

Solving Historical Dictionary Codes with a Neural Language Model

Christopher Chu, Raphael Valenti, and Kevin Knight

DiDi Labs

4640 Admiralty Way

Marina del Rey, CA 90292

{chrischu, kevinknight}@didiglobal.com

Abstract

We solve difficult word-based substitution codes by constructing a decoding lattice and searching that lattice with a neural language model. We apply our method to a set of enciphered letters exchanged between US Army General James Wilkinson and agents of the Spanish Crown in the late 1700s and early 1800s, obtained from the US Library of Congress. We are able to decipher 75.1% of the cipher-word tokens correctly.

1 Introduction

Cryptography has been used since antiquity to encode important secrets. There are many unsolved ciphers of historical interest, residing in national libraries, private archives, and recent corpora collection projects (Megyesi et al., 2019; Pettersson and Megyesi, 2019). Solving classical ciphers with automatic methods is a needed step in analyzing these materials.

In this work, we are concerned with automatic algorithms for solving a historically-common type of *book code*, in which word tokens are systematically replaced with numerical codes. Encoding and decoding are done with reference to a dictionary possessed by both sender and recipient. While this type of code is common, automatic decipherment algorithms do not yet exist. The contributions of our work are:

- We develop a algorithm for solving dictionary-based substitution codes. The algorithm uses a known-plaintext attack (exploiting small samples of decoded material), a neural language model, and beam search.
- We apply our algorithm to decipher previously-unread messages exchanged between US Army General James Wilkinson and agents of the Spanish Crown in the late 1700s and early 1800s, obtaining 72.1% decipherment word accuracy.

	Table-based key	Book-based key
Cipher (character)	Caesar cipher Simple substitution Zodiac 408 Copiale cipher	Beale cipher
Code (word)	Rossignols' Grand Chiffre	Mexico-Nauen code Scovell code Wilkinson code (<i>this work</i>)

Figure 1: Simplified typology of substitution-based cryptosystems, with some examples. Ciphers involve character-level substitutions (e.g., $f \rightarrow q$), while codes involve word-level substitutions (e.g., *forest* \rightarrow *5731*).

2 Related Work

Figure 1 gives a simplified typology of classical, substitution-based cryptosystems.¹

Table-based Ciphers involve character-based substitutions. The substitution may take the form of a simple offset, as in the Caesar substitution system, e.g., $(a \rightarrow d)$, $(b \rightarrow e)$, $(c \rightarrow f)$, etc. The Caesar cipher can be easily solved by algorithm, since there are only 26 offsets to check. The algorithm need only be able to recognize which of the 26 candidate plaintexts form good English. Since 25 of the candidates will be gibberish, even the simplest language model will suffice.

A *simple substitution cipher* uses a substitution table built by randomly permuting the alphabet. Since there are $26! \approx 4 \cdot 10^{26}$ possible tables, algorithmic decipherment is more difficult. However, there are many successful algorithms, e.g., (Hart, 1994; Knight and Yamada, 1999; Hauer et al., 2014; Olson, 2007; Ravi and Knight, 2008; Corlett and Penn, 2010). Many of these systems

¹Our typology is geared toward explaining our contribution in the context of related systems. For a fuller picture of classical cryptology, the reader is directly to Kahn (1996) and Singh (2000). For example, we do not discuss here systems in which a substitution key evolves during the encoding process, such as the Vigenère cipher or the German Enigma machine.

search for substitution tables that result in candidate plaintexts that score well according to a character n-gram language model (Shannon, 1951), and they use search techniques like hill-climbing, expectation-maximization, beam search, and exact search. The main practical challenge is to decipher short messages. In a very long ciphertext, it is easy to “spot the q ” because it is always followed by the same cipher character, which we can immediately guess stands for plaintext u , and so forth.

More sophisticated ciphers use *homophonic* substitution, in which plaintext characters are replaced non-deterministically. By applying high nondeterminism to frequent characters, the cryptographer can flatten out ciphertext character frequencies. Homophonic ciphers occur frequently in historical collections. The Copiale cipher (Knight et al., 2011) is a well-known example from a German secret society in the 1700s. These ciphers can also be attacked successfully by algorithm. For example, the homophonic Zodiac 408 cipher can be solved with EM with restarts (Berg-Kirkpatrick and Klein, 2013), Bayesian sampling (Ravi and Knight, 2011), or beam search (Nuhn et al., 2014) (all with n-gram character language models). Kambhatla et al. (2018) employ a more powerful character-based neural language model to break short ciphers more accurately. In the present work, we use a word-based neural language model.

