scribe (“hand”). It will be assumed that we are dealing with five different scribes, i.e. in the present context, one for each codex number. Four pages (from Cod. Sang. 557) without any minuscule writing were removed from the page dataset. There are between 2 and 24 lines of minuscule writing on each manuscript page.

IV. PROPOSED METHOD

A. General considerations

The present method is inspired by traditional palaeography, where writing styles and scribal hands are analyzed through inspection of letter forms. Its mode of operation consequently allows us to generate justifications for the predictions which would be close to the reasoning of a traditional philologist. Another motivation for this approach is that it fits within a

framework for textual search and interpretation (OCR), where human supervision through letter data would be necessary for predicting which character each proposed letter instance represents.

The system involves scaling and binarization as preprocessing steps, writing identification, letter segmentation, and uses the letter dataset to predict the hand and character of an arbitrary letter image. The hand behind a codex page is predicted by a voting procedure operating on a toplisted subset of the letter-level proposals generated for that page. A Java implementation was used to develop and evaluate the method.

B. Data and development

Letter instances (of the four characters b , d , p , and q) were cut out manually from one or two suitable (representative and visually high-contrast) pages from each of the five codicological units. These characters were assumed to be distinctive of different hands. It was also assumed that instances of them would be relatively easy targets for the segmentation procedure. The selection of character instances and determination of bounding boxes around them consequently represent a manual input to the handwriting analysis. Only letters of the “normal” minuscule form were used for hand prediction. (Majuscule letters in different styles, used for emphasis and in titles, are also common in the Caroline books.) Furthermore, only the “half-uncial” d with straight upright ascender (like the d examples in table I) was used, excluding its “uncial” allograph. The letter dataset used for the present study contained a total of 436 items, 59–108 letter images for each hand, and 9–33 instances of each hand-character combination.

The data subset formed by every tenth page (position 1, 11, 21, etc., in the file name alphabetical order) of the codices was used during the development process for design and parameter tuning. The remaining pages were left unseen until the final evaluation. Table I shows the distribution of the pages into the letter source subset, the tuning dataset, and the evaluation dataset.

C. Preprocessing steps

The hand (and character) prediction process takes a manuscript image as input without requiring any human intervention. In the experiments reported here, image data were processed at a resolution of 8 pixels/mm, i.e. resized from the 13–21 pixels/mm of the original images. Different choices of processing resolution might generate different rounding effects.

The text lines of the manuscript pages are roughly parallel to the x-axis of the images, and the image processing preserves the coordinate systems given by the source images. Pixels belonging to the black (camera table) margins of each manuscript image are automatically identified as such (i.e. as large dark areas) and ignored in the further analysis targeting the parchment page area. Writing foreground and parchment background are separated (“binarized”) by means of a customized version of the Otsu [8] algorithm, which works quite well for manuscripts with as good contrast as the ones selected for the experiments reported here.

D. Letter segmentation

Connected components of foreground, presumed to form letters and letter sequences, are identified by means of connected component labeling. These components are split into pieces presumed to correspond to single letters. Vertical cuts are proposed, where the horizontal pixel projection profile is thinnest, looking from a window of a certain size (corresponding to 2 mm), but not thicker than a certain amount of foreground (corresponding to 1 mm). Any segment between these cuts is proposed as a letter candidate (to be classified), if its width is between 90% of the narrowest and 111% of the widest item in the letter dataset. The letter segmentation module essentially maps each manuscript image to a set of bounding boxes.

This somewhat naive method works because the next step in the classification process will discard bad proposals. Furthermore, with their clearly separated lines and the good contrast between ink and parchment, the Caroline manuscripts studied here are less challenging segmentation-wise than many other examples of writing.

E. Classification (hand, character) of letter instances

Each letter-size image is represented by a feature vector, which is computed with reference to the minimal bounding box enclosing the foreground pixels.

All features carry values in the interval $[0.0, 1.0]$, and can be understood as quotients. Experiments with different feature schemes indicated that this set of 24 features ($p, d, r_1, \dots, r_{20}, e_1, e_2$) would be among the optimal ones for the present purpose, characterizing the image in the following terms:

- Bounding box proportions: $p = w/(w + h)$, where w and h are the width and height, respectively, of the bounding box.
- Foreground density: $d = f/wh$, where f is the number of foreground pixels.
- Distribution of pixels as captured by a grid of 5×4 subrectangles over the bounding box, with the total number of foreground pixels as divisor: $r_n = f_n/f$, where f_n is the number of foreground pixels in the subrectangle n , $1 \leq n \leq 20$. See figure 1.
- Distribution of enclosed (“hole”) background pixels as captured by a grid of 2×1 subrectangles over the bounding box, with the subrectangle area as divisor: $e_n = 2h_n/wh$, where h_n is the number of enclosed background pixels (given the whole bounding box) in the subrectangle n , $1 \leq n \leq 2$. See figure 1.

