

Decipherment of Historical Manuscript Images

Xusen Yin & Nada Aldarrab
Information Sciences Institute
University of Southern California
Marina del Rey, California, USA
{xusenyin,aldarrab}@isi.edu

Beáta Megyesi
Department of Linguistics and Philology
Uppsala University
Uppsala, Sweden
Beata.Megyesi@lingfil.uu.se

Kevin Knight
DiDi Labs
Marina del Rey, California, USA
kevinknight@didiglobal.com

Abstract—European libraries and archives are filled with enciphered manuscripts from the early modern period. These include military and diplomatic correspondence, records of secret societies, private letters, and so on. Although they are enciphered with classical cryptographic algorithms, their contents are unavailable to working historians. We therefore attack the problem of automatically converting cipher manuscript images into plaintext. We develop unsupervised models for character segmentation, character-image clustering, and decipherment of cluster sequences. We experiment with both pipelined and joint models, and we give empirical results for multiple ciphers.

Keywords—decipherment; historical manuscripts; image segmentation; character recognition; unsupervised learning; zero-shot learning

I. INTRODUCTION

European libraries and book collections are filled with undeciphered manuscripts dating from ca 1400 to 1900. These are often of historical significance, but historians cannot read them. Over recent years, a large number of ciphers are being scanned, collected, and put online for experimentation [1]. Figure 1 shows examples of cipher manuscripts. These ciphers are all considerable low-resource datasets from aspects of using one-off alphabets and glyphs. From the fraction of manuscripts that have been deciphered, cipher systems include simple substitution, homophonic substitution (where there may be many ways to encipher a given plaintext letter), substitution-transposition, nomenclators (where symbols may stand for whole words), or a combination of those. Plaintext languages include Latin, English, French, German, Italian, Portuguese, Spanish, Swedish, and so on.

Manual decipherment requires three major steps (Fig. 2):

- **Segmentation.** First, we decide where each character begins and ends. Even though ciphers often employ novel one-off alphabets, human analysts are quite good at segmenting lines into individual characters. However, problems do arise. For example, in the Borg cipher (Figure 1a), should \mathfrak{B} be segmented as one character or two?
- **Transcription.** Next, we convert the written characters into editable text, suitable for automatic analysis, such as character frequency counting. As Figure 2 shows, this may involve inventing nicknames for characters

that cannot be typed (e.g., zzz for \mathfrak{Z}). A human analyst can do this quite accurately, though mistakes happen. For example, in the Copiale cipher, should \mathfrak{g} and \mathfrak{g} be transcribed the same way, or are they actually distinct ciphertext symbols?

- **Decipherment.** Finally, we guess a cipher key that (when applied to the transcription) yields sensible plaintext. This is the hardest step for a human analyst, requiring intuition, insight, and grunt work.

Segmentation and transcription can be performed either manually or by semi-automatically, with post-editing. Given the number of historical encrypted manuscripts, manual transcription is infeasible, because it is too time-consuming, expensive, and prone to errors [2].

We would like to automate *all* of these steps, delivering a camera-phone decipherment app that a historian could use directly in the field. Automation efforts to date, however, have focused primarily on the decipherment step. Therefore, the problem we attack in this paper is *automatic decipherment directly from scanned images*.¹

Existing optical-character recognition (OCR) techniques are challenged by cipher manuscripts. The vast bulk of modern handwritten OCR requires large supervised datasets [3]–[5], whereas ciphers often use one-off alphabets for which no transcribed data exists. Back before supervised datasets were available, early OCR research proposed unsupervised identification and clustering of characters [6]–[9]. This is the general approach we follow here. Also in the unsupervised realm, recent work on historical documents focuses on printed, typeset texts [10], [11]. Though these methods model various types of noise, including ink bleeds and wandering baselines, they expect general consistency in font and non-overlapping characters.

