



**Decrypting Classical Cipher Text
Using Markov Chain Monte Carlo**

by

Jian Chen
Department of Statistics
University of Toronto

and

Jeffrey S. Rosenthal
Department of Statistics
University of Toronto

Technical Report No. 1005 May 22, 2010

TECHNICAL REPORT SERIES

University of Toronto
Department of Statistics

Decrypting Classical Cipher Text Using Markov Chain Monte Carlo

Jian Chen and Jeffrey S. Rosenthal*
Department of Statistics, University of Toronto

May, 2010

We investigate the use of Markov Chain Monte Carlo (MCMC) methods to attack classical ciphers. MCMC has previously been used to break simple substitution ciphers. Here, we extend this approach to transposition ciphers and to substitution-plus-transposition ciphers. Our algorithms run quickly and perform fairly well even for key lengths as high as 40.

1 Introduction

Cryptography (e.g. [11]) is the study of algorithms to encrypt and decrypt messages between senders and receivers. And, Markov chain Monte Carlo (MCMC) algorithms (e.g. [5, 10]) are popular methods of approximately sampling from complicated probability distributions. Traditionally these two subjects have been quite distinct.

However, recently MCMC algorithms have been used to iteratively converge to solutions of how to break simple substitution codes. This approach was first introduced by Marc Coram and Phil Beineke in the Stanford statistical consulting service (see the Introduction to Diaconis [2]), and later studied more systematically by Connor [1]. Table 1 shows output from a typical run of this algorithm (in this case, decrypting the first line of the Project Gutenberg [9] copy of *Oliver Twist*).

Iteration #	First Line of Decrypted Text
0	LIW PSKMWCL YNLWRDWSY WDKKJ KH KGUXWS LAUEL DQ CIVSGWE FUCJWRE
200	RAS KINJSBR MDRSEHSIM SHNNV NW NGUZSI RPUOR HY BATIGSO LUBVSEO
400	ARS HUNJSPA GDASEBSUG SBNNV NW NMIKSU AFIOA BY PRTUMSO LIPVSEO
600	ARE HLNJEP A KOAESBELK EBNNW NG NMIVEL AFIDA BY PRULMED TIPWESD
800	IME KNSJEPI HUIETBENH EBSSG SW SLOVEN ICODI BY PMANLED ROPGETD
1000	IME GNOJEPI HUIETBENH EBOOK OF OLSVEN ICSDI BY PMANLED RSPKETD
1200	SME GNOJECS HUSETBENH EBOOK OF OLIVEN SPIIDS BY CMANLED RICKETD
1400	SME PNOJECS HUSETBENH EBOOK OF OLIVEN SWIDS BY CMANLED RICKETD
1600	SHE MROJECS GUSETBERG EBOOK OF OLIVER SWIDS BY CHARLED NICKETD
1800	SHE PROJECS GUSETBERG EBOOK OF OLIVER SWINS BY CHARLEN DICKETN
2000	SHE PROJECS GUSELBERG EBOOK OF ONIVER SWITS BY CHARNET DICKELT
2200	THE PROJECT GUTENBERG EBOOK OF OLIVER TWIST BY CHARLES DICKENS

Table 1: A sample run of a simple MCMC decryption algorithm.

*Supported in part by NSERC of Canada.

We see that the algorithm begins with encrypted text that looks like gibberish, and then gradually (in this case, after 2200 iterations) breaks the code and recovers the correct original text.

In this paper, we significantly extend the use of MCMC in decryption, to the more complicated cases of transposition ciphers, and substitution-plus-transposition ciphers. Our key innovations include combining multiple independent runs, cycling between different cipher attacks, and (for substitution-plus-transposition ciphers) using a uni-gram attack as an initialization point for a sequence of bi-gram attacks.

Extensive computer simulations indicate that our algorithms run quickly, and work quite well even for key lengths as large as 40. These results appear to improve upon existing decryption methods [3], and suggest that MCMC algorithms can be of genuine use for decrypting encoded text.

This paper is organized as follows. We present background on cryptography and on MCMC below. The use of MCMC for decryption is outlined in Section 2. We then present detailed algorithms and simulation results for attacking substitution ciphers (Section 3), transposition ciphers (Section 4), and substitution-plus-transposition ciphers (Section 5). All of the software used for our simulations is freely available at probability.ca/decipher.