As a step before the computation of these feature values, each separate connected element containing less than 20% of the foreground pixels in the letter bounding box is removed. The purpose of this is to eliminate, for instance, overhanging parts from adjoining letters, which might have been cut off by the bounding box, and other small foreground noise items.

If we compare these features with traditional scholarly criteria for hand identification, guided by the overview given

Fig. 1. The grid of 5×4 rectangles corresponding to the 20 features used to capture the distribution of foreground (ink) in the bounding boxes (enclosing letters) (left) and the grid of 2×1 rectangles used to capture the amount and location of background enclosed by foreground (right). The image is binarized as a step in the computation of these features.

TABLE II. AN EXAMPLE OF THE TOP 24 LETTER HITS FOR A MANUSCRIPT PAGE (COD. SANG. 186, P. 103). (CHARACTER CLASSIFICATION CORRECT IN ALL CASES, AND HAND IN 14 CASES.)

by Stokes [10], we can note that they in particular will capture aspects such as:

- Form: the morphology of the letters.
- Modulus: the proportions of the letters.
- Weight: the differences in thickness between different kinds of line.

Cosine similarity, defined by Jones and Furnas [7] as in equation 1, is used as the similarity metric (with $n = 24$, here). The features are given equal weight.

$$\text{similarity}(I, J) = \frac{\sum_{i=1}^n I_i \times J_i}{\sqrt{\sum_{i=1}^n I_i^2} \times \sqrt{\sum_{i=1}^n J_i^2}} \quad (1)$$

Given the feature model and the similarity measure, the hand (and character) prediction is performed by instance-based nearest neighbor classification. This means that it returns the prediction that an image is of the same type (hand and character) as the most similar instance in the letter dataset.

TABLE III. THE PERFORMANCE OF THE HAND CLASSIFIER ON TOPLISTED, AUTOMATICALLY EXTRACTED LETTERS, BY CODICOLOGICAL UNIT/HAND, GENERATED IN THE ANALYSIS OF THE UNSEEN SUBSET OF THE DATA.

Hand/unit	Total	Correct	Incorrect	Accuracy
Cod. Sang. 112	8273	7369	904	89.1%
Cod. Sang. 186	3711	2433	1278	65.6%
Cod. Sang. 557	6649	5132	1517	77.2%
Cod. Sang. 562	2291	1956	335	85.4%
Cod. Sang. 565	5683	5038	645	88.7%
All five	26607	21928	4679	82.4%

TABLE IV. THE PERFORMANCE OF THE HAND CLASSIFIER ON TOPLISTED, AUTOMATICALLY EXTRACTED LETTERS, LIKE TABLE III, BUT BY PREDICTED CHARACTER.

Character (pred.)	Total	Correct (hand)	Incorrect	Accuracy
<i>b</i>	5586	4303	1478	77.0%
<i>d</i>	8888	7947	941	89.4%
<i>p</i>	8390	6912	1478	82.4%
<i>q</i>	3743	2766	977	73.9%
All four	26607	21928	4679	82.4%

F. Predicting the hand of a manuscript page

As mentioned above, the system uses only the characters *b*, *d*, *p*, and *q* of the “normal” minuscule form for hand prediction. The letter recognition process generates, for each manuscript page, a set of letter predictions (character, hand, bounding box) ranked by the similarity score. See table II for an example. Only letter predictions whose score is 0.985 or higher are considered. The hand that receives the largest number of votes from the top 29 letter predictions (or all of them if their number is smaller than that) is returned as the hand prediction for the page. If a toplist gives the same largest number of votes to several hands voting based on a one item shorter toplist will decide instead. This procedure will consequently produce a hand prediction for a manuscript page if at least one letter is recognized.

V. EVALUATION

A. Performance of the classifier on the letter dataset

Different letter classification models were evaluated during the development stage using the letter dataset. The accuracy of the letter classifier (based on the feature scheme and the similarity metric) with respect to hand and character prediction can be computed by applying it in a leave one out manner, i.e. by predicting the type of each letter instance by comparing it to every other letter in the letter dataset (containing 436 items). The letter image classifier decided on (and presented here) yields an accuracy rate of 90.8% (396 correct) for hand prediction and 99.8% (435 correct) for character (*b*, *d*, *p*, or *q*) prediction, when evaluated in this way.

B. Performance of the classifier on automatically extracted letters

As can be expected, if we remember that the letter dataset contains manually selected and excerpted “good” character instances, the performance of the letter-level classifier for hand prediction on single, automatically extracted letters (see table II for examples), gives us lower accuracy scores. The performance on those letters extracted for the voting procedure

TABLE V. THE PERFORMANCE OF THE HAND PREDICTION SYSTEM APPLIED TO SINGLE PAGES OF THE FIVE CODICOLOGICAL UNITS/HANDS, EVALUATED ON THE UNSEEN SUBSET OF THE DATA.

