

The Effects of Character-Level Data Augmentation on Style-Based Dating of Historical Manuscripts

Lisa Koopmans^a, Maruf A. Dhali^b and Lambert Schomaker^c

Department of Artificial Intelligence, University of Groningen, The Netherlands

Keywords: Data Augmentation, Document Analysis, Historical Manuscript Dating, Self-Organizing Maps, Neural Networks, Support Vector Machines.

Abstract: Identifying the production dates of historical manuscripts is one of the main goals for paleographers when studying ancient documents. Automatized methods can provide paleographers with objective tools to estimate dates more accurately. Previously, statistical features have been used to date digitized historical manuscripts based on the hypothesis that handwriting styles change over periods. However, the sparse availability of such documents poses a challenge in obtaining robust systems. Hence, the research of this article explores the influence of data augmentation on the dating of historical manuscripts. Linear Support Vector Machines were trained with k-fold cross-validation on textural and grapheme-based features extracted from historical manuscripts of different collections, including the Medieval Paleographical Scale, early Aramaic manuscripts, and the Dead Sea Scrolls. Results show that training models with augmented data improve the performance of historical manuscripts dating by 1% - 3% in cumulative scores. Additionally, this indicates further enhancement possibilities by considering models specific to the features and the documents' scripts.

1 INTRODUCTION


Handwritten accounts, letters, and similar documents provide essential information about history. To understand such historical manuscripts' social and cultural contexts, paleographers seek to identify their script(s), author(s), location, and production date. Traditionally, paleographers study manuscripts by their writing materials, content, and handwriting styles. However, these methods require specific domain knowledge, are timely processes, and lead to subjective estimations. Additionally, repetitive physical handling leads to further degradation of valuable documents.


The digitization of historical manuscripts has contributed to their preservation and allowed for the development of automatized methods through machine learning. These tools are more objective than traditional methods and can aid paleographers in assessing their hypotheses. Historical manuscript dating, in particular, can benefit from this, as it can be required to resort to physical methods, which have limited reliability and can be destructive.


Dates of digitized historical manuscripts have been commonly predicted based on the hypothesis that handwriting styles change over a period (He et al., 2014). Thus, manuscripts could be dated by identifying common characteristics in handwriting specific to periods.

Due to the limited availability of historical manuscripts, research has mainly focused on statistical feature-extraction techniques. These statistical methods extract the handwriting style by capturing attributes such as curvature or slant or representing the general character shapes in the documents (Bulacu and Schomaker, 2007). However, for reliable results, manuscripts need a sufficient amount of handwriting to extract the handwriting styles.

Both traditional and automatized methods must deal with data sparsity and the degradation of ancient materials; new data can only be obtained by digitizing or discovering more manuscripts. A possible solution to this issue is data augmentation. Data augmentation is commonly used in machine learning to generate additional realistic training data from existing data to obtain more robust models. However, information on the handwriting styles is lost using standard techniques, such as rotating or mirroring the images. Character-level data augmentation could generate realistic samples simulating an author's variability

^a  <https://orcid.org/0000-0001-6556-2600>

^b  <https://orcid.org/0000-0002-7548-3858>

^c  <https://orcid.org/0000-0003-2351-930X>

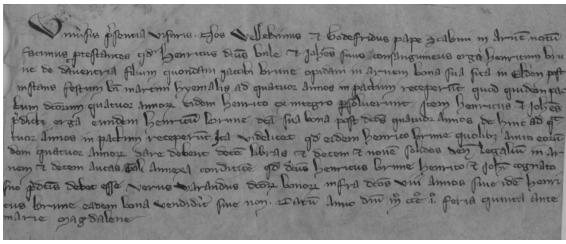


Figure 1: A document image from the Medieval Paleographical Scale (MPS) collection.

in handwriting.

Research on the style-based dating of digitized historical manuscripts using data augmentation techniques still needs to be done. Hence, the current research will explore the effects of character-level data augmentation on the style-based dating of digitized historical manuscripts. Manuscript images taken from the Medieval Paleographical Scale (MPS) collections, the Bodleian Libraries of the University of Oxford, the Khalili collections, and the Dead Sea Scrolls were augmented with an elastic rubber-sheet algorithm (Bulacu et al., 2009a). The first collection, MPS, has medieval charters produced between 1300 and 1550 CE in four cities: Arnhem, Leiden, Leuven, and Groningen. A number of early Aramaic, Aramaic, and Hebrew manuscripts were taken from the last three collections. Several statistical feature-extraction methods on the textural and character level were used to train linear Support Vector Machines (SVM) with only non-augmented images and with both non-augmented and augmented images.

2 RELATED WORKS

The main challenge in style-based dating is the selection of feature-extraction techniques. Each script has its own characteristics, which may not be represented well by every feature. Collections of historical manuscripts written in various languages and scripts have been digitized. For example, the Medieval Paleographical Scale (He et al., 2016d) and the Svenskt Diplomatariums huvudkartotek (SDHK) data sets¹ are written in Roman script, consisting of medieval Dutch and Swedish manuscripts respectively. Moreover, the early Aramaic and Dead Sea Scrolls collections (Shor et al., 2014) contain ancient texts in Hebrew, Aramaic, Greek, and Arabic, dating from the fifth century BCE (Before the Common Era) until the Crusader Period (12th–13th centuries CE).

Statistical feature-extraction methods are commonly divided into textural-based features that cap-

¹<https://sok.riksarkivet.se/SDHK>




Figure 2: An Early Aramaic (EA) manuscript from the Bodleian Libraries, University of Oxford (Pell. Aram. I).

ture textural information of the handwriting across an entire image and grapheme-based features that capture character-shape information. Graphemes extracted from a set of documents are used to train a clustering method. The cluster representations form a codebook, from which a probability distribution of grapheme usage is computed for each document to represent the handwriting styles.