1.1 Background on Cryptography

In cryptography, the original text is called the plain text, and the encrypted text is called the cipher text. The algorithms to perform encryption and decryption are referred to as ciphers. Usually a cipher contains one or two keys. In a symmetric key algorithm (e.g. DES), the decryption key is the same as the encryption key (or just the inverse function of it). In an asymmetric key algorithm (e.g. RSA), two different keys are used. The public key is used for encryption and a private key is for the decryption.

Ciphers can also be categorized in a different way, as classical ciphers and modern ciphers. Classical ciphers, such as the substitution and transposition ciphers considered herein, perform encryption and decryption text manipulations at the byte level. Modern ciphers, such as DES (symmetric key) and RSA (asymmetric key), perform encryption and decryption at the bit level, and are correspondingly more complicated and secure than the classical ciphers, and we do not consider them here. (Note, however, that a simplified version of DES, called SDES [8, 4], can be regarded as a special case of substitution cipher and is thus included in our results below.)

A simple substitution cipher works by replacing each letter with another one. In this paper, we only substitute alphabetic letters; spaces are left untouched and all other non-alphabetic characters are removed. So, the number of the possible keys is equal to $26! \doteq 4 \times 10^{26}$. Table 2 illustrates an encryption and decryption example of a simple substitution cipher. For the encryption, all 'A's in the plain text are replaced by letter 'X', 'B's replaced by 'E', etc. For the decryption, all 'A's in the cipher text are replaced by 'I', 'B's replaced by 'C' etc. Note that the encryption key is the inverse function of the decryption key.

plain text	THE PROJECT GUTENBERG EBOOK OF OLIVER TWIST
encryption key	XEBPROHYAUFTIDSJLKZMWVNGQC
cipher text	MYR JKSURBM HWMRDERKH RESSF SO STAVRK MNAZM
decryption key	ICZNBKXGMPRQTFWDYEOLJVUAHS
decrypted text	THE PROJECT GUTENBERG EBOOK OF OLIVER TWIST

Table 2: A simple example of a substitution cipher encryption and decryption.

Another classical cipher is the transposition cipher (also called the permutation cipher). The letters in the plain text stay the same but their positions are rearranged in a different order. A simple transposition cipher works by splitting the plain text into fixed sized blocks. The length of the key (also called the period) is the same as the size of the block. Letters in each block are permuted according to a same pattern(the key). Table 3 illustrates an example of encryption and decryption by a transposition cipher with key length 10. Note the encryption key is the inverse function of the decryption key.

plain text	T	H	E		P	R	O	J	E	C
encryption key	1	9	3	7	0	4	5	8	6	2
cipher text	H	C		J	T	P	R	E	O	E
decryption key	4	0	9	2	5	6	8	3	7	1
decrypted text	T	H	E		P	R	O	J	E	C

Table 3: A simple example of a transposition cipher encryption and decryption.

A *product* cipher combines a sequence of simple transformations such as substitution, transposition and other arithmetic. SP-network (Substitution-permutation network) is an example of a product cipher, involving repeated applications of substitutions (S-box) and permutations (P-box), which is very common in the design of modern ciphers such as DES. Herein, we consider the special case of a simple substitution-plus-transposition (or, substitution-transposition) cipher, in which the plain text is encrypted by a substitution cipher followed by a transposition cipher, causing additional challenges.

Frequency analysis (e.g. [12]) is the study of the frequencies of letters or combination of letters in the cipher text. In a particular language (e.g. English) certain letters and there combinations occurs more frequently than others. The frequencies of letters are also called n -gram, i.e. unigram stand for single letter frequencies, bigram for combination of 2 letters, trigram for 3 letters, etc. For example, in English 'E' is the most used single letter while 'Z' is the least used single letter. 'TH' and 'ER' are pairs which arise frequently. Classical ciphers are often broken by comparing the letter frequencies of the cipher text to a reference text (usually a large text such as *War and Peace*).

When faced with a simple substitution cipher, the simplest form of frequency analysis is a uni-gram attack, which involves simply replacing the most frequent letter in the cipher text by the most frequent occurred letter in the reference text, and the second-most-frequent letter in the cipher text by the second-most-frequent letter in the reference text, and the third-most by the third-most, and so on. As a first experiment, we tried this simple uni-gram attack on the novel *Oliver Twist* after a random simple substitution cipher was applied to it. This attack does not succeed very well, revealing just 16 out of 26 letters in the cipher text (Table 4). This is because some letters (e.g. 'R' and 'S', or 'C' and 'W') have similar frequencies and are thus likely to be interchanged in such an attack.

plain text	THE PROJECT GUTENBERG EBOOK OF OLIVER TWIST
decrypted text	THE PSOJEWTFUTEIBESF EBOOK OGNLVES TCNRT

Table 4: Attempted decryption of *Oliver Twist* with a uni-gram attack.