Hand/unit	Total	Correct	Incorrect	Accuracy
Cod. Sang. 112	289	289	0	100.0%
Cod. Sang. 186	128	108	14 (as 562) + 6 (as 565)	84.3%
Cod. Sang. 557	241	240	1 (as 565)	99.6%
Cod. Sang. 562	79	79	0	100.0%
Cod. Sang. 565	196	196	0	100.0%
All five	933	912	21	97.7%

(i.e. those which have been classified by the nearest neighbor classifier with a similarity score of at least 0.985 and which are among the 29 top hits for a page) is shown in table III, with the scores computed separately for the five codicological units. We see some interesting differences, which will be discussed below. Table IV shows the accuracy scores for each predicted character in the corresponding way. As character location ground truth data were not available, an exact quantitative evaluation of the performance of the letter classifier for character prediction can not be provided. However, a manual inspection of the extracted letters, which were saved as images, suggests – some caution should be observed when it comes to reading letters deprived of their context – that its accuracy rate lies above 99%. Most errors seem to relate to overprediction of *d*, for instance, when *a*, *cl*, or *il* were the correct readings. Nevertheless, as table IV somewhat surprisingly shows, hand prediction gives a considerably higher accuracy score for instances of *d* than for the other characters. This could, if some speculation is excused, indicate that the execution of some characters is more sensitive to the personal writing practice of individual scribes.

C. Performance of the hand classifier for manuscript pages

The performance scores for the page-level hand classification system are based on its predictions for the “unseen” manuscript pages left as evaluation data. The results are summarized in table V. For three of the five codicological units the classifier gives a 100.0% accuracy rate. For Cod. Sang. 557 there is just one error. For the pages from Cod. Sang. 186 the performance is considerably worse, with an error rate around 15%. This unit exhibits a striking variation in the thickness of the letter’s lines (see e.g. p. 15), and this might explain why the classifier runs into considerably more difficulties there. The errors for Cod. Sang. 186 and 557 corresponds, as we could expect, to lower letter-level accuracy rates, as was seen in table III.

D. Using a smaller and balanced letter dataset

Given that the letter dataset (see section IV-B) has a composition reflecting an opportunistic compilation procedure, there is a possibility that the effects reported in the previous section are due to its containing an unbalanced number of instances for different hand-character combinations. In order to investigate this possibility, the evaluation procedure was repeated using a dataset containing exactly 9 instances of each hand-character combination, i.e. a subset of 180 letters of the original 436. Interestingly, this improves the performance, giving us 13 (10 as 562 and 3 as 565) misclassifications of pages from Cod. Sang. 186 (leading to an accuracy rate of

89.8%), which is a better than the 20 (14 as 562 and 6 as 565) we saw before (in table V, whose numbers otherwise remain as they are with the smaller letter dataset). However, the tendency is still pronounced, that the hand behind Cod. Sang. 186 is more difficult to classify, and that the classifier in the error cases is inclined to identify it with Cod. Sang. 562.

VI. CONCLUSION

The main purpose of the system described here is to predict which scribe, among a set of known scribes, has produced a manuscript page. The system does this by using a partial OCR technique for extracting letters. A nearest neighbor classifier based on a letter shape feature model and the cosine similarity measure performs both hand and character prediction for the extracted letters. The process is supervised by relying on a set of cut-out letters representing known characters and scribal hands. The hand prediction for whole pages is based on voting, using the letter-level hand proposals.

The page-level hand prediction method was evaluated on manuscript pages from five different hands and reached an accuracy rate above 99% for four of them and 84% (which was improved to 90% by reducing the size of the letter dataset) for the “hardest” one. Even if it remains to be seen how well a classifier based on these principles would perform on manuscript collections representing a larger number of hands, the present performance figures are promising. The remarkable accuracy rate level of 70%–90% with which the hand of single, automatically extracted letters can be predicted is a new kind of result.

The feature model used here is very simple, yet sensitive to both the individual variation among scribes executing the Caroline script and to features distinguishing its different characters. No learning process, apart from the computation of feature vectors, is involved in the classifier construction, since it is based on a nearest neighbor search procedure.

The Java implementation of the classifier requires, on an ordinary laptop, a few seconds to process a manuscript page. Most of that time is spent on the connected component labeling and letter segmentation, for which only a naive implementation has been used.