A widely used textural feature is the 'Hinge' feature, which captures a handwriting sample's slant and curvature information. The features are extensions of the Hinge feature, which describes the joint probability distribution of two hinged edge fragments (Bulacu and Schomaker, 2007). In addition, Hinge is extended to i.e., co-occurrence features QuadHinge and CoHinge, which emphasize curvature and shape information respectively (He and Schomaker, 2017b). Other features, such as curvature-free and chain code features, have also been proposed (He and Schomaker, 2017c), (Siddiqi and Vincent, 2010).

Connected Component Contours (CO3) (Schomaker and Bulacu, 2004) is a grapheme-based feature that describes the shape of a fully connected contour fragment. As cursive handwriting has large connected contour fragments, the feature was extended to Fraglets (Bulacu and Schomaker, 2007), which parts the connected contours based on minima in the fragments. Moreover, k contour fragments (k CF) and k stroke fragments (k SF) features were proposed that partition CO3 in k contour and stroke fragments respectively (He et al., 2016b). Finally, Junclets (He et al., 2015) represents junctions in characters, which are constructed differently in varying writing styles.

Much research on historical manuscript dating has been done on the MPS data set, specifically by He et al. In (He et al., 2014), they predicted dates with a technique combining local and global Support Vector Regression, using Fraglets and Hinge features. They later extended this work, proposing new features such as k CF, k SF, and Junclets. In addition, they

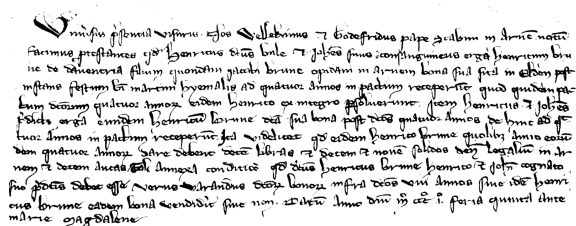


Figure 3: The binarized version of the image from Figure 1 with Otsu thresholding.

proposed the temporal pattern codebook (He et al., 2016a), which maintains temporal information lost in the commonly used Self-Organizing Map (SOM) (Kohonen, 1990) to train codebooks. Finally, various statistical feature-extraction methods were compared for historical manuscript dating in (He and Schomaker, 2017a).

While the MPS data set is relatively clean, it is not representative of many other historical manuscripts. In early works (Dhali et al., 2020), an initial framework was proposed for the style-based dating of the Dead Sea Scrolls. Unfortunately, the manuscripts from this collection are heavily degraded; many scrolls are fragmented, and ink traces have eroded due to aging. Additionally, the number of labeled manuscripts is small. Therefore, this collection poses a challenge for automatized dating of historical manuscripts.

Deep learning approaches have applied transfer learning, meaning pre-trained neural networks were fine-tuned using new data on a different task than initially trained for. This approach requires less data than standard deep learning methods, enabling its use for historical manuscript dating. For example, (Wahlberg et al., 2016) used the Google ImageNet-network and fine-tuned it using 11000 images from the SDHK collection. However, this is large for a data set of historical manuscripts. In (Hamid et al., 2019), a group of pre-trained neural networks was fine-tuned on the 3267 images from the MPS data set. The best-performing model was shown to outperform statistical methods.

While deep learning approaches show promising results, it is still relevant to consider statistical methods. To train a neural network, the manuscripts' images need to be partitioned into patches, possibly leading to loss of information. To solve this problem, (Hamid et al., 2019) ensured that each patch contained "3 to 4 lines of text with 1.5 to 2 words per line" to extract the handwriting style. While this was a solution for the MPS data set, it may not be for smaller and more degraded collections, such as the Dead Sea Scrolls. In contrast, statistical feature ex-

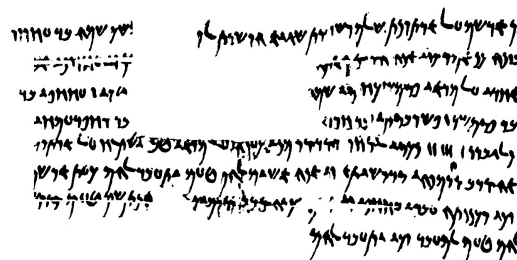


Figure 4: The binarized version of the image from Figure 2 using BiNet (Dhali et al., 2019).

traction does not require image resizing and considers the handwriting style over the entire image.

3 METHODS

This section will present the dating model along with data description, image processing, and feature extraction techniques.

3.1 Data

3.1.1 MPS

The current research uses the MPS data set (He et al., 2014), (He et al., 2016c), (He et al., 2016b), (He et al., 2016d). Non-text content, such as seals, supporting backgrounds, color calibrators, etc., have been removed. Consequently, this data set provides relatively clean images. However, some images have been degraded or still contain a small part of a seal or ribbon. The data set is publicly available via Zenodo².

The MPS data set contains 3267 images of charters collected from four cities signifying four corners of the medieval Dutch language area. Figure 1 shows an example image. Charters were commonly used to document legal or financial transactions or actions. Additionally, their production dates have been recorded. For these charters, usually parchment and sometimes paper was used.

The charters date from 1300 CE to 1550 CE. Due to the evolution of handwriting being slow and gradual, documents from 11 quarter century key years with a margin of \pm five years were included in the data set. Hence, the data set consists of images of charters from the medieval Dutch language area in the periods 1300 ± 5 , 1325 ± 5 , 1350 ± 5 , up to 1550 ± 5 . Table 1 contains the number of charters in each key year.

²<https://zenodo.org/record/1194357#.YrLU-OxBy3I>

Table 1: The number of samples over the key years of the MPS data set.

key year	1300	1325	1350	1375	1400	1425	1450	1475	1500	1525	1550
number of charters	106	164	199	386	311	323	501	423	372	241	241

3.1.2 Early Aramaic and Additional (EAA) Manuscripts

In addition to the MPS data set, 30 images from the early Aramaic, Aramaic, and Hebrew manuscripts were used. For ease of refereeing to this second dataset, EAA is used in the rest of the article, even though EAA contains Aramaic and Hebrew in addition to early Aramaic scripts. A list of the EAA images used in this study can be found in the appendix (see Table 5). For these selected manuscripts from the EAA dataset, the dates were directly inferred from dates or events recorded in the manuscripts (i.e., internally dated), and they are publicly available through the Bodleian Libraries, University of Oxford³, the Khalili collections⁴, and the Leon Levy Digital Library⁵. Their dates span from 456 BCE to 133 CE. An example image is shown in Figure 2. In addition, the data set contains several degraded manuscripts with missing ink traces or only two or three lines of text.