Because of results like this, more complicated attacks involving pair frequencies have to be employed. Since pairs cannot be directly substituted as in uni-gram attacks, this leads to more complicated algorithms, as we discuss herein.

1.2 Background on MCMC

MCMC algorithms have long been used by physicists and statisticians to sample from complicated high-dimensional probability distributions. Let $\pi(\cdot)$ be an important possibly-unnormalised density (for example, the posterior distribution from a Bayesian inference problem) on a state space \mathcal{X} (often \mathcal{X} is an open subset of \mathbf{R}^d). MCMC proceeds by defining an iterative sequence X_0, X_1, X_2, \dots of \mathcal{X} -valued random variables such that

$$\lim_{n \rightarrow \infty} \mathbf{P}(X_n \in A) = \int_A \pi(x) dx. \quad (1)$$

It follows that for large n , the value X_n is approximately a “sample” from $\pi(\cdot)$. Repeating or continuing this process leads to multiple samples, which can then be used to estimate probabilities and expected values with respect to $\pi(\cdot)$.

The simplest version of MCMC is the full-dimensional Metropolis algorithm [7], which proceeds as follows:

- Choose an initial state $X_0 \in \mathcal{X}$.
- For $n = 1, 2, 3, \dots$,
 - Propose a new state $Y_n \in \mathcal{X}$ from some symmetric proposal density $q(X_{n-1}, \dots)$.
 - Let $U_n \sim \text{Uniform}[0, 1]$, independently of X_0, \dots, X_{n-1}, Y_n .
 - If $U_n < (\pi(Y_n)/\pi(X_{n-1}))$, then “accept” the proposal by setting $X_n = Y_n$, otherwise “reject” the proposal by setting $X_n = X_{n-1}$.

Thus, the acceptance probability of each proposal is equal to $\min(1, \pi(Y_n)/\pi(X_{n-1}))$. This probability is chosen precisely so that the resulting Markov chain X_0, X_1, X_2, \dots will be *reversible* with respect to $\pi(\cdot)$, so that $\pi(\cdot)$ is a stationary density, and under the mild assumptions of irreducibility and aperiodicity, the probabilities will converge to those of $\pi(\cdot)$ as in (1). (It is not essential the the proposal density $q(x, \cdot)$ be symmetric, but if it is not then the acceptance probability must be appropriately modified, so for simplicity we do not consider that case here.)

It is also possible to replace $\pi(x)$ by a power, $(\pi(x))^p$, so that the acceptance condition above is replaced by $U_n < ((\pi(Y_n)/\pi(X_{n-1}))^p)$. This is a *tempering* modification in which p plays the role of *inverse temperature*, and we shall refer to p as a *scaling parameter*. Such a modification changes and flattens (for $0 < p < 1$) the density $\pi(\cdot)$, potentially changing the corresponding probabilities and expected values, but leaving the *mode* (argmax) of $\pi(\cdot)$ unchanged. Thus, such tempering modifications can help the chain escape from local modes, while preserving the same mode; we shall make use of them herein.

2 Using MCMC to Break Classical Ciphers

We now discuss the use of MCMC for breaking classical ciphers.

For this application, the relevant Markov chain has state space \mathcal{X} consisting of all possible decryption keys (a large but finite state space). That is, each possible decryption key is a possible

state of the Markov Chain. Following [2, 1], we make use of a long reference text such as *War and Peace*. For each pair of characters β_1 and β_2 (e.g. $\beta_1 = \text{T}$ and $\beta_2 = \text{H}$), we let $r(\beta_1, \beta_2)$ record the number of times that specific pair (e.g. “TH”) appears consecutively in the reference text. Similarly, for a putative decryption key $x \in \mathcal{X}$, we let $f_x(\beta_1, \beta_2)$ record the number of times that pair appears when the cipher text is decrypted using the decryption key x . To avoid problems from zeroes, we also add one to each of $r(\beta_1, \beta_2)$ and $f_x(\beta_1, \beta_2)$.