The present “literate” system represents a new approach to the issue of hand prediction. Its letter-based mode of operation makes it potentially useful for the purposes of digital palaeography. An analysis which processes handwriting as a linear sequence of linguistic forms (graphemes, characters) allows us to derive a justification for each hand prediction which would be comprehensible to a traditional palaeographer. Such an account would point to the pairwise similarities between specific letters, one in the manuscript under scrutiny, and one in a reference (letter source) manuscript. As a by-product, it automatically excerpts character instances which could be subjected to further palaeographical analysis, for instance, by tools of the kind proposed by Ciula [3] and Stokes [10].

Admittedly, several components of the present proposal would have a hard time coping with examples of handwriting which are more difficult than the well-preserved Caroline

codices studied here. Discolored parchment, faded ink, bleed-through, and similar problems are challenges to any binarization method, and they necessitate more sophisticated solutions than Otsu's algorithm. Furthermore, more cursive handwriting styles and manuscripts exhibiting overlapping between the lines would make letter extraction a more demanding business than it is for the typically well-spaced and clearly articulated Caroline minuscule style.

The main disadvantage of the present kind of approach, if we look at things from the perspective of the "illiterate" systems discussed in section II, is that it requires a manually produced dataset of letters. A set of 436 images (covering all combinations of five codicological units and four characters) like the one used here would take a couple of hours to compile. This suggests that it would be worthwhile to explore the possibility of combining the letter-based approach for hand classification with automatic extraction of reference letters. One way of doing this would be to use, for instance, the letters already collected, to extract letters from a new collection of manuscripts. A system based on the hand prediction method proposed here could then be used to predict which pages were most likely to have been written by the same hand. It could also be used to search for those manuscript pages which are script-wise most similar to a given page. The extraction-voting classification method could also be used with more arbitrary writing elements of the kind identified in the codebook-based methods mentioned in section II, without requiring that these elements correspond to characters.

ACKNOWLEDGMENTS

The author would like to thank Anders Brun, Fredrik Wahlberg, Lasse Mårtensson, Tomas Wilkinson, and Kalyan Ram Ayyalasomayajula for helpful discussions on the system reported here.

REFERENCES

- [1] A. A. Brink, J. Smit, M. L. Bulacu, and L. R. B. Schomaker, Writer identification using directional ink-trace width measurements, *Pattern Recognition* 45, 162–171, 2012.
- [2] M. Bulacu and L. Schomaker, Text-independent writer identification and verification using textural and allographic features, *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)* 29(4), 701–717, 2007.
- [3] A. Ciula, The Palaeographical Method under the Light of a Digital Approach, in F. Fischer, C. Fritze, and G. Vogeler (Eds.), *Kodikologie und Paläographie im digitalen Zeitalter – Codicology and Palaeography in the Digital Age*, Schriften des Instituts für Dokumentologie und Editorik, 2. BoD, Norderstedt, 219–234, 2009.
- [4] C. De Stefano, F. Fontanella, M. Maniaci, and A. Scotto di Freca, A method for scribe distinction in medieval manuscripts using page layout features, in G. Maino and G.L. Foresti (Eds.), *ICIAP 2011, Part I, LNCS 6978*, Springer-Verlag, Berlin, Heidelberg, 393–402, 2011.
- [5] A. Fischer, V. Frinken, A. Fornés, and H. Bunke, Transcription Alignment of Latin Manuscripts using Hidden Markov Models, *Proc. 1st Int. Workshop on Historical Document Imaging and Processing*, 29–36, 2011.
- [6] R. Jain and D. Doerman, Writer Identification Using an Alphabet of Contour Gradient Descriptors, *12th International Conference on Document Analysis and Recognition (ICDAR)*, 550–554, 2013.
- [7] W. P. Jones and G. W. Furnas, Pictures of Relevance: A Geometric Analysis of Similarity Measures, *Journal of the American Society for Information Science and Technology* 38(6), 420–442, 1987.
- [8] N. Otsu, A threshold selection method from gray-level histograms, *IEEE Transactions on Systems, Man and Cybernetics* 9(1), 62–66, 1979.
- [9] L. Schomaker, M. Bulacu, and K. Franke, Automatic writer identification using fragmented connected-component contours, in F. Kimura and H. Fujisawa (Eds.), *Ninth IWFHR*, Tokyo, 185–190, 2004.
- [10] P. Stokes, P., Computer-Aided Palaeography, Present and Future, in F. Fischer, C. Fritze, and G. Vogeler (Eds.), *Kodikologie und Paläographie im digitalen Zeitalter – Codicology and Palaeography in the Digital Age*, Schriften des Instituts für Dokumentologie und Editorik, 2. BoD, Norderstedt, 309–338, 2009.
- [11] Y. Tang, X. Wu, and W. Bu, Offline Text-independent Writer Identification Using Stroke Fragment and Contour Based Features, *2013 International Conference on Biometrics (ICB)*, 2013.