3.2 Preprocessing

3.2.1 Label Refinement

The set of images from the EAA collections did not contain sufficient samples for each year. Therefore, the samples were manually classified based on historical periods identified by historians⁶. The time periods and the corresponding number of samples are shown in Table 2.

The Persian Period contained two groups of samples spread apart for more than 30 years. Under the speculation that handwriting styles changed during this time, these samples were split into two periods: the Early and Late Persian Periods. These were not based on defined historical periods but on the samples' production years. Images from the upper bound of the year range in Table 2 were included in the classes. The manuscripts from the Roman Period were excluded as there were insufficient samples. The images were relabeled according to the median of their corresponding year ranges.

³<https://digital.bodleian.ox.ac.uk/>

⁴<https://www.khalilicollections.org/all-collections/aramaic-documents/>

⁵<https://www.deadseascrolls.org.il/>

⁶<https://www.deadseascrolls.org.il/learn-about-the-scrolls/>

3.2.2 Data Augmentation

To augment the data such that new samples simulate a realistic variability of an author's handwriting, the Imagemorph program (Bulacu et al., 2009b) was used. The program applies random elastic rubber-sheet transforms to the data through local non-uniform distortions, meaning that transformations occur on the components of characters. Consequently, the Imagemorph algorithm can generate a large number of unique samples. For the augmented data to be realistic, a smoothing radius of 8 and a displacement factor of 1 were used, measured in units of pixels. As images of the MPS data set required high memory, three augmented images were generated per image. Since the EAA data sets were small, 15 images were generated per image.

3.2.3 Binarization

To extract only the handwriting, the ink traces in the images were extracted through binarization. This resulted in images with a white background representing the writing surface, and a black foreground representing the ink of the handwriting. Otsu thresholding (Otsu, 1979) was used for binarizing the MPS images, as the MPS data set is relatively clean, and it has been successfully used in previous research with the data set (He et al., 2014), (He and Schomaker, 2017a), (He et al., 2016b). Otsu thresholding is an intensity-based thresholding technique where the separability between the resulting gray values (black and white) is maximized. Figure 3 shows Figure 1 after binarization.

The EAA images were more difficult to binarize using threshold-based techniques. So, for the EAA images, we used BiNet: a deep learning-based method designed specifically to binarize historical manuscripts (Dhali et al., 2019). Figure 4 shows Figure 2 after binarization.

3.3 Feature Extraction

The handwriting styles of manuscripts were described by five textural features and one grapheme-based feature. Since the MPS and the EAA data sets are written in different scripts, features were chosen that perform well across different scripts.

Table 2: Division of EAA manuscripts across historical time periods. Note that these dates may not exactly be the same as defined by historians.

Time period	Year range	Median year	Number of samples
Early Persian Period	540 BCE - 400 BCE	470 BCE	12
Late Persian Period	400 BCE - 330 BCE	365 BCE	11
Hellenistic Period	330 BCE - 65 BCE	198 BCE	5
Roman Period	65 BCE - 325 CE	195 CE	2

3.3.1 Textural Features

Textural-based feature-extraction methods contain statistical information on handwriting in a binarized image by considering its texture. Textural-based features capture handwriting attributes like slant, curvature, and the author's pen grip, represented in a probability distribution.

He et al. proposed the joint feature distribution (JFD) principle, describing how new, more robust features can be created (He and Schomaker, 2017a). They identified two groups of such features: the spatial joint feature distribution (JFD-S) and the attribute joint feature distribution (JFD-A). The JFD-S principle derives new features by combining the same feature at adjacent locations, describing a larger area. The JFD-A principle derives new features from different features at the same location and consequently captures various properties.

Hinge (Bulacu and Schomaker, 2007): is obtained by taking orientations α and β with $\alpha < \beta$ of two contour fragments attached at one pixel and computing their joint probability distribution. The Hinge feature captures the curvature and orientation in the handwriting. 23 angle bins were used for α and β .

CoHinge (He and Schomaker, 2017b): follows the JFD-S principle, combining two Hinge kernels at two different points x_i, x_j with a Manhattan distance l , and is described by:

$$\text{CoHinge}(x_i, x_j) = [\alpha_{x_i}, \beta_{x_i}, \alpha_{x_j}, \beta_{x_j}] \quad (1)$$

This shows that the CoHinge kernel over contour fragments can be quantized into a 4D histogram. The number of bins for each orientation α and β was set to 10.

QuadHinge (He and Schomaker, 2017b): follows the JFD-A principle, combining the Hinge kernel with the fragment curvature measurement $C(f_c)$. Although Hinge also captures curvature information, it focuses on the orientation due to the small lengths of the contour fragments or lengths of the hinge edges. The fragment curvature measurement is defined as:

$$C(F_c) = \frac{\sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2}}{s} \quad (2)$$

F_c is a contour fragment with length s on an ink trace with endpoints $(x_1, y_1), (x_2, y_2)$. In addition, the

QuadHinge feature is scale-invariant due to agglomerating the kernel with multiple scales. The QuadHinge kernel can then be described through the Hinge kernel and the fragment curvature measurement on contour fragments F_1, F_2 :

$$H(x_i, s) = [\alpha_{x_i}, \beta_{x_i}, C(F_1), C(F_2)] \quad (3)$$

The number of bins of the orientations was set to 12, and that for the curvature to 6, resulting in a dimensionality of 5184.