For a particular decryption key x , we then define its *score function* as follows:

$$\pi(x) = \prod_{\beta_1, \beta_2} r(\beta_1, \beta_2)^{f_x(\beta_1, \beta_2)}. \quad (2)$$

This function can be thought of as multiplying, for each consecutive pair of letters in the decrypted text, the number of times that pair occurred in the reference text. Intuitively, the score function is higher when the pair frequencies in the decrypted text most closely match those of the reference text, and the decryption key is thus most likely to be correct. (In our computer programs, we compute (2) on a log scale for easy calculation and to avoid numerical errors.)

In terms of this score function, we use the following general MCMC algorithm to break the classical ciphers:

- Choose an initial state (initial decryption key), and a fixed scaling parameter p .
- Repeat the following steps for many iterations (e.g. 10,000 iterations).
 - Given the current state x , propose a new state y from some symmetric density $q(x, y)$.
 - Sample $u \sim \text{Uniform}[0, 1]$ independently of all other variables.
 - If $u < (\frac{\pi(y)}{\pi(x)})^p$ then accept the proposal y by replacing x with y , otherwise reject y by leaving x unchanged.

By the usual Markov chain convergence theorem, this Markov chain will converge in probability to its stationary distribution, which in this case means it will converge to the distribution with density proportional to $(\pi(x))^p$ with $\pi(\cdot)$ as in (2). So, intuitively, after many iterations, the algorithm is likely to be at a decryption key which gives decryption text pair frequencies close to those of the reference text, and is thus more likely to be correct.

2.1 Testing methodology

To test our algorithms, we shall primarily use the four texts listed in Table 5. During the programming and initial testing we used *War and Peace* as the reference text and *Oliver Twist* as the plain text. All four texts were then used to test our final attack algorithms. (A systematic investigation of MCMC decryption results with many different choices of texts was undertaken by Connor [1], so we do not repeat that here.)

Text	Author	Publication Date
<i>War and Peace</i>	Leo Tolstoy	1869
<i>Oliver Twist</i>	Charles Dickens	1838
<i>Pride and Prejudice</i>	Jane Austen	1813
<i>Ice Hockey (Wikipedia Page)</i> [6]	Wikipedia	2010

Table 5: Cipher texts and reference texts used in our attacks.

For simplicity, we first convert all letters to upper case, and remove or convert to spaces all non-alphabetic characters. So in total we have 27 characters (26 upper case English alphabet letters plus one space character), which we number from 0 to 26.

For each attack algorithm we consider, we run the encryption and the decryption process 100 separate times. In each such run, a random key is generated to encrypt the plain text, and the attack is then performed on the cipher text. At the end of the attack, we compare the decrypted text with the plain text. If the decrypted text is the same as the cipher text, it is a successful run. Obviously, the more successful runs out of 100, the better has our algorithm performed.

Even if a run is not completely successful, it is still true that if we successfully guessed “most” of the letters, i.e. our decryption was “mostly” successful, then this may still be helpful because the remaining cipher text can probably then be determined by human intervention. For this reason, we also want a definition of “accuracy” to measure how close the decrypted text is to the plain text.

For a substitution ciphers, the accuracy is defined as $\frac{m_s}{n_s}$, where m_s is the number of letters correctly revealed, and n_s is the number of available letters in the plain text (usually n_s is 26, but it may be less than 26 for short cipher text). A letter is said to be correctly revealed if the position of its first appearance in the plain text is the same as that of the decrypted text.

For a transposition ciphers, we define accuracy as $\frac{m_t}{n_t}$, where n_t is equal to the key length minus 2, and m_t is the number of letters correctly placed in one period (the key length). A letter is said to be correct positioned if it has the same neighbors in the decrypted text as in the plain text. We do not count the letters in the start and end positions as they only have one neighbor.

We also measure how long it takes for our attacks to run, since a good attack should finish within a reasonable time. Our program is written in C++ and was run on a MacBook Pro with the system configuration as in Table 6.

CPU	2.26 GHz Intel Core 2 Duo
Memory	4GB 1067 MHz DDR3 Memory
OS version	Mac OS X Version 10.6.3
Compiler	g++ i386-apple-darwin10-g++-4.2.1

Table 6: System configuration of the machine running the attacks.

3 Attacks on Substitution Ciphers

We now consider attacks on substitution ciphers, in which an unknown permutation is applied to the 26 letters of the English alphabet. Following [2, 1], we use MCMC algorithms as in the previous section, and find good results.