Book-based ciphers increase homophony, and also avoid physical substitution tables that can be stolen or prove incriminating. In a book-based cipher, sender and recipient verbally agree up front on an innocuous-looking shared document (the “book”), such as the US Declaration of Independence, or a specific edition of the novel *Moby Dick*. When enciphering a plaintext letter token like f , the sender selects a random letter f from the shared document—if it is the 712th character in the document, the plaintext f might be enciphered as 712. The next plaintext f might be enciphered differently.

Nuhn et al. (2014) solve one of the most well-known book ciphers, part two of the Beale Cipher (King, 1993). Surprisingly, they treat the cipher as a regular homophonic cipher, using the same beam-search algorithm as for the table-based Zodiac 408 cipher, together with an 8-gram character language model. One might imagine exploiting the fact that the book is itself written in En-

glish, so that if ciphertext unit 712 is known to be f , then ciphertext unit 713 is probably not h , as fh is unlikely to appear in the book. Nuhn et al. (2014)’s simple, effective algorithm ignores such constraints. Other methods have been proposed for attacking book ciphers, such as *crib dragging* Churchhouse (2002).

Codes, in contrast to ciphers, make substitutions at the whole-word level.² A large proportion of the encrypted material in Megyesi et al. (2019) consists of **table-based codes**. A famous example is Antoine Rossignol’s *Grand Chiffre*, used during the reign of Louis XIV. The sender and receiver each own copies of huge specially-prepared tables that map words onto numbers (e.g., *guerre* \rightarrow 825). If the enciphering tables are kept secret, this type of code is very hard to break. One might guess that the most frequent cipher token stands for the word *the*, but it quickly becomes challenging to decide which number means *practice* and which means *paragon*. Even so, Dou and Knight (2012) take on the task of automatically deciphering newswire encrypted with an arbitrary word-based substitution code, employing a slice-sampling Bayesian technique. Given a huge ciphertext of $\sim 50,000$ words, they can decipher $\sim 50\%$ of those tokens correctly. From one billion ciphertext tokens, they recover over 90% of the word tokens. However, this method is clearly inapplicable in the world of short-cipher correspondence.

In the present work, we consider **book-based codes**. Instead of using specially-prepared tables, the sender and receiver verbally agree to use an already-existing book as a key. Because it may be difficult to find a word like *paragon* in a novel like *Moby Dick*, the sender and receiver often agree on a shared pocket dictionary, which has nearly all the words. If *paragon* were the 10,439th word in the dictionary, the sender might encode it as 10439.

Such codes have been popular throughout history, employed for example by George Scovell during the Napoleonic Wars (Urban, 2002), and by John Jay during the US Revolutionary War (Blackwood, 2009). They were used as late as World War II, when German diplomats employed the *Langenscheidt’s Spanish-German Pocket Dictionary* as a key to communicate between the cities of Chapultepec, Mexico and Nauen, Germany (NSA,

²Commonly, a single historical system will mix letter substitutions and word substitutions. Such a system is called a *nomenclator*.

	Word tokens	Word types
Deciphered letters	approx. 800	326
Evaluation set	483	226
Test set	341	192

Table 1: Summary of transcribed data. Deciphered and evaluation sets have gold-standard decipherments; the evaluation set is held out. The test set has no gold-standard decipherment.

2011). In that case, the US Coast Guard intercepted messages and was able to make a bit of headway in deciphering them, but the real breakthrough came only when they obtained the applicable dictionary (key).

Unfortunately, there appear to be no automatic algorithms for solving book-based codes without the key.³ According to [Dunin and Schmeh \(2020\)](#):

“So far, there are no computer programs for solving codes and nomenclators available. This may change, but in the time being, solving a code or nomenclator message is mainly a matter of human intelligence, not computer intelligence.”

In this paper, we develop an algorithm for automatically attacking book-based codes, and we apply it to a corpus of historically-important codes from the late 1700s.