DeltaHinge (He and Schomaker, 2014): is a rotation-invariant feature generalizing the Hinge feature by computing the first derivative of the Hinge kernel over a sequence of pixels along a contour. Consequently, it captures the curvature information of the handwriting contours. The Delta-n-Hinge kernel is defined as:

$$\begin{cases} \Delta^n \alpha(x_i) = \frac{\Delta^{n-1} \alpha(x_i) - \Delta^{n-1} \alpha(x_i + \delta l)}{\delta l} \\ \Delta^n \beta(x_i) = \frac{\Delta^{n-1} \beta(x_i) - \Delta^{n-1} \beta(x_i + \delta l)}{\delta l} \end{cases} \quad (4)$$

Where n is the n th derivative of the Hinge kernel. When used for writer identification, performance decreased for $n > 1$, implying that the feature's ability to capture writing styles decreased. Hence, the current research used $n = 1$.

Triple Chain Code (TCC) (Siddiqi and Vincent, 2010): captures the curvature and orientation of the handwriting by combining chain codes at three different locations along a contour fragment. The chain code represents the direction of the next pixel, indicated by a number between 1 to 8. TCC is defined as:

$$\text{TCC}(x_i, x_{i+l}, x_{i+2l}) = [CC(x_i), CC(x_{i+l}), CC(x_{i+2l})] \quad (5)$$

Where $CC(x_i)$ is the chain code at location x_i , and Manhattan distance $l = 7$.

3.3.2 Grapheme-Based Features

Grapheme-based features are allograph-level features that partially or fully overlap with allographs in handwriting, described by a statistical distribution. The handwriting style is then represented by the probability distribution of the grapheme usage over a document, computed with a common codebook.

Junctlets (He et al., 2015): represents the crossing points, i.e., junctions, in handwriting. Junctions are

categorized into ‘L’, ‘T’, and ‘X’ junctions with 2, 3, and 4 branches, respectively. In different time periods, the angles between the branches, the number of branches, and the lengths of the branches can differ, making the feature appropriate for dating. Compared to other grapheme-based features, this feature does not need segmentation or line detection methods. A junction is represented as the normalized stroke-length distribution of a reference point in the ink over a set of $N = 120$ directions. The stroke lengths are computed with the Euclidean distance from a reference point in a direction until the edge of the ink. The feature is scale-invariant and captures the ink-width and stroke length.

3.3.3 Codebook

Previous research commonly used the Self-Organizing Map (SOM) (Kohonen, 1990) unsupervised clustering method to train the codebook (He and Schomaker, 2017a). By using this, however, temporal information in the input patterns is lost. The partially supervised Self-Organizing Time Map (SOTM) (Sarlin, 2013) maintains this information. In (He et al., 2016a), SOTM showed an improved performance for a grapheme-based feature compared to SOM. Hence, the codebook was trained with SOTM.

SOTM trains sub-codebooks D_t for each time period using the standard SOM (Kohonen, 1990), with handwriting patterns $\Omega(t)$ from key year $y(t)$. The key years for the MPS (in CE) and the EAA (in BCE) data sets were defined as $y(t) = \{1300, 1325, 1350, \dots, 1550\}$, and $y(t) = \{470, 365, 198\}$ respectively. The final codebook D , is composed of the sub-codebooks D_t : $D = \{D_1, D_2, \dots, D_n\}$, with n key years. To maintain the temporal information, the sub-codebooks are trained in ascending order. The initial sub-codebook D_1 is randomly initialized as no prior information exists in the data set. The succeeding sub-codebooks are initialized with D_{t-1} and then trained. Algorithm 1 shows the pseudo-code obtained from (He et al., 2016a).

To train the sub-codebooks, the Euclidean distance measure was used as it significantly decreased training times. Each sub-codebook was trained for 500 epochs to ensure sufficient training took place. The learning rate α^* decayed from $\alpha = 0.99$ following (6). The sub-codebooks were trained on a computer cluster⁷.

$$\alpha^* = \alpha \cdot \left(1 - \frac{\text{current epoch}}{\text{max epoch}}\right) \quad (6)$$

⁷<https://wiki.hpc.rug.nl/peregrine/start>

Algorithm 1: SOTM (He et al., 2016a).

```

 $t \leftarrow 1$ 
Randomly initialize  $D_t$ 
Train  $D_t$  using  $\Omega(t)$  by the standard SOM
while  $t \leq n$  do
   $t \leftarrow t + 1$ 
  Initialise  $D_t$  using  $D_{t-1}$ 
  Train  $D_t$  using  $\Omega(t)$  by the standard SOM
end while
Output  $D = \{D_1, D_2, \dots, D_t, \dots, D_n\}$ 

```

A historical manuscript’s feature vector was obtained by mapping its extracted graphemes to their most similar elements in the trained codebook, computed via the Euclidean distance, and forming a histogram. Finally, the normalized histogram formed the feature vector.

3.4 Post-Processing

The feature vectors of all features were small decimal numbers, varying between 10^{-2} and 10^{-6} . To emphasize the differences between the feature vectors of a type of feature, the feature vectors were normalized between 0 and 1 based on the range of a feature’s feature vectors. A feature vector f is scaled according to the following equations:

$$f_{std} = \frac{f - \min(f)}{\max(f) - \min(f)} \quad (7)$$

$$f_{scaled} = f_{std} \cdot (\max - \min) + \min \quad (8)$$

Here, \max and \min are the maximum and minimum values over the whole set of feature vectors of a certain feature, while $\max(f)$ and $\min(f)$ are the maximum and minimum values of the feature vector f (Pedregosa et al., 2011).

3.5 Dating

3.5.1 Model

Historical manuscript dating can be regarded as a classification or a regression problem. As the MPS data set was divided into 11 classes (or key years) with clear borders, and the EAA data set was partitioned into classes, it was regarded as a classification problem. Following previous research on the MPS data set (He and Schomaker, 2017a), linear Support Vector Machines (SVM) were used for date prediction with a one-versus-all strategy.

3.5.2 Measures

The Mean Absolute Error (MAE) and the Cumulative Score (CS) are two commonly used metrics to evalu-




Figure 5: MAE over sub-codebook size on non-augmented MPS data from 10-fold cross-validation.