The novel contributions of our paper are:

- Automatic algorithms for character segmentation, character clustering, and decipherment for handwritten cipher manuscripts.
- Evaluation on image data from multiple ciphers, measuring accuracy of individual steps as well as end-to-

¹Our code is available at <https://github.com/yinxusen/decipherment-images>

Input cipher	
Segmentation	
Transcription	v x. zzz bar ih lam s. k sqp ki arr bar w n
Decipherment	_ g e s e t z _ b u ch s _ _

Figure 2: Steps in decipherment a cipher manuscript.

Cipher	# pages	# characters	alphabet size
Courier	1	653	22
Arial	1	653	22
Borg	3	1054	23
Copiale	10	6491	79

Table I: Statistics of cipher image datasets used in this paper.

end decipherment accuracy.

- Improved techniques for joint inference, merging transcription and decipherment.

II. DATA

We perform experiments on two historical manuscripts—Borg (Figure 1a) and Copiale (Figure 1b)—and two synthetic ciphers (see Table I). All images are black-and-white in PNG format.

Borg. This is a 408-page manuscript from the 17th century, automatically deciphered by [12]. Its plaintext language is Latin. A few pages of the Borg cipher contain plaintext Latin fragments, which we remove from the images. In our experiments, we choose three consecutive pages and trim margins of each page (872x1416 pixels on average).

Copiale. This is a 105-page cipher from the 18th century, deciphered by [13]. The plaintext of Copiale is German. Copiale uses eight nomenclators that map to whole plaintext words, so we remove them from cipher images. In our experiments, we use the first 10 pages (1160x1578 pixels on average).

Courier (synthetic). We encipher a 653-character English text (simple substitution, no space), print it with fixed-width Courier font, then scan it into a PNG file (1956x2388 pixels).

Arial (synthetic). We create a similar image using variable-width Arial font (1976x2680 pixels).

We now turn to automatic methods for segmenting, transcribing, and deciphering.

III. AUTOMATIC SEGMENTATION

We define the upper-left corner of a page image as its origin, and the upper boundary as x-axis, left boundary as y-axis.

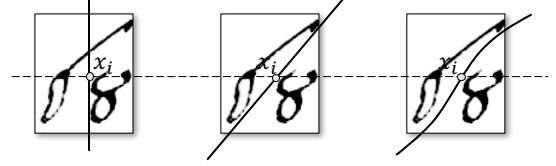


Figure 3: Vertical, slant, and cubic character-segmentation curves in solid lines cutting through the same point x_i . In this example, we choose the slant line as our cutting curve, and the number of intersected black pixels is $b_i = 0$.

- (T_1) We draw horizontal lines $y = c$, to split the manuscript into rows of characters;
- (T_2) We then draw vertical lines $x = c$, slant lines $y = bx$, or cubic curves $y = ax^3 + bx + c$ in each row of image to split characters.

Taking Task T_2 as an example, for an image row with m characters, we need to find cutting points c_1, c_2, \dots, c_m on the x-axis to draw curves.² Cutting points and curves we choose on the x-axis should meet the following (conflicting) requirements:

- (R_1) the number of cutting points should be m .
- (R_2) the widths of characters should be as similar as possible.
- (R_3) curves drawing across cutting points should intersect with as few black pixels as possible.

We use a generative model to formulate the requirements. At every point x_i on the x-axis of a row image, we use a set of pre-defined curves to cut through it (see Figure 3), and choose one with the minimum intersected black pixels b_i . When curves tie, we choose the simplest curve. If the row has W total pixel columns, our observed data is the sequence of black pixel numbers b_1, b_2, \dots, b_W collected from all cutting points on the x-axis x_1, x_2, \dots, x_W . Our goal is to choose m cutting points out of x_1, x_2, \dots, x_W to meet the requirements.