3 Wilkinson Letters

Our cipher corpus consists of letters to and from US General James Wilkinson, who first served as a young officer in the US Revolutionary War. He subsequently served as Senior Officer of the US Army (appointed by George Washington) and first Governor of the Louisiana Territory (appointed by Thomas Jefferson). Wilkinson also figured in the Aaron Burr conspiracy ([Isenberg, 2008](#)).

Long after his death, letters in a Cuban archive revealed the famous Wilkinson to be an agent of the Spanish Crown during virtually his entire service, and his reputation collapsed ([Linklater, 2009](#)).

Table 1 summarizes our Wilkinson correspondence data.⁴ We transcribe scans of manuscripts in the US Library of Congress. We have 73pp of

³[Kahn \(1996\)](#) suggests that national security services have long ago digitized all published books and applied brute-force to find the book that renders a given code into natural plaintext.

⁴All data is included with our released code (<https://github.com/c2huc2hu/wilkinson/>).

undeciphered text (Figure 3a) and 28pp of deciphered text (Figure 3b), with some overlap in content. Deciphered correspondence, with plaintext above ciphertext, likely resulted from manual encryption/decryption carried out at the time.

4 Encryption Method

As is frequent in book codes, there are two types of substitutions. Some plaintext words are enciphered using a large shared table that maps words onto numbers (**table-based code**). Other words are mapped with a shared dictionary (**book-based code**). Despite serious efforts, we have not been able to obtain the dictionary used in these ciphers.

In our transcription, we mark entries from the table portion with a caret over a single number, e.g., [123]^ . Before [160]^, the table seems to contain a list of people or place names; between [160]^ (“a”) and [1218]^ (“your”), a list of common words in alphabetic order; and finally more common words. The last block was likely added after the initial table was constructed, suggesting that the table was used to avoid having to look up common words in the dictionary.

The ciphertext for the dictionary code has two numbers that mark a word’s dictionary page and row number, respectively, plus one or two bars over the second number indicating the page column. For example, 123.[4]= refers to the fourth row of the second column of the 123rd page in the dictionary. From the distribution of cipher tokens, the dictionary is about 780 pages long with 29 rows per column, totaling about 45,000 words.

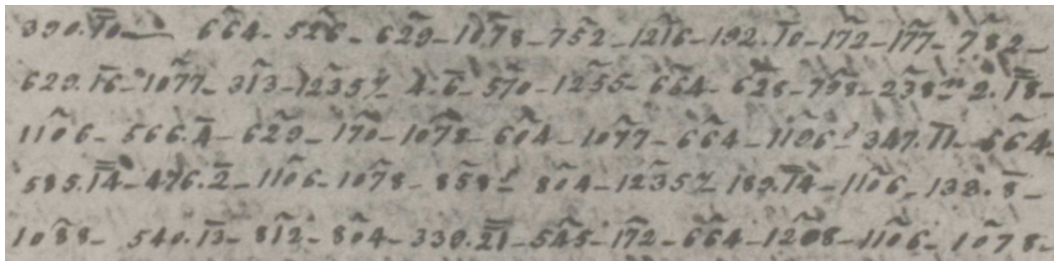
The cipher usually does not contain regular inflected forms of words, though inflections are sometimes marked with a superscript (e.g., ^{ing}). Numbers and some words are left in plaintext. Long horizontal lines mark the ends of sentences, but other punctuation is not marked.

5 Automatic Decryption Method

As the corpus includes a handful of deciphered pages, we employ a known-plaintext attack ([Kahn, 1996](#)).

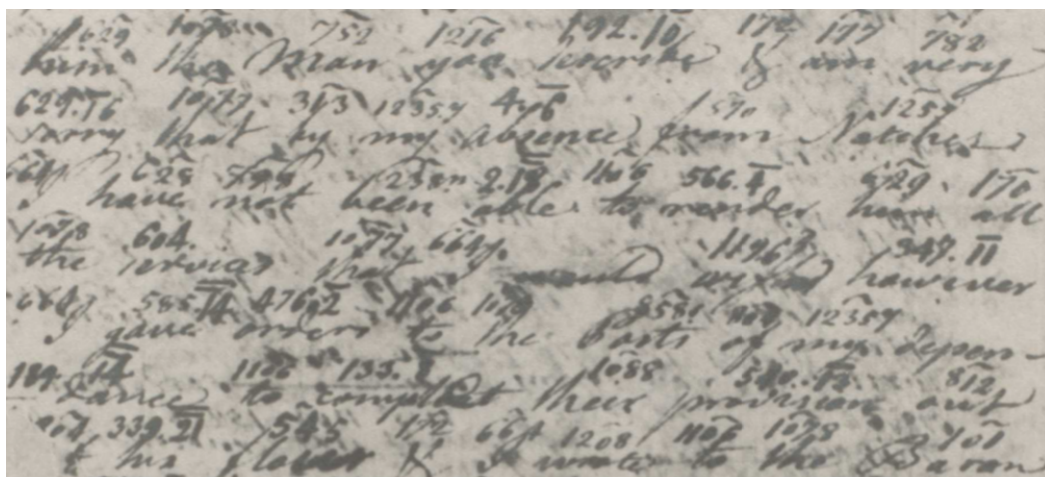
We first extract a small *wordbank* of known mappings, shown in Figure 3. Next, we apply the wordbank to our held-out evaluation ciphertext. We find that 40.8% of word tokens can be deciphered, mainly common words. After this step, we render a ciphertext as:

I ? he the man you ? and ? very ? that



390.[10]= . [664]^ [526]^ [629]^ [1078]^ [752]^ [1216]^ 192.[10]- [172]^ [177]^ [782]^
 629.[16]- [1077]^ [313]^ [1235]^ +y 4.[6]- [570]^ [1255]^ [664]^ [628]^ [798]^ [238]^ +n 2.[18]=
 [1106]^ 566.[4]- [629]^ [170]^ [1078]^ [604]^ [1077]^ [664]^ [1106]^ 347.[11]- [664]^
 585.[14]= 476.[2]- [1106]^ [1078]^ [858]^ [804]^ [1235]^ +y 189.[14]= [1106]^ 133.[8]-
 [1088]^ 540.[13]- [812]^ [804]^ 339.[21]= [545]^ [172]^ [664]^ [1208]^ [1106]^ [1078]^

(a)



... him the man you described and am very
 sorry that by my absence from Natches
 I have not been able to render him all
 the service that I wished however
 I gave orders to the ports of my dependence
 to complete their provisions over
 of his ? and I wrote to the ...

(b)

Figure 2: Sample encoded letters from US General James Wilkinson, with transcriptions: (a) portion of a ciphertext letter, (b) a recovered intermediate version of this letter with both ciphertext and plaintext.

Table wordbank		Dictionary wordbank		
Cipher	Plain	Cipher	Plain	Index
[13]^	Wilkinson	7.[24]-	acquisition	24
[33]^	Kentucky	15.[21]-	after	485
[92]^	Philadelphia	29.[29]-	answer	1305
[160]^	a	44.[28]-	attache	2174
[172]^	and	47.[21]-	attachment	2341
[229]^	be	59.[19]-	bearer	3035
[231]^	bear	65.[17]=	better	3410
[313]^	by	75.[29]-	bosom	3973
...	...	103.[40]=	chamber	5637
[1218]^	your	113.[4]-	cipher	6152
[1235]^	me	114.[20]-	civil	6226
[1249]^	policy

Figure 3: Table wordbank (left) and dictionary wordbank (right) extracted from five deciphered letters. The table contains proper nouns, an alphabetic word list (a–your) and other common words. Dictionary codes are of the form “page.[row].col”, where “-” indicates column 1 and “=” column 2. The Index guesses that “answer” is the 1305th word in the dictionary.

by ?? from ? I have not ?? to ? he ?
the service that I [...]

This is not yet a useful result. However, the wordbank also helps us to recover the rest of the plaintext. Since both the table and dictionary are in alphabetical order, we use the wordbank to constrain the decipherment of unknown words.

For example, given cipher word [163]^, we know from Figure 3 that its plaintext must lie somewhere between the two anchor-words [160]^ (which stands for “a”) and [172]^ (which stands for “and”). Moreover, it is likely to be closer to “a” than “and”. Repeating this for every cipher word in an undeciphered document, we construct a word lattice of all possible decipherments, shown in Figure 4. Our goal is then to search for the most fluent path through this lattice. Following are the details of our method:

Anchor. To propose candidate words between two anchors, we use a modern lemma-based dictionary with 20,770 entries.⁵ In this dictionary, for example, there are 1573 words between “attachment” and “bearer”.