Our Markov chain state space now consists of all the $26! \doteq 4 \times 10^{26}$ possible permutations of 26 letters. We let the initial state be the identity permutation 'ABCD...XYZ' (so the decrypted text using this key is identical to the cipher text itself).

A key part of the MCMC algorithm is to define a proposal so the chain is detailed balanced and guaranteed to converge to its stationary distribution. Similar to [2, 1], we propose a new key by swapping 2 randomly selected letters in the current key. So, each such swap has proposal probability $1/n^2$. Note that these proposals are symmetric. (For simplicity, we allow our program

to propose swapping a letter with itself, e.g. swapping 'A' and 'A', even though such proposals will not change the chain's state.)

For this algorithm, we next try adjusting various parameters to see which tuning allows the algorithm to perform optimally.

3.1 Number of Iterations

Table 7 shows that by increasing number of iterations, we improve the accuracy and the number of successful runs. But the accuracy doesn't change much after 10,000 iterations. At this point, although the accuracy is quite high (greater than 90%) which mean most of the runs were very close to the correct result, the number of completely successful runs is fairly low (around 50–60 out of 100). Next we try to improve the algorithm by tuning different parameters.

iterations	accuracy	no. of successful runs
1,000	0.5196	0
2,000	0.7732	19
5,000	0.9060	47
10,000	0.9064	51
20,000	0.9348	59
50,000	0.8932	54

Table 7: Results from initial attempt to decrypt substitution ciphers using bi-grams, with different numbers of MCMC iterations.

3.2 Tuning the Scaling Parameter

The scaling parameter is very important in the MCMC algorithm. Larger scaling parameters give lower acceptance rates. But if the acceptance rate is too low, the chain is moving too slowly, and it will take too long to converge. Smaller scaling parameters gives higher acceptance rates. But if acceptance rate is too high, the chain will move too often and may not always stay in the stationary distribution.

We see from Table 8 that in this case, the best result is when the scaling parameter is set to 1, leading to an acceptance rate of 0.04. If we increase the acceptance rate by lowering the scaling parameter, the chain will not converge well (on average, it only revealed 26% of the cipher text after 10,000 iterations).

scaling parameter	acceptance rate	accuracy	no. of successful runs
0.05	0.27	0.2664	0
0.1	0.12	0.6184	0
1	0.04	0.9064	51

Table 8: Results of bi-gram attacks for substitution ciphers after 10,000 iterations, for different choices of the scaling parameter.

3.3 Remembering the Best Score Function

The above runs had our algorithm return the *final* decryption key from the run, i.e. whatever key the Markov chain ends up at after a full run of (say) 10,000 iterations.

However, we found that many of our runs revealed the plain text (e.g. "THE PROJECT") in the middle of the run, but then later jumped away from it (e.g. "THE PROZECT"). We know that larger score functions usually indicate better solutions. So, instead of having our algorithm return the *final* decryption key from the run, we instead have it return whichever decryption key gave the largest log score function. The results are presented in Table 9. Comparison with Table 7 shows that the new modification leads to significantly better results.

iterations	accuracy	no. of successful runs
1,000	0.5300	1
2,000	0.7716	20
5,000	0.9172	87
10,000	0.9312	90
20,000	0.9148	87
50,000	0.9488	93

Table 9: Results of bi-gram attacks on substitution ciphers, when we return whichever key maximizes the score function.

3.4 How Much Cipher Text is Needed

Usually we use the entire cipher text when computing the score function (2) at each iteration. We can ask whether it is more efficient, and of comparable accuracy, to compute the score function at each iteration using just a (random) subset of the cipher text. Table 10 indicates that in this case, the time spend on the decryption is essentially independent of the length of the cipher text used. On the other hand, the accuracy is already quite high (over 93%) when using just 2,000 characters of cipher text.

cipher text	accuracy	no. of successful runs	duration (in seconds)
1,000	0.7143	0	0.4441
2,000	0.9312	90	0.4442
full cipher text	0.9831	97	0.4381

Table 10: Results from attacks on substitution ciphers using bi-gram, when using different amounts of cipher text. Each run uses 10,000 iterations, and returns whichever key maximize the log score.