The generative story is first to choose the number of characters m of the row image according to a Gaussian distribution $m \sim P_{\phi_1, \sigma_1}(m)$. Then starting from the beginning of the row image, we generate the width of the next char-

²The last cutting point c_m is unnecessary, we use it to write clear equation.

acter from another Gaussian distribution $w_i \sim P_{\phi_2, \sigma_2}(w_i)$.³ Subsequently we use a geometric distribution p to generate the “observed” number of black pixels $b_i \sim P_p(b_i)$. We repeat for all m characters. We manually set parameters of the three distributions $\phi_1, \sigma_1, \sigma_2, p$ for each cipher.

We use Viterbi decoding to find the sequence c_1, c_2, \dots, c_m that best satisfies R_1, R_2 , and R_3 .

$$\arg \max_{m, c_1 \dots c_m} P_{\phi_1, \sigma_1}(m) \prod_{i=1}^m P_{\sigma_2}(w_i) P_p(b_i)$$

Figure 4 shows automatic segmentation results on snippets of manuscripts shown in Figure 1a and Figure 1b.

We manually create gold segmentations by cropping characters with a mouse. However, we evaluate our segmenter only as it contributes to accurate transcription.

IV. TRANSCRIPTION

For transcription, we:

- 1) Scale all character images to 105x105 pixels.
- 2) Convert each character image into a low-dimensional feature-vector representation.
- 3) Cluster feature vectors into similar groups.
- 4) Output a sequence of cluster IDs.

We implement two methods to transform a character image x into a fixed-length feature vector g with $g = F(x)$.

First, we propose a pairwise similarity matrix (SimMat). Given a sequence of character images $X = \{x_1, x_2, \dots, x_n\}$, SimMat computes the similarity between every pair of images as $s(x_i, x_j)$. Image x_i is then transformed into a n -dim vector according to the following equation,

$$F_{SimMat}(x_i; X) = [s(x_i, x_j)]_{j=1}^n$$

The similarity function $s(x_i, x_j)$ is the maximum of cross correlate matrix of two images. We use the `signal.correlate2d` function in Scipy package.

SimMat is a non-parametric feature extractor with a $O(n^2)$ time complexity, which makes it hard to apply to long ciphers. A sample cluster is shown in Figure 5.

Our second strategy exploits Omniglot, a labeled character-image dataset [14], containing 50 different alphabets, about 25 unique characters each, and 20 handwritten instances per.

- 1) We follow [15] to train a Siamese Neural Network (SNN) on pairs of Omniglot images. The SNN outputs 0 if two input images represent the same character.
- 2) We feed cipher character images into the SNN to extract feature representations.

The SNN architecture [15] is shown in Figure 6. The SNN has a visual feature extraction part $f(\cdot)$, which is

³Note that according to R_2 , given ϕ_1 , we have $\phi_2 = W/\phi_1$. So we can omit ϕ_2 in practice.

a convolution neural network, plus a single-layer neural classifier. Given an input image pair (x_1, x_2) , the output is

$$y = \text{sigmoid}(w * (f(x_1) - f(x_2)) + b).$$

We turn the classifier into a feature extractor by removing its classification part:

$$F_{SNN}(x_i) = w * f(x_i).$$

For clustering feature vectors, we use a standard Gaussian mixture model (GMM).⁴ Finally, as our transcription, we output a sequence of cluster IDs.

Evaluating Automatic Transcription. We manually create gold-standard transcriptions for all our ciphers. To judge the accuracy of our automatic transcription, we cannot simply use edit distance, because cluster IDs types do match human-chosen transcription symbols. Therefore, we map cluster ID types into human transcription symbols (many to one), transform the cluster IDs sequence accordingly, and then compute normalized edit distance. There are many possible mappings—we choose the one that results in the minimal edit distance. We have two methods to accomplish this, one based on integer-linear programming, and one based on expectation-maximization (EM). We call this **Normalized Edit Distance over Alignment** (NEDoA). Table II gives an example.

Table III compares NEDoA transcription accuracies under the SimMat and SNN feature extractors. SNN outperforms SimMat. The table also compares transcriptions from gold segmentation and automatic segmentation. Automatic segmentation on Borg degrades transcription accuracy.