Probabilities. We assign a probability to each candidate based on its distance from the ideal candidate. For example, in the table code, [163]^ is 30% of the way from [160]^ to [172]^ (Figure 5), so our ideal candidate decipherment of [172]^ will be 30% of the way between “a” and “and” in our modern dictionary. To apply this method to the dictionary code, we convert each cipher word’s

⁵www.manythings.org/vocabulary/lists/l (core ESL)

page/column/row to a single number n (the “Index” in Figure 3), which estimates that the cipher word corresponds to the n th word in the shared dictionary.

We use a beta distribution for assigning probabilities to candidate words, because the domain is bounded. We parameterize the distribution $B'(x; m, \beta)$ with mode m and sharpness parameter $\beta=5$. This is related to the standard parameterization, $B(x; \alpha, \beta)$, by:

$$B'(x; m, \beta) = B\left(x; \alpha = \frac{m\beta - 2m + 1}{1 - m}, \beta\right)$$

The sample space (0 to 1) is divided equally between the M words in the modern dictionary, so the i^{th} word gets probability:

$$\int_{\frac{i-1}{M}}^{\frac{i}{M}} B'(x; m, \beta) dx$$

In our example, because [163]^ is 30% of the way from [160]^ to [172]^, we have $m = 0.3$. There are $M = 650$ words in the modern dictionary between these two anchors, so the $i = 105^{\text{th}}$ word (“access”), gets probability 0.00231.

Inflections. We expand our lattice to include inflected forms of words (e.g., “find” → “find”, “found”, “finding”). We generate inflections with the *Pattern* library (De Smedt and Daelemans, 2012). Some words are generated more than once, e.g., “found” is both a base verb and the past tense of “find”. *Pattern* inflects some uncommon words incorrectly, but such inflections are heavily penalized in the best-path step. Inflections divide the probability of the original word equally.

Table edge cases. We replace unknown entries before the first anchor in the table with an arbitrary proper noun (“America”), and words outside the alphabetic section of the table with equal probability over a smaller vocabulary containing the 1000 most common words.⁶

Scoring lattice paths. After we have constructed the lattice, we automatically search for the best path. The best path should be fluent English (i.e., assigned high probability by a language model), and also be likely according to our wordbank (i.e., contain high-probability lattice transitions).

To score fluency, we use the neural GPT2 word-based language model (Radford et al., 2019), pre-trained on ~40GB of English text. We use the

⁶simple.wiktionary.org/wiki/Wiktionary:BNC_spoken_freq_01HW

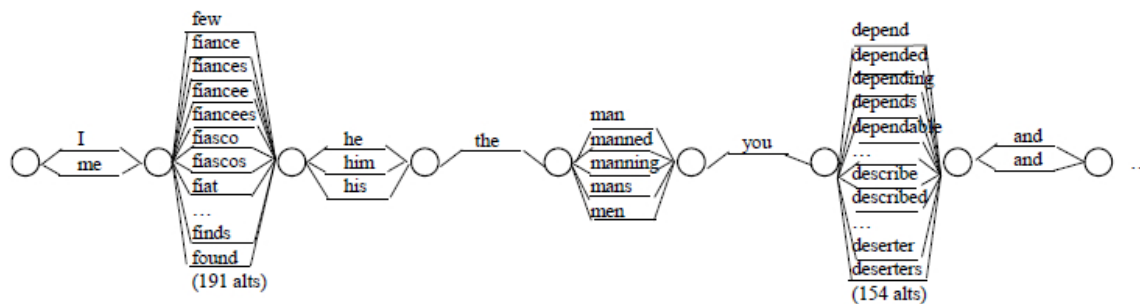


Figure 4: Turning an encrypted letter into a lattice of possible decipherments. Segments with few alternatives come from wordbank substitutions (and their automatically-produced morphological variants), while other segments come from interpolation-based guesses. Each link has an associated probability (not shown here). On average, we supply 692 alphabetically-close alternatives per segment, but supply fewer than ten for most.