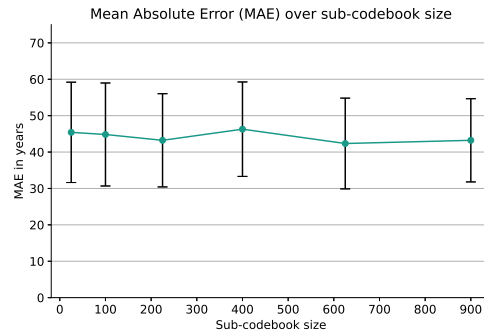


Figure 7: MAE over sub-codebook size on non-augmented EAA data from 4-fold cross-validation.

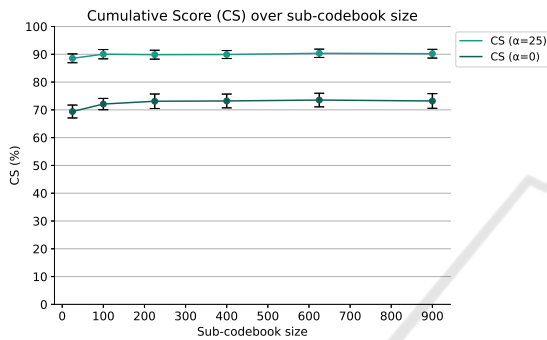


Figure 6: CS with $\alpha = 25$ and $\alpha = 0$ years over sub-codebook size on non-augmented MPS data from 10-fold cross-validation.

ate model performance for historical manuscript dating. The MAE is defined as follows:

$$MAE = \frac{\sum_{i=0}^{N-1} |y_i - \bar{y}_i|}{N} \quad (9)$$

Here, y_i is a query document’s ground truth, and \bar{y}_i is its estimated year. N is the number of test documents. The CS is defined in (Geng et al., 2007) as

$$CS = \frac{N_{e \leq \alpha}}{N} \cdot 100\% \quad (10)$$

The CS describes the percentage of test images that are predicted with an absolute error e no higher than a number of years α . At $\alpha = 0$ years, the CS is equal to the accuracy.

For both the MPS and the EAA data sets, CS with $\alpha = 0$ years was used. Since paleographers generally consider an absolute error of 25 years acceptable, and the MPS set has key years spread apart by 25 years, CS with $\alpha = 25$ years was also used for this data set.

3.5.3 Experiments

The MPS images were randomly split into a test and training set, containing 10% and 90% of the data, respectively. The EAA images were split into a test set

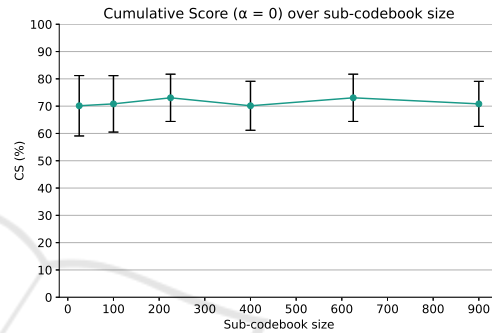


Figure 8: CS with $\alpha = 0$ years over sub-codebook size on non-augmented EAA data from 4-fold cross-validation.

of 5 images and a training set of 23 images. 2 samples were included from classes 470 and 365 BCE each. As the class 198 BCE contained only five images, one image from this class was considered in the test set. The images were sorted based on their labels, and the first images of each class were selected for testing.

The models were tuned with stratified k-fold cross-validation for both data sets, as they were imbalanced. For the MPS data set, $k = 10$. Since the training set of the EAA data set contained only four images from 198 BCE, $k = 4$ for this set. To prevent a randomized split in each iteration of the k-fold cross-validation from affecting the selection of hyper-parameters, hyper-parameters were selected using the mean results of stratified k-fold cross-validation across six random seeds, ranging from 0 to 250 with steps of 50. The set of values considered for the hyper-parameters were 2^n , $n = -7, -6, -5, \dots, 10$. During the process, the augmented images of those in the validation and test sets were excluded from the training sets.

Models were trained in two conditions. In the non-augmented condition only non-augmented images were used, and in the augmented condition both augmented and non-augmented images were used for training.

Codebook. Different sub-codebook sizes can result in different model performances. Hence, various sub-codebook sizes were tested to obtain the size for the Junclets feature. A codebook’s size is its total number of nodes, i.e., $n_{columns} \cdot n_{rows}$. The full codebook D is the concatenation of the sub-codebooks D_i , and thus its size will be $size_{D_i} \cdot n_{classes}$. The set of sub-codebook sizes $s = \{25, 100, 225, 400, 625, 900\}$ with $n_{columns} = n_{rows}$ were considered. These conditions were the same for the MPS and the EAA images. Since different codebook sizes result in different features, the sub-codebook sizes were determined based on the validation results of models trained on only non-augmented images.

The code used for the experiments and the SOTM is publicly available⁸.

4 RESULTS

To explore the effects of data augmentation on the style-based dating of historical manuscripts, five textural features and one grapheme-based feature were used. Linear SVMs were trained using only non-augmented data in the ‘non-augmented’ condition, and using both augmented and non-augmented data in the ‘augmented’ condition. The models were tuned with stratified 10-fold (MPS) and 4-fold (EA) cross-validation and tested on a hold-out set containing only non-augmented data. The test set of the MPS data set contained 10% of the data, and that of the EAA dataset contained 17.8% (5 images) of the data.

The models were evaluated with the MAE and CS with $\alpha = 0$ years (i.e. accuracy). In addition, the MPS data set was also evaluated with CS with $\alpha = 25$ years.

4.1 Sub-codebook Size

To investigate Junclets, first, an optimal sub-codebook size needed to be selected. Results of k-fold cross-validation for sub-codebook sizes 25, 100, 225, 400, 625, and 900 were evaluated on non-augmented data.