V. DECIPHERMENT FROM TRANSCRIPTION

We can decipher from auto-transcription with the noisy channel model [16].

This generative model (Figure 7) first creates a sequence of plaintext characters $E = e_1 e_2 \dots e_n$ with a character n-gram model, then uses a channel model $P(C|E)$, transforms E into cipher text $C = c_1 c_2 \dots c_n$ character-by-character. The probability of our observation C is

$$P(C) = \sum_E P(E) P(C|E)$$

We can find the optimal channel model with the EM algorithm:

$$P(C|E) = \arg \max_{P(C|E)} P(C)$$

After we get the trained channel model, we use Viterbi decoding to find out the plaintext:

$$E = \arg \max_E P(E|C) \propto \arg \max_E P(E) P(C|E)$$

⁴We set GMM covariance type as diagonal, spherical, and fix-cov={1, 0.1, 0.01, 0.001} and choose the best one for each dataset.

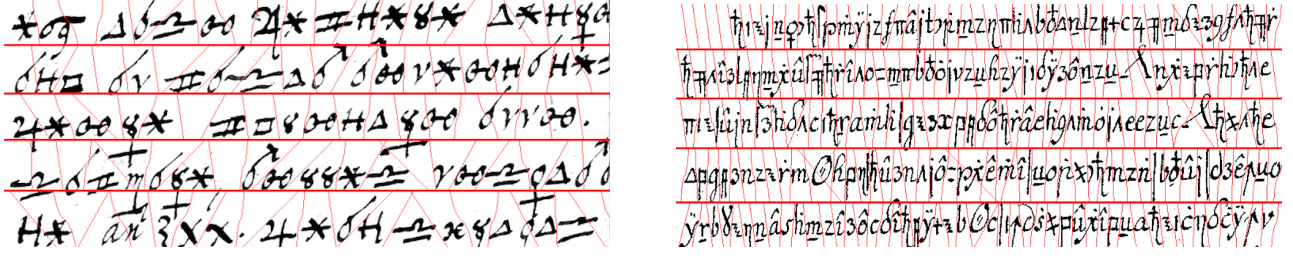


Figure 4: Results of automatic segmentation for part of Borg (left) and Copiale (right). We first use horizontal lines to segment rows, then use curves to segment characters in each row.

Gold transcription	z o d i a c k i l l e r
Automatic transcription	c0 c1 c2 c3 c3 c4 c5 c3 c6 c6 c7 c8
Best alphabet mapping	c0 - z, c1 - o, c2 - d, c3 - i, c4 - c, c5 - k, c6 - l, c7 - e, c8 - r
Edit distance after mapping	1 (a - c3)
NEDoA	$1 / \text{len}(\text{zodiackiller}) = 1 / 12 = 0.083$

Table II: Computing automatic transcription accuracy using the NEDoA metric. We map cluster ID types to gold transcription symbol types, make substitutions on the automatic transcription, then compute edit distance with gold. We search for the mapping that leads to the minimum normalized edit distance.

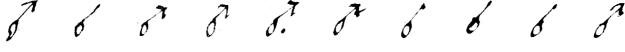


Figure 5: Part of a cluster from the Borg dataset using SimMat features. Here, two different cipher symbols are conflated into a single cluster.

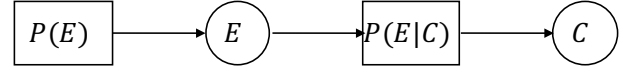


Figure 7: Noisy channel model for decipherment. $P(E)$ is a character language model and $P(E|C)$ is a channel model.