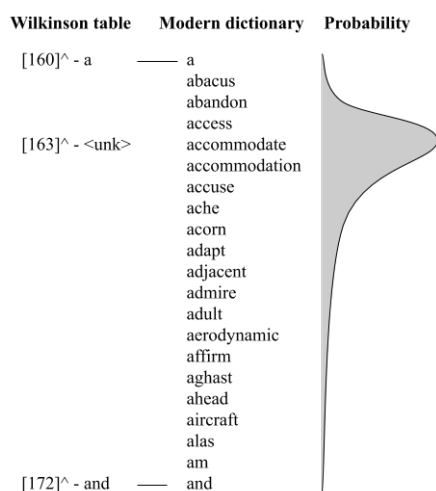


Figure 5: Interpolating a ciphertext word not present in the wordbank. When deciphering [163]^, we list candidate decipherments using a modern dictionary. We assign probabilities to candidates based on interpolation between anchor words, in this case “a” and “and”.

HuggingFace implementation⁷ with 12 layers, 768 hidden units, 12 heads, and 117M parameters.

Neural language models have significantly lower perplexity than letter- or word-based n-gram language models. For example, Tang and Lin (2018) benchmark WikiText-103 results for a Kneser-Ney smoothed 5-gram word model (test perplexity = 152.7) versus a quasi-recurrent neural network model (test perplexity = 32.8). This gives neural language models a much stronger ability to distinguish good English from bad.

Beam search. We search for the best lattice path using our own beam search implementation. To score a lattice transition with GPT, we must first tokenize its word into GPT’s subword vocabu-

⁷huggingface.co/transformers/pretrained_models.html

lary. Since alphabetically-similar words often start with the same subword, we create a subword trie at each lattice position; when a trie extension falls off the beam, we can efficiently abandon many lattice transitions at once.

6 Evaluation

To evaluate, we hold out one deciphered document from wordbanking. Some plaintext in that document is unreadable or damaged, so decryptions are added when known from the wordbank or obvious from context.

Table 2 gives our results on per-word-token decipherment accuracy. Our method is able to recover 73.8% of the word tokens, substantially more than using the wordbank alone (40.8%). We also outperform a unigram baseline that selects lattice paths consisting of the most popular words to decipher non-wordbank cipher tokens (46.9%).

The maximum we could get from further improvements to path-extraction is 91.3%, as 8.7% of correct answers are outside the lattice. This is due to unreadable plaintext, limitation of our modern dictionary, use of proper names [1]^ to [159]^, transcription errors, etc.

Table 3 details the effect of beam size on decipherment accuracy, runtime, and path score (combining GPT log probability with lattice scores). Increasing beam size leads us to extract paths with better scores, which correlate experimentally with higher task accuracy.

Figure 6 shows a portion of our solution versus the gold standard.

For tokens where our system output does not match the original plaintext, we asked an outside annotator to indicate whether our model captures

Method	Token accuracy
Apply wordbank	40.8
Unigram baseline	46.9
Our method, beam size = 1	68.3
Our method, beam size = 4	73.0
+ domain-tuning	73.2
+ self-learning	75.1
Our method, beam size = 16	73.8
+ domain-tuning + self-learning	75.1
Oracle	91.3

Table 2: Token decipherment accuracy. The Unigram baseline selects the most frequent word at each point in the lattice. Our method uses beam search to extract the path with highest GPT and lattices scores. Oracle selects the path through the lattice that most closely matches the gold decipherment.

Beam size	Runtime	Internal Best Path Score	Task Accuracy
1	1m 38s	-3714.9	68.3
4	6m 19s	-3569.6	73.0
16	25m 31s	-3544.9	73.8

Table 3: Effect of beam size on decipherment. A larger beam lets us extract paths with better internal scores, which correlate with better task accuracy.

the same meaning. For example, when our system outputs “I am much sorry” instead of “I am very sorry,” the annotator marks all four words as *same-meaning*. Under this looser criterion, accuracy rises from 73.0% to 80.1%.

We also decipher a Wilkinson ciphertext letter for which we have no plaintext. Transcription is less accurate, as we cannot confirm it using decipherment. The letter also includes phrases in plaintext French, which we translate to English before adding them to the lattice. Despite these challenges, the model still outputs relatively fluent text, including, for example: “... as may tend most powerfully and most directly to dissolve the whole America of the first states from the east and to cease the intercourse of the west.” This passage is consistent with Wilkinson’s plan to seek independence for parts of America.

7 Additional Methods

We experiment with three additional decipherment methods.

Weighted scoring. When scoring paths, we sum log probabilities from GPT and the lattice transitions, with the two sources equally weighted. This turns out to be optimal. Table 4 gives results when we multiply the lattice-transition score by a .

a	Accuracy	β	Accuracy
0.2	70.5	1	67.8
0.5	71.6	3	71.6
1	73.0	5	73.0
2	72.2	10	70.7

Table 4: Effects on decipherment accuracy of a (weight applied to lattice scores vs. GPT scores for each path) and β (sharpness parameter for candidate-word probabilities).