These results suggest that for simple substitution ciphers, it does not much matter (for either speed or accuracy) whether we use just 2,000 characters of cipher text, or the entire cipher text. However, since our main interest is in transposition-related ciphers for which speed is much more effected (see below), for our final attack we use just 2,000 characters of cipher text for simple substitution ciphers as well.

3.5 Independent Repetitions

Experimentation indicates mixed result when using just 2000 randomly-chosen consecutive characters from the cipher text for the attack (Table 11). That is, some selections from the cipher text are better for decryption than others.

cipher text	accuracy	no. of successful runs	duration (in seconds)
574798	0.9492	0	0.3906
416031	0.9488	90	0.3933
243158	0.9840	97	0.3932
551774	0.9452	0	0.3940
223511	0.9596	94	0.3939

Table 11: Results of attacks on substitution ciphers using bi-gram, depending on the position of cipher text used in the attack. Each run uses 10,000 iterations, and 2000 characters of cipher text, and returns whichever key maximize the log score.

This suggests that our final attack should, instead of using just one run, use several independent repeated runs, and return whichever final result has the largest score function (from (2) computed using the *entire* cipher text). We use this approach in our final attack below.

3.6 Tri-gram Attack

As a final check, we tried modifying the previous MCMC algorithm to use tri-grams (triple letters frequencies) instead of bi-grams. That is, we replace the score function (2) by:

$$\pi(x) = \prod_{\beta_1, \beta_2, \beta_3} f_r(\beta_1, \beta_2, \beta_3)^{f_d(\beta_1, \beta_2, \beta_3)},$$

where $\beta_1, \beta_2, \beta_3$ are all possible three-characters combinations. And, we replace $r(\beta_1, \beta_2)$ and $f_x(\beta_1, \beta_2, \beta_3)$ by the corresponding triple letter frequencies of the reference text and the decrypted text respectively.

This new tri-gram attack also works, but the result is not as good as the attack using the bi-grams (Table 12). Therefore, we stick with bi-grams for the final version of our attack.

iterations	accuracy	no. of successful runs	duration (in seconds)
1,000	0.5336	1	0.9089
2,000	0.7652	46	1.6892
5,000	0.7896	75	4.0435
10,000	0.8920	87	7.9283

Table 12: Results of attacks on substitution ciphers by tri-gram, for different numbers of iterations. Each run returns whichever key maximize the log score.

3.7 Attack for Substitution Ciphers – Final Version

Based on the above experimentation, we take the final version of our attack to involve 10,000 iterations, with scaling parameter 1, and with cipher text length 2,000.

To investigate how our program works, we apply this attack to different combinations of cipher text and reference text. The results are presented in Table 13.

cipher text	reference text	accuracy	no. of successful runs
<i>Oliver Twist</i>	<i>War and Peace</i>	1.0000	100
<i>Pride and Prejudice</i>	<i>War and Peace</i>	1.0000	100
<i>Ice Hockey (Wikipedia Page)</i>	<i>War and Peace</i>	1.0000	100
<i>Pride and Prejudice</i>	<i>Oliver Twist</i>	1.0000	100
<i>War and Peace</i>	<i>Oliver Twist</i>	0.9977	97
<i>Ice Hockey (Wikipedia Page)</i>	<i>Oliver Twist</i>	0.9869	83
<i>War and Peace</i>	<i>Pride and Prejudice</i>	0.9977	97
<i>Ice Hockey</i>	<i>Pride and Prejudice</i>	0.9977	97
<i>Oliver Twist</i>	<i>Pride and Prejudice</i>	0.9985	98
<i>Pride and Prejudice</i>	<i>Ice Hockey (Wikipedia Page)</i>	1.0000	100
<i>War and Peace</i>	<i>Ice Hockey (Wikipedia Page)</i>	0.9938	92
<i>Oliver Twist</i>	<i>Ice Hockey (Wikipedia Page)</i>	0.9750	74

Table 13: Results of our final attack on substitution ciphers with key length 20, for different choices of cipher text and reference text. Each run uses 5 repetitions of 10,000 iterations each, with 2000 characters of cipher text, and scaling parameter 1, and returns whichever key gives the highest score function.

We see from the table that our final attack algorithm performed very well, often achieving perfect or near-perfect scores. The only sub-par performances arose when using the *Ice Hockey* Wikipedia page, which is much shorter than the three novels (less than 8,000 words) and thus provides insufficient text for our algorithm to perform well. (Furthermore it was written in the modern era so it may have somewhat different language usage as well.)