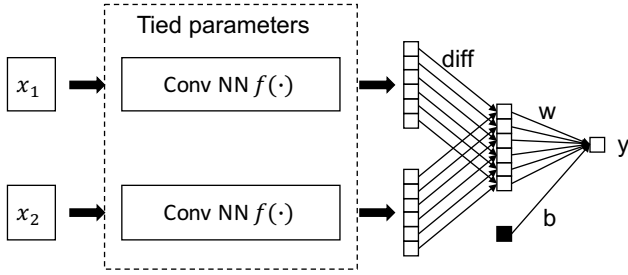


Figure 6: The architecture of Siamese Neural Network. x_1 and x_2 are a pair of character image inputs. Output $y = 0$ means x_1 and x_2 represent the same character type.

	Gold segmentation		Auto-segmentation
Feature extractor	SimMat	SNN	SNN
Courier	0.03	0.03	0.03
Arial	0.03	0.04	0.08
Borg	0.63	0.22	0.57
Copiale	0.87	0.37	0.44

Table III: Automatic transcription error rates (using NEDoA).

We call the first step **deciphering**, and the second step **decoding**. Combining segmentation, transcription, and decipherment, we create a pipeline to decipher from a scanned image, which we call **3-stage decipherment**.

Results for 3-stage Decipherment. We build pre-trained bigram character language models of English, Latin, and German. Since our cipher datasets have been deciphered, for each dataset we generate gold plaintext from gold transcription. We evaluate decipherment with normalized edit distance (NED).

Table IV compares decipherment error rates under gold segmentation and automatic segmentation. We also study decipherment under gold transcription. Our fully automatic system deciphers Copiale at 0.51 character error. While high, this is actually remarkable given that our transcription has 0.44 error. It seems that our fully-connected noisy channel decipherment model is able to overcome transcription mistakes by mapping the same cluster ID onto different plaintext symbols, depending on context.

Even so, we notice substantial degradation along the pipeline. Human analysts also revise transcriptions once decipherments are found.

VI. DECIPHERMENT FROM CHARACTER IMAGES

We propose **2-stage decipherment**, which models transcription and decipherment as a single integrated step.

	Auto-transcription		Gold transcription
	Gold seg	Auto-seg	
Courier	0.08	0.08	0.08
Arial	0.07	0.16	0.08
Borg	0.35	0.76	0.01
Copiale	0.46	0.51	0.20

Table IV: Error rates of 3-stage decipherment (NED) based on gold segmentation and auto-segmentation, compared with deciphering from gold transcription.

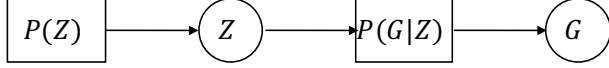


Figure 8: GMM. $P(Z)$ is a multinomial distribution over cluster IDs, and $P(G|Z)$ is a mixture of Gaussian distributions.

A. Language Model Constrained Gaussian Mixture Model (LM-GMM)

GMM. Given a sequence of feature vectors $G = \{g_1, g_2, \dots, g_n\}$ generated by the SNN feature extractor, GMM generates G by first using a multinomial distribution $P(Z)$ to choose cluster assignments, then uses the mixture of Gaussian distributions $P(G|Z)$ to generate feature vectors, as shown in Figure 8.

$$P(G) = \sum_Z P(Z)P(G|Z)$$

LM-GMM. Instead of using the multinomial distribution $P(Z)$ to choose clusters, LM-GMM uses the decipherment language model to choose appropriate character sequences. Since we do not have cipher language model, we use the noisy channel model as the **cipher language model** as shown in Figure 9.

$$P(G) = \sum_E P(E) \sum_C P(C|E)P(G|C)$$

Simplified LM-GMM. LM-GMM can be simplified for simple substitution ciphers. Since simple substitution ciphers use one-to-one and onto mappings between plaintext alphabet and cipher alphabet, the channel model $P(C|E)$ is not necessary. We imagine, for example, that Borg is written in Latin, but the author writes Latin characters strangely. The simplified model is

$$P(G) = \sum_E P(E)P(G|C)$$

B. LM-GMM Model Error

LM-GMM is a combination of discrete distributions (LM, channel) and a continuous one (GMM). Figure 11 (left) shows results of 5,000 random restarts (+). The x-axis

	Gold seg		Auto-seg	
	3-stage	2-stage	3-stage	2-stage
Courier	0.08	0.06	0.08	0.06
Arial	0.07	0.04	0.16	0.11
Borg	0.35	0.20	0.76	0.69
Copiale	0.46	0.41	0.51	0.50

Table V: Error rates of 2-stage decipherment (NED) on both gold segmentation and auto-segmentation, compared with 3-stage decipherment.