Figures 5 and 6 show the MAE and CS for the MPS data set over sub-codebook size, respectively. The MAE shows a minimum at the sub-codebook size of 625. Moreover, CS with $\alpha = 25$ and $\alpha = 0$ years show a maximum at sub-codebook size 625. Therefore, Junclets features were obtained with sub-codebooks of size 625 on the MPS data.

Figure 7 displays the MAE over the sub-codebook size on validation results for the EAA data. The MAE

decreases until the sub-codebook size is 225, after which it fluctuates. This is reflected in the CS with $\alpha = 0$ years (Figure 8), which displays an increase until size 225, after which it fluctuates. In addition, the standard deviations for the MAE and CS ($\alpha = 0$) appear the smallest here. Hence, a sub-codebook size of 225 was chosen for the EAA data.

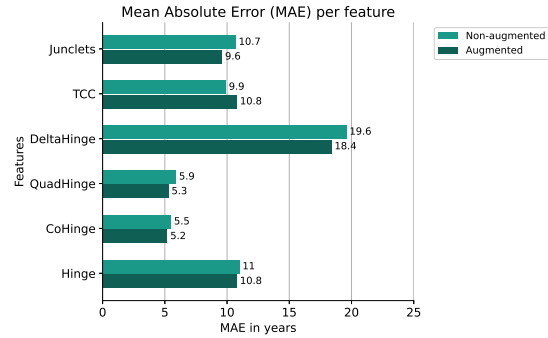


Figure 9: MAE on MPS (unseen) test data across non-augmented and augmented conditions.




Figure 10: CS with $\alpha = 25$ years on MPS (unseen) test data across non-augmented and augmented conditions.




Figure 11: CS with $\alpha = 0$ years on MPS (unseen) test data across non-augmented and augmented conditions.

⁸<https://github.com/Lisa-dk/Bachelor-s-thesis.git>

Table 3: k-Fold cross-validation results on the MPS data set.

Feature	Non-augmented MAE	Augmented MAE	Non-augmented CS ($\alpha=25$)	Augmented CS ($\alpha=25$)	Non-augmented CS ($\alpha=0$)	Augmented CS ($\alpha=0$)
Junclets	10.93 ± 1.31	9.15 ± 1.29	90.35 ± 1.49	92.39 ± 1.51	73.53 ± 2.45	77.56 ± 2.31
TCC	9.47 ± 1.08	8.95 ± 1.17	91.37 ± 1.39	92.16 ± 1.47	77.00 ± 1.85	77.98 ± 2.10
DeltaHinge	20.08 ± 1.88	18.35 ± 1.55	81.59 ± 1.96	83.00 ± 1.78	61.60 ± 2.25	63.91 ± 2.10
QuadHinge	5.76 ± 0.97	5.74 ± 0.97	95.38 ± 1.16	95.44 ± 1.17	84.65 ± 1.89	84.53 ± 1.94
CoHinge	6.81 ± 0.96	6.48 ± 0.88	94.32 ± 1.23	94.59 ± 1.17	82.13 ± 1.93	82.64 ± 1.95
Hinge	11.55 ± 1.44	11.28 ± 1.38	89.42 ± 1.74	89.36 ± 1.76	73.60 ± 2.55	73.76 ± 2.52

Table 4: k-Fold cross-validation results on the EAA data set.

Feature	Non-augmented MAE	Augmented MAE	Non-augmented CS ($\alpha=0$)	Augmented CS ($\alpha=0$)
Junclets	43.22 ± 12.80	40.40 ± 20.78	73.05 ± 8.66	70.28 ± 13.68
TCC	47.26 ± 12.52	57.92 ± 18.77	71.94 ± 10.59	65.42 ± 12.53
DeltaHinge	46.72 ± 9.13	45.54 ± 22.79	65.83 ± 8.20	75.55 ± 12.80
QuadHinge	38.97 ± 13.53	29.92 ± 15.72	76.67 ± 8.10	82.08 ± 9.41
CoHinge	48.18 ± 8.74	38.95 ± 20.93	64.44 ± 7.97	75.28 ± 12.69
Hinge	33.84 ± 13.43	26.17 ± 20.63	79.86 ± 7.07	84.44 ± 11.28

4.2 Augmentation

4.2.1 MPS

Figure 9 shows the MAE for each feature across the augmented and non-augmented conditions. The MAE for TCC increased in the augmented condition compared to the non-augmented condition. All other features displayed a decrease in the augmented condition.

Figure 10 shows the CS with $\alpha = 25$ years for both non-augmented and augmented conditions. An increase occurred in the augmented condition compared to the non-augmented condition for all features, except for TCC and Hinge, which display a decrease. Additionally, Junclets did not change in performance across conditions.

As displayed in Figure 11, all features showed an increase in CS with $\alpha = 0$ years in the augmented condition compared to the non-augmented condition with the exception of DeltaHinge. This feature showed no change in performance on test data.

These results denote an overall increase in performance for all features, with the exception of TCC. However, the changes in performances are small, which is reflected in the validation results displayed in Table 3, where changes between the non-augmented and augmented conditions are insignificant. This is indicated by means of the measures in augmented conditions falling within the ranges denoted by the standard deviations of the non-augmented conditions.

4.2.2 EAA Collections

Figures 12 and 13 show the MAE and CS with $\alpha = 0$ years across all features for the EAA data set. Performance increased for Junclets in the augmented condition compared to the non-augmented condition, in-




Figure 12: MAE on EAA (unseen) test data across non-augmented and augmented conditions.

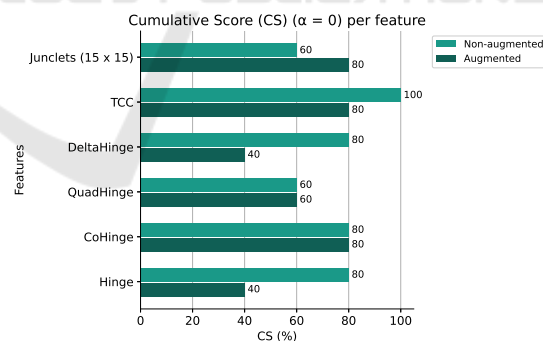


Figure 13: CS with $\alpha = 0$ years on EAA (unseen) test data across non-augmented and augmented conditions.

dicated by the decrease in MAE and increase in accuracy. QuadHinge also showed an increase in performance as the MAE decreased in the augmented condition. A decrease in performance for TCC, DeltaHinge, and Hinge features is denoted by an increase in MAE and a reduction in accuracy. CoHinge displayed no change across conditions.