Halving the weight of the lattice scores degrades accuracy from 73.0 to 71.6 (-1.4 for beam=4), while doubling it degrades from 73.0 to 72.2 (-0.8 for beam=4). Table 4 also shows the impact of the sharpness parameter β on accuracy.

Domain-tuned language model. We collect letters written by Wilkinson⁸ totalling 80,000 word tokens, and fine-tune the GPT language model for one epoch on this data. The domain-tuned GPT increases decipherment accuracy from 73.0 to 73.2 (+0.2 for beam=4). Fine tuning for more than one epoch degrades decipherment accuracy. We found experiments with COFEA⁹ (American English sources written between 1765 and 1799) to be fruitless. We fine-tune a language model on a COFEA subset consisting of 3.8 million word tokens for one epoch, but this degrades accuracy from 73.0% to 65.6%.

Iterative self-learning. We apply iterative self-learning to improve our decipherment. After extracting the best path using beam search, we take the words with the smallest increases in perplexity on the lattice and language models, and we add them to the wordbank. The new wordbank provides tighter anchor points. We then construct a new lattice (using the expanded wordbank), search it, and repeat. This further improves decoding accuracy to 75.1 (+1.9 for beam=4).

8 Synthetic Data Experiments

We next experiment with synthetic data to test the data efficiency of our method. To create arbitrary amounts of parallel plaintext-ciphertext data, we encipher a book from Project Gutenberg,¹⁰ using a different machine readable dictionary.¹¹ We build wordbanks from parallel documents and use them to decipher a separately-enciphered book by the

⁸founders.archives.gov

⁹<https://lcl.byu.edu/projects/cofea/>

¹⁰*Wisdom of Father Brown*, <https://www.nltk.org/>

¹¹<https://packages.ubuntu.com/xenial/wamerican>

Parallel plain-cipher tokens	Wordbank size	Coverage of test tokens	Decipherment accuracy
500	227	53.8	66.0
800	351	58.3	69.6
2000	691	68.7	73.5
20000	2904	76.6	79.6

Table 5: Experiments with synthetic data. By enciphering material from Project Gutenberg, we produce arbitrary-sized wordbanks from arbitrary amounts of parallel plaintext-ciphertext. We then test how well those wordbanks support decipherment of new material.

same author.¹² The results are shown in Table 5.

9 Conclusion and Future Work

In this work, we show that it is possible to decipher a book-based cipher, using a known-plaintext attack and a neural English language model. We apply our method to letters written to and from US General James Wilkinson, and we recover 75.1% of the word tokens correctly.

We believe word-based neural language models are a powerful tool for decrypting classical codes and ciphers. Because they have much lower perplexities than widely-used n-gram models, they can distinguish between candidate plaintexts that resemble English at a distance, versus candidate plaintexts that are grammatical, sensible, and relevant to the historical context.

Acknowledgments

We would like to thank Johnny Fountain and Kevin Chatupornpitak of Karga7, and the staff who transcribed data from the Library of Congress, who provided scans of the original documents. We would also like to thank the anonymous reviewers for many helpful suggestions.

References

Taylor Berg-Kirkpatrick and Dan Klein. 2013. Decipherment with a million random restarts. In *Proc. EMNLP*.

Gary Blackwood. 2009. *Mysterious Messages: A History of Codes and Ciphers: A History of Codes and Ciphers*. Penguin.

Robert Churchhouse. 2002. *Codes and Ciphers*. Cambridge.

¹²*The Man Who Was Thursday*, <https://www.nltk.org/>

Eric Corlett and Gerald Penn. 2010. An exact A* method for deciphering letter-substitution ciphers. In *Proc. ACL*.

Tom De Smedt and Walter Daelemans. 2012. *Pattern for Python*. *J. Mach. Learn. Res.*, 13:2063–2067.

Qing Dou and Kevin Knight. 2012. Large-scale decipherment for out-of-domain translation. In *Proc. EMNLP*.

Elonka Dunin and Klaus Schmeh. 2020. *Codebreaking & Cryptograms: A Practical Guide*. Taschen. (draft of forthcoming).