	3-stage decipherment	2-stage decipherment
Bigram	0.35	0.20
Trigram	0.24	0.16

Table VI: Decipherment error rates (NED) with bigram and trigram language models for Borg, using gold segmentation.

is the log-likelihood of observed feature vectors. The y-axis is decipherment NED. We also plot the gold model by generating the Gaussian Mixture $P(G|C)$ with gold plaintext, as the solid dot. Training from the gold model, we can reach the solid square.

The gold model does not receive the highest model score. This modeling error is caused by the strong GMM multiplying with the character language model—we tend to choose the result satisfying the GMM part. To fix this, we use $P(G) = \sum_E P(E)^3 P(G|C)$ to highlight the importance of the language model during deciphering phase, leading to a result shown in Figure 11 (right).

C. LM-GMM Search Error

Now we observe that even after many EM restarts, we cannot reach a model that scores as well as the gold model.

To fix this search problem, we randomly restart our EM training from a GMM model computed from plaintext from 3-stage training. We illustrate the initialization method on the Borg dataset (Figure 12). The initialization point comes from 3-stage decipherment (NED=0.35), which 2-stage decipherment improves to NED=0.20. This approaches the retrained gold model (NED=0.15), with only 50 restarts.

D. Evaluation of 2-stage Decipherment

Decipherment results are shown in Table V, which compares 2- and 3-stage decipherment under auto- and gold segmentation. All results are trained with bigram character language model. Results of 2-stage decipherment use the initialization method described above. From the results we can see that 2-stage decipherment outperforms 3-stage decipherment, especially for Borg and Copiale.

Instead of using bigram language model, we also use a trigram language model on Borg, as shown in Table VI. Both 3-stage and 2-stage decipherment get better results.

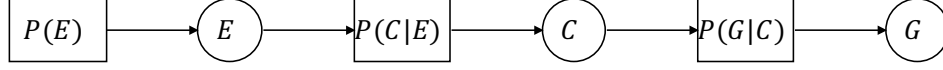


Figure 9: LM-GMM. $P(E)$ is character language model, $P(C|E)$ is a channel model, and $P(G|C)$ is a mixture of Gaussian distributions.

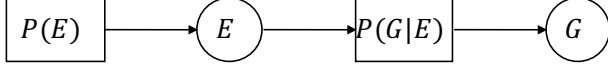


Figure 10: Simplified LM-GMM for simple substitution ciphers.

VII. CONCLUSION AND FUTURE WORK

In this paper, we build an end-to-end system to decipher from manuscript images. We show that the SNN feature extractor with a Gaussian mixture model can be good for unseen character clustering. We fix our EM search problem for LM-GMM by using a better initialization method.

Interesting future work can include 1-stage decipherment. Can we use our cipher language model to improve the image segmentation? Can we merge image segmentation into the whole EM training framework? How much benefit can we get?

Finally, to realize a fully-automatic camera-phone decipherment app, we need to lift several assumptions we made in the paper. These include knowing the plaintext language and cipher system [12], [17], [18], pre-processing images to remove margins and non-cipher text, knowing the cipher alphabet size, and cipher-specific setting of segmentation parameters.

ACKNOWLEDGMENTS

The authors thank all people who have deciphered these manuscripts to make this work possible. This work was sponsored by DARPA (grant HR0011-15-C-0115), by the Swedish Research Council (grant E0067801), by the USC Annenberg Graduate Fellowship, and by a gift from Google, Inc.