4 Attacks on Transposition Ciphers

We now turn our attention to Transposition Ciphers. Since Transposition Ciphers only move letters around, there is no change to the frequencies of single letters, so we certainly can't use a uni-gram attack to break it. Instead, we concentrate on bi-gram attacks (i.e., again using the frequencies of pairs of letters).

The state space depends on the key length of the transposition cipher. For key length k , there are $k!$ possible decryption keys, corresponding to all possible permutations of $0, 1, 2, \dots, k - 1$. We again choose the initial decryption key to be the identity permutation, so the decrypted text using this key is identical to the cipher text.

The score function is thus again the same as in (2), and the algorithm and acceptance rate are still the same as in Section 2. The only potential difference from the bi-gram attack for the substitution cipher concerns the proposal distribution, as we now discuss.

4.1 Swap Moves versus Slide Moves

For the proposal, first we tried the same swap moves as in our substitution cipher attacks. But we found the swap moves are not very efficient in some cases. For example, suppose the plain text is “PROJECT”, and the current decryption key gives a decrypted text “ROJECTP”. Then this is very close to the correct answer, but we need at least 6 swap moves find the correct decryption key from here.

Instead of swapping characters, we can instead use a “slide move” of randomly take one character out and insert it back to a random position of the remaining text. We call this proposal a slide move. For example, suppose the key length $k = 8$, and the letter at position 3 is taken out and inserted back in position 6 in the remaining text. Then the text “ABCDEFGH” will become “ABCEFGDH” by this slide move.

Sliding moves of a single letter work pretty well with small key length. But as the key length gets larger, it also becomes less efficient. For example, suppose the plain text is “THE PROJECT”, but the current decryption key generates a decrypted text “PROJECT THE”. To get the correct key using single letter slide moves, we need at least 3 moves (move each character in “THE” to the left of “PROJECT”). But each such move may lower the score function since we are breaking the word “THE”, so it will more likely be rejected by the algorithm, making this a difficult feat for our algorithm to perform.

This can be solved by using slide move involving entire blocks of letters. That is, we select a random block of letters, which we remove and insert back somewhere within the remaining text. For the “PROJECT THE” example, the word “HE” can be moved to the left of “PROJECT” by just one move. Formally speaking, for a key length k , the new proposal is to slide move a block of n letters from position k_1 to k_2 , where $n \sim \text{Uniform}\{0, \dots, p - 2\}$, $k_1 \sim \text{Uniform}\{0, \dots, p - n + 1\}$, and $k_2 \sim \text{Uniform}\{0, \dots, p - n + 1\}$.

For example, with key length $k = 8$, we could propose to move 2 letters from position 3 to position 6 if $k = 8, n = 2, k_1 = 3, k_2 = 6$, in which case “01234567” will become “01256347” (since ‘34’ is moved to after ‘6’).

Tables 14 and 15 compare algorithms using swap moves, single letter slide moves, and block slide moves, each for 1,000 iterations. With key length $k = 10$, we see some improvement of the accuracy and success rate by switching from swap move to slide move algorithm. As the key length increases to $k = 20$, the benefit of using slide moves becomes more significant: swap moves and single letter slide moves can’t achieve a single successful run, but block letter slide move performs very well with an accuracy of 95% and complete success in 90 out of 100 runs.

move	accuracy	no. of successful runs
swap move	0.6550	17
single letter slide move	0.9525	83
block letter slide move	0.9587	90

Table 14: Comparison of results for attacks on transposition ciphers of key length 10 when the proposals are swap moves, single-letter slide moves, and block-letter slide moves. Each run uses 1000 iterations and 1000 characters of cipher text, with scaling parameter 1.

move	accuracy	no. of successful runs
swap move	0.4383	0
single letter slide move	0.6333	0
block letter slide move	0.8961	59

Table 15: Comparison of results for attacks on transposition ciphers of key length 20 when the proposals are swap moves, single-letter slide moves, and block-letter slide moves. Each run uses 1000 iterations and 1000 characters of cipher text, with scaling parameter 1.