George W. Hart. 1994. To decode short cryptograms. *Communications of the Association for Computing Machinery*, 37(9).

Bradley Hauer, Ryan Hayward, and Grzegorz Kondrak. 2014. Solving substitution ciphers with combined language models. In *Proc. COLING*.

Nancy Isenberg. 2008. *Fallen Founder: The Life of Aaron Burr*. Penguin.

David Kahn. 1996. *The Codebreakers: The Comprehensive History of Secret Communication from Ancient Times to the Internet*. Scribner.

Nishant Kambhatla, Anahita Mansouri Bigvand, and Anoop Sarkar. 2018. Decipherment of substitution ciphers with neural language models. In *Proc. EMNLP*.

John C. King. 1993. A reconstruction of the key to Beale cipher number two. *Cryptologia*, 17(3).

Kevin Knight, Beáta Megyesi, and Christiane Schaefer. 2011. The Copiale cipher. In *ACL workshop on Building and Using Comparable Corpora (BUCC)*.

Kevin Knight and Kenji Yamada. 1999. A computational approach to deciphering unknown scripts. In *Proc. ACL Workshop on Unsupervised Learning in Natural Language Processing*.

Andro Linklater. 2009. *An Artist in Treason: The Extraordinary Double Life of General James Wilkinson*. Walker.

Beáta Megyesi, Nils Blomqvist, and Eva Pettersson. 2019. The DECODE database: Collection of ciphers and keys. In *Proc. 2nd International Conference on Historical Cryptology (HistoCrypt19)*.

NSA. 2011. *Cryptologic Aspects of German Intelligence Activities in South America during World War II*. Center for Cryptologic History National Security Agency.

Malte Nuhn, Julian Schamper, and Hermann Ney. 2014. Improved decipherment of homophonic ciphers. In *Proc. EMNLP*.

Edwin Olson. 2007. Robust dictionary attack of short simple substitution ciphers. *Cryptologia*, 31(4).

Ciphertext	Gold	Wordbank only	Unigram Baseline	Our method
[229]^+ing	being	be	be	being
[186]^	at	at	at	at
[1049]^	such	such	such	such
[160]^	a	a	a	a
212.[20]=	distance		distant	distance
[570]^	from	from	from	from
[90]^	North Carolina	America	America	America
[172]^	and	and	and	and
286.[14]=	for	for	for	for
[509]^	fear		father	fear
[804]^	of	of	of	of
446.[1]- +ing	missing		mistress	missing
[1218]^	your	your	your	your
294.[20]=	garrison		from	friend
[1084]^	there		therefore	therefore
286.[14]=	for	for	for	for
[1078]^	the	the	the	the
153.[5]-	conveyance		could	convenience
[804]^	of	of	of	of
678.[6]=	this	this	this	this
[664]^	I		I	I
[177]^	am		any	am
467.[24]- +d	obliged		of	obliged
[1106]^	to	to	to	to
[1206]^	write	write	written	write
[1106]^	to	to	to	to
[1216]^	you	you	you	you
[807]^	in	on	in	in
[160]^	a	a	a	a
349.[1]=	hurry		I	hurry
572.[5]- +ing	resuming		said	requiring
251.[6]=	every	every	every	every
[852]^	point		people	precaution

Figure 6: Decipherment of a portion of our evaluation set, compared to the gold standard.

- Eva Pettersson and Beáta Megyesi. 2019. Matching keys and encrypted manuscripts. In *Proc. 22nd Nordic Conference on Computational Linguistics*.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Sujith Ravi and Kevin Knight. 2008. Attacking decipherment problems optimally with low-order n-gram models. In *Proc. EMNLP*.
- Sujith Ravi and Kevin Knight. 2011. Bayesian inference for Zodiac and other homophonic ciphers. In *Proc. ACL*.
- Claude E. Shannon. 1951. Prediction and entropy of printed english. *Bell System Technical Journal*, 30(1).
- Simon Singh. 2000. *The Code Book: The Science of Secrecy from Ancient Egypt to Quantum Cryptography*. Anchor.
- Raphael Tang and Jimmy Lin. 2018. Progress and tradeoffs in neural language models. *CoRR*, abs/1811.00942.
- Mark Urban. 2002. *The Man Who Broke Napoleon's Codes: The Story of George Scovell*. Gardners.