REFERENCES

- [1] E. Pettersson and B. Megyesi, “The histcorp collection of historical corpora and resources,” in *Proceedings of the Third Conference on Digital Humanities in the Nordic Countries*, 2018.
- [2] A. Fornés, B. Megyesi, and J. Mas, “Transcription of encoded manuscripts with image processing techniques,” in *Proceedings of Digital Humanities*, 2017.
- [3] R. Smith, “An overview of the tesseract OCR engine,” in *Proceedings of the Ninth International Conference on Document Analysis and Recognition*, 2007, pp. 629–633.
- [4] A. Kae, G. Huang, C. Doersch, and E. Learned-Miller, “Improving state-of-the-art OCR through high-precision document-specific modeling,” in *IEEE Conference on Computer Vision and Pattern Recognition*, June 2010, pp. 1935–1942.
- [5] V. Kluzner, A. Tzadok, D. Chevion, and E. Walach, “Hybrid approach to adaptive OCR for historical books,” in *IEEE International Conference on Document Analysis and Recognition*, Sept 2011, pp. 900–904.
- [6] G. Nagy, S. C. Seth, and K. Einspahr, “Decoding substitution ciphers by means of word matching with application to OCR,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 9, no. 5, pp. 710–715, May 1987. [Online]. Available: <http://dx.doi.org/10.1109/TPAMI.1987.4767969>
- [7] T. K. Ho and G. Nagy, “OCR with no shape training,” in *Proceedings of 15th International Conference on Pattern Recognition*, vol. 4, 2000, pp. 27–30.
- [8] T. M. Breuel, “Modeling the sample distribution for clustering OCR,” in *Proc. SPIE, Document Recognition and Retrieval VIII*, vol. 4307, 12 2000, pp. 193–200.
- [9] G. Huang, E. Learned-Miller, and A. McCallum, “Cryptogram decoding for ocr using numerization strings,” in *Ninth International Conference on Document Analysis and Recognition*, vol. 1, Sept 2007, pp. 208–212.
- [10] T. Berg-Kirkpatrick, G. Durrett, and D. Klein, “Unsupervised transcription of historical documents,” in *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics*, 2013, pp. 207–217.
- [11] T. Berg-Kirkpatrick and D. Klein, “Improved typesetting models for historical OCR,” in *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics*, 2014, pp. 118–123. [Online]. Available: <http://aclweb.org/anthology/P/P14/P14-2020.pdf>
- [12] N. Aldarrab, “Decipherment of historical manuscripts,” Master’s thesis, University of Southern California, Los Angeles, California, 2017.
- [13] K. Knight, B. Megyesi, and C. Schaefer, “The copiale cipher,” in *Proceedings of the 4th Workshop on Building and Using Comparable Corpora: Comparable Corpora and the Web*, 2011, pp. 2–9. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2024236.2024239>
- [14] B. M. Lake, R. Salakhutdinov, J. Gross, and J. B. Tenenbaum, “One shot learning of simple visual concepts,” in *Proceedings of the 33th Annual Meeting of the Cognitive Science Society*, 2011. [Online]. Available: <https://mindmodeling.org/cogsci2011/papers/0601/index.html>