These results are not reflected in the validation results (Table 4), where Junclets and TCC displayed

a decrease in performance with a reduction in mean MAE and an increase in mean accuracy in the augmented condition compared to the non-augmented condition. DeltaHinge, QuadHinge, CoHinge, and Hinge, however, displayed the opposite. Additionally, standard deviations increased significantly in the augmented condition compared to the non-augmented condition.

4.2.3 Significance

A statistical test (ANOVA, (Cuevas et al., 2004)) was performed to see if the results showed significant improvement. For the MPS data, the results from Junclets feature were statistically significant for both MAE and CS, with p-values much smaller than 0.005. However, for the EAA data, it did not show any significance for any of the feature extraction techniques.

5 DISCUSSION

The current study explores the effects of character-level data augmentation on the style-based dating of historical manuscripts using images from the MPS and EAA collections. Images were augmented with the Imagemorph program (Bulacu et al., 2009b) and then binarized. Linear SVMs were trained on five textural features and one grapheme-based feature. The grapheme-based feature Junclets was obtained by mapping extracted junction representations to a codebook trained with SOTM (Sarlin, 2013). Experiments were conducted to determine the sub-codebook sizes. SVMs were trained in ‘non-augmented’ and ‘augmented’ conditions where only non-augmented images and both non-augmented and augmented images were used, respectively. Models were evaluated through the MAE and CS with α -values of 0 and 25 years.

5.1 Key Findings

5.1.1 MPS

Test results showed that linear SVMs trained on MPS data in the augmented condition displayed an overall increased performance compared to the non-augmented condition for all features except TCC. TCC showed a decrease in performance. However, these increases and decreases were small, and changes in validation results were insignificant, with the ranges of the standard deviations and means overlapping across conditions.

The MPS images require much computer memory and, consequently, long running times to acquire

the features and models. Specifically, obtaining the Junclets features required several days. Hence, only three augmented images per MPS image were generated. Were more images generated, results might have shown a clearer picture of the influence of data augmentation on historical manuscript dating.

Another possible explanation for the small changes in performance shown by the MPS data set results is that MPS images were augmented before binarization. The Imagemorph program applies a Gaussian filter over local transformations. Consequently, if it is applied before binarization, the background’s influence leads to less severe distortions than if it is applied after binarization. Although the distortions were noticeable, they might have been too light to produce samples with natural within-writer variability. Whether this significantly affected the results is uncertain and should be considered in the future.

5.1.2 EAA Collections

Models trained on the EAA data set showed increased performance in the augmented condition compared to the non-augmented condition for Junclets and QuadHinge on test data. On the other hand, models for TCC, DeltaHinge, and Hinge showed a decreased performance in the augmented condition, and CoHinge showed no change in performance on test data. However, this is not reflected in the validation results (Table 4). Instead, validation results show a decrease in performance in the augmented condition for Junclets and TCC compared to the non-augmented condition, and an increase in performance for the remaining features.

The results of the EAA data could be explained by the increase in standard deviations across all features for models trained on both augmented and non-augmented data compared to models trained on only non-augmented data. This increase indicates that models were less robust to new data in the augmented condition, which may have led to diverging test results. Additionally, the differences between test results and validation results within the conditions, e.g., QuadHinge, indicate overfitting. This likely follows from the small size of the data set.

A possible reason why models trained with the EAA data set were less robust in the augmented condition is that linear SVMs were inappropriate for the data. While they previously worked well for the Roman script on the MPS data set, temporal information in the features extracted from EAA may follow non-linear patterns. Data augmentation could have emphasized these non-linear patterns, making linear models too rigid.

5.2 Future Research

Scripts have different characteristics, possibly resulting in differing distributions of extracted features. Likewise, individual features capture varying attributes of handwriting. Therefore, temporal information on handwriting styles might follow different trends across various features. While linear SVMs performed well on the MPS data set for the features used in the current research, these potential differences in distributions were not considered. This could lead to a decrease in performance for models trained on augmented data. Hence, other kernels should be studied to obtain optimal models for individual features and scripts.

One of the risks with historical manuscript dating is that the majority of the samples from a period, or a year, originate from one writer. Rather than learning to distinguish between characteristics in handwriting styles specific to a particular period or year, models would learn traits specific to writers for these years. Data augmented to simulate variability between writers within time periods might lead to more robust models than when data is augmented to simulate a realistic within-writer variability.

As mentioned in Section 2, deep learning approaches outperformed statistical approaches on the MPS data set. Considering this, it would be interesting to investigate whether data augmentation might positively affect historical manuscript dating on smaller and heavier degraded manuscripts, such as the EAA collections. Moreover, using the shape evolution of individual characters with grapheme-based statistical features might bypass the issue of limited data and loss of information due to the resizing of images.

ACKNOWLEDGEMENT

The study for this article collaborated with several research outcomes from the European Research Council (EU Horizon 2020) project: The Hands that Wrote the Bible: Digital Palaeography and Scribal Culture of the Dead Sea Scrolls (HandsandBible 640497), principal investigator: Mladen Popović. Furthermore, for the high-resolution, multi-spectral images of the Dead Sea Scrolls, we are grateful to the Israel Antiquities Authority (IAA), courtesy of the Leon Levy Dead Sea Scrolls Digital Library; photographer: Shai Halevi. Additionally, we express our gratitude to the Bodleian Libraries, University of Oxford, the Khalili collections, and the Staatliche Museen zu Berlin (photographer: Sandra Steib) for the early Aramaic im-

ages. We also thank Petros Samara for collecting the Medieval Paleographical Scale (MPS) dataset for the Dutch NWO project. Finally, we thank the Center for Information Technology of the University of Groningen for their support and for providing access to the Peregrine high-performance computing cluster.