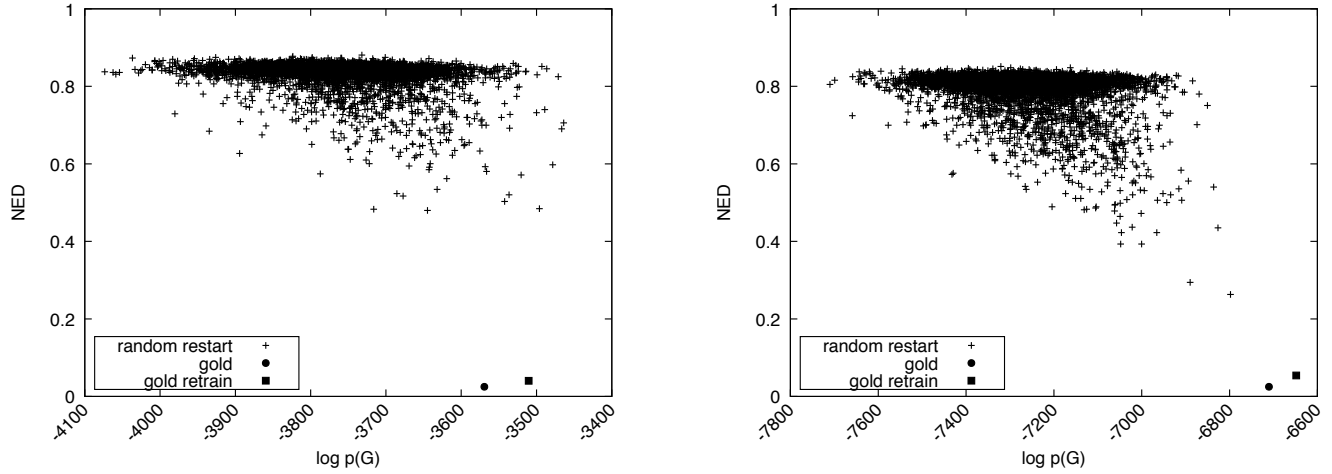


Figure 11: 5000 random restarts of training LM-GMM (left), and LM-GMM with cubic LM (right) on Courier dataset marked with plus (+). X-axis is the likelihood of the Courier dataset according to LM-GMM, y-axis is NED between deciphered text and gold plaintext. The solid dot is the gold model, and the solid square is the LM-GMM training result initialized from the gold model.

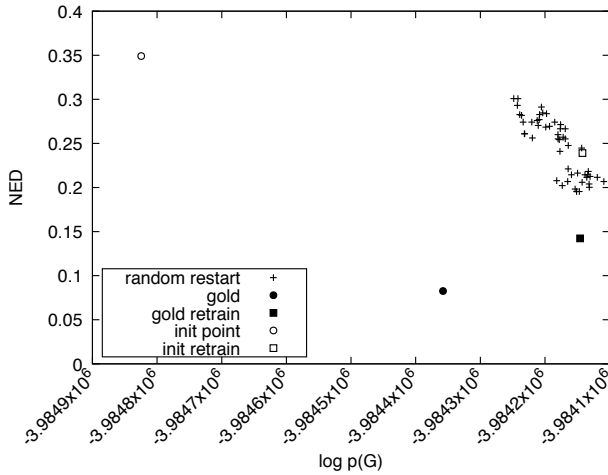


Figure 12: 50 restarts of LM-GMM on Borg dataset initialized from 3-stage decipherment with random noise (+). The x-axis is the log-likelihood of Borg vectors according to LM-GMM, and the y-axis is NED between deciphered text and gold plaintext. The hollow circle is the 3-stage decipherment result mapped onto this plot, and the hollow square is the LM-GMM training result initialized from the hollow circle. The solid dot is the gold model, and the solid square is the LM-GMM training result starting from the solid dot.

deciphering unknown scripts,” in *Proceedings of the ACL Workshop on Unsupervised Learning in Natural Language Processing*, 1999.

- [17] M. Nuhn and K. Knight, “Cipher type detection,” in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, 2014, pp. 1769–1773.
- [18] B. Hauer and G. Kondrak, “Decoding anagrammed texts written in an unknown language and script,” *Transactions of the Association for Computational Linguistics*, vol. 4, pp. 75–86, 2016.

[15] G. Koch, R. Zemel, and R. Salakhutdinov, “Siamese neural networks for one-shot image recognition,” in *Proceedings of the 32nd International Conference on Machine Learning*, 2015.

[16] K. Knight and K. Yamada, “A computational approach to