REFERENCES

- Bulacu, M., Brink, A., Van Der Zant, T., and Schomaker, L. (2009a). Recognition of handwritten numerical fields in a large single-writer historical collection. In *2009 10th international conference on document analysis and recognition*, pages 808–812. IEEE.
- Bulacu, M., Brink, A., Zant, T., and Schomaker, L. (2009b). Recognition of handwritten numerical fields in a large single-writer historical collection. pages 808–812. Note: this is a peer-reviewed conference paper on an important international conference series, ICDAR; 2009 10th International Conference on Document Analysis and Recognition ; Conference date: 26-07-2009 Through 29-07-2009.
- Bulacu, M. and Schomaker, L. (2007). Text-independent writer identification and verification using textural and allographic features. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(4):701–717.
- Cuevas, A., Febrero, M., and Fraiman, R. (2004). An anova test for functional data. *Computational statistics & data analysis*, 47(1):111–122.
- Dhali, M., Wit, J., and Schomaker, L. (2019). Binet: Degraded-manuscript binarization in diverse document textures and layouts using deep encoder-decoder networks. *ArXiv*. 26 pages, 15 figures, 11 tables.
- Dhali, M. A., Jansen, C. N., de Wit, J. W., and Schomaker, L. (2020). Feature-extraction methods for historical manuscript dating based on writing style development. *Pattern Recognition Letters*, 131:413–420.
- Geng, X., Zhou, Z.-H., and Smith-Miles, K. (2007). Automatic age estimation based on facial aging patterns. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(12):2234–2240.
- Hamid, A., Bibi, M., Moetesum, M., and Siddiqi, I. (2019). Deep learning based approach for historical manuscript dating. In *2019 International Conference on Document Analysis and Recognition (ICDAR)*, pages 967–972.
- He, S., Samara, P., Burgers, J., and Schomaker, L. (2014). Towards style-based dating of historical documents. In *14th International Conference on Frontiers in Handwritten Recognition*. IEEE. 14th International Conference on Frontiers in Handwriting Recognition ; Conference date: 01-09-2014 Through 04-09-2014.
- He, S., Samara, P., Burgers, J., and Schomaker, L. (2016a). Historical manuscript dating based on temporal pattern codebook. *Computer Vision and Image Understanding*, 152:167–175.
- He, S., Samara, P., Burgers, J., and Schomaker, L. (2016b). Image-based historical manuscript dating using con-

- tour and stroke fragments. *Pattern Recognition*, 58:159–171.
- He, S., Samara, P., Burgers, J., and Schomaker, L. (2016c). A multiple-label guided clustering algorithm for historical document dating and localization. *IEEE Transactions on Image Processing*, 25(11):5252–5265.
- He, S. and Schomaker, L. (2014). Delta-n hinge: Rotation-invariant features for writer identification. In *22th International Conference on Pattern Recognition (ICPR)*, pages 2023–2028. IEEE (The Institute of Electrical and Electronics Engineers). 22nd International Conference on Pattern Recognition (ICPR) 2014 ; Conference date: 24-08-2014 Through 28-08-2014.
- He, S. and Schomaker, L. (2017a). Beyond OCR: Multi-faceted understanding of handwritten document characteristics. *Pattern Recognition*, 63:321–333.
- He, S. and Schomaker, L. (2017b). Co-occurrence features for writer identification. In *Proceedings of International Conference on Frontiers in Handwriting Recognition, ICFHR*, pages 78–83. Institute of Electrical and Electronics Engineers Inc.
- He, S. and Schomaker, L. (2017c). Writer identification using curvature-free features. *Pattern Recognition*, 63:451–464.
- He, S., Schomaker, L., Samara, P., and Burgers, J. (2016d). MPS Data set with images of medieval charters for handwriting-style based dating of manuscripts.
- He, S., Wiering, M., and Schomaker, L. (2015). Junction detection in handwritten documents and its application to writer identification. *Pattern Recognition*, 48(12):4036–4048.
- Kohonen, T. (1990). The self-organizing map. *Proceedings of the IEEE*, 78(9):1464–1480.
- Otsu, N. (1979). A threshold selection method from gray-level histograms. *IEEE Transactions on Systems, Man, and Cybernetics*, 9(1):62–66.
- Predregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Sarlin, P. (2013). Self-organizing time map: An abstraction of temporal multivariate patterns. *Neurocomputing*, 99:496–508.
- Schomaker, L. and Bulacu, M. (2004). Automatic writer identification using connected-component contours and edge-based features of uppercase western script. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(6):787–798.
- Shor, P., Manfredi, M., Bearman, G. H., Marengo, E., Boydston, K., and Christens-Barry, W. A. (2014). The leon levy dead sea scrolls digital library: The digitization project of the dead sea scrolls. *Journal of Eastern Mediterranean Archaeology and Heritage Studies*, 2(2):71–89.
- Siddiqi, I. and Vincent, N. (2010). Text independent writer recognition using redundant writing patterns with contour-based orientation and curvature features. *Pattern Recognition*, 43(11):3853–3865.
- Wahlberg, F., Wilkinson, T., and Brun, A. (2016). Historical manuscript production date estimation using deep convolutional neural networks. In *2016 15th International Conference on Frontiers in Handwriting Recognition (ICFHR)*, pages 205–210.

APPENDIX

Table 5: The list of EAA images used in this research.

A6_11R	A6.8	NS_A1r
A6_12R	B3.1	NS_A2r
A6_13R	IA01	NS_A4r
A6_14	IA03	NS_A5r
A6_15	IA04	NS_A6r
A6_16	IA06	NS_C1r
A6_3	IA17	NS_C4r
A6_4	IA21	WDSP1_1
A6_5	Mur24_1	WDSP2
A6_7	Mur24_2	Maresha Ostracon