4.2 Optimizing the Scaling Parameter

We again experimented with different choices of the scaling (inverse temperature) parameter. We found that small choices of this parameter lead to poor performance, while values equal to or greater than 1 lead to approximately equally good performance (and acceptance rates around 0.44), see Table 16. We again choose our final scaling parameter to be 1 since that gives the highest accuracy and produces the highest percentage of successful runs.

power	acceptance rate	accuracy	no. of successful runs
0.01	0.975	0.013	0
0.1	0.690	0.084	0
1	0.439	0.9694	87
10	0.437	0.9389	77
1,000	0.431	0.9489	81
100,000	0.440	0.9489	81

Table 16: Results of attacks on transposition ciphers with key length 20, for different choices of the scaling parameter. Each run uses 1,000 iterations, an 1000 characters of cipher text.

Recall that with substitution ciphers, we improved our attack algorithm by remembering whichever key gave the highest score function. For transposition ciphers, this turns out to be unimportant, since with scaling parameter 1 it is very rare for the chain to ever jump from higher to lower score functions. Indeed, in each of our runs above, the last key was also the key which gives the highest log score. However, we still choose to remember the highest log score, since this doesn't cost much overhead and it guarantees that we will always return the key which gave the best result.

4.3 Amount of Cipher Text Needed

We next investigated the extent to which the accuracy and success rate of the algorithm are affected by the length of the cipher text used to compute the score function. This question is more relevant here than for substitution ciphers, since now the speedup from using less cipher text is much more significant.

We found (Table 17) that we certainly need at least 500 characters of cipher text to break a transposition cipher with key length 20 in 10,000 iterations. More precisely, it appears that the 2000 characters of cipher text is the best choice, since that leads to very high accuracy (over 99%) and successful runs (95%), and using more cipher text requires significantly more time to process but leads to very marginal benefits.

cipher text length	accuracy	no. of successful runs	duration (in seconds)
100	0.0789	0	0.4640
200	0.2844	0	0.5093
500	0.9250	72	0.6526
1,000	0.9694	87	0.8856
2,000	0.9933	95	1.3500
5,000	0.9917	96	2.7557
10,000	0.9861	92	5.079

Table 17: Results of attacks on transposition ciphers with key length 20, when using different numbers of characters of cipher text. Each run uses 10,000 iterations, with scaling parameter 1.

We also found that using a different section of cipher text of the same length leads to very similar results (Table 18), showing a certain stability of this approach.

cipher text starting position	accuracy	no. of successful runs
830080	0.9917	97
254640	0.9933	96
568780	0.9906	96
634220	0.9972	98
366660	0.9928	96

Table 18: Results of attacks on transposition ciphers with key length 20, when using 2000 characters of cipher text starting from different positions in the text. Each run uses 10,000 iterations, with scaling parameter 1.

4.4 Number of Iterations

With the above optimal block slide move proposal, and scaling parameter 1, and cipher text size 2000 for transposition key length 20, we next investigate the extent to which we can increase the accuracy by running more iterations. Table 19 shows that the accuracy and success rate increase steadily as we use longer runs up to about 20,000 iterations, after which there is little further gain. So, we choose 20,000 iterations as our optimal run length.

no. of iterations	accuracy	no. of successful runs
1,000	0.6878	3
2,000	0.7944	17
5,000	0.9544	73
10,000	0.9911	95
20,000	1.0000	100
50,000	1.0000	100

Table 19: Results of attacks on transposition ciphers with key length 20, when using 2000 characters of cipher text and scaling parameter 1, for different numbers of iterations.

More generally, for different transposition key lengths, we tried increasing the number of iterations to get a reasonable accuracy rate. Our results are presented in Table 20.

key length	no. of iterations	accuracy	no. of successful runs	duration (in seconds)
10	2,000	0.9962	97	0.3424
20	10,000	0.9911	95	1.3736
30	50,000	0.9957	98	6.4034
40	50,000	0.9613	70	6.3969
50	100,000	0.9648	68	12.8497

Table 20: Results of attacks on transposition ciphers with different key lengths, when using 2000 characters of cipher text and scaling parameter 1, when using different numbers of iterations.

On the other hand, increasing the number of iterations alone is not sufficient to overcome all difficulties. We illustrate this more precisely using just 1000 characters of cipher text. In this case, we already know that the accuracy will not be great. However, it is also true that this accuracy will not increase very quickly as we run more iterations (Table 21).

no. of iterations	accuracy	no. of successful runs
1,000	0.6333	1
2,000	0.7583	12
5,000	0.8961	59
10,000	0.9694	87
20,000	0.9872	95
50,000	0.9856	93

Table 21: Results of attacks on transposition ciphers with key length 20, when using just 1000 characters of cipher text for different numbers of iterations (still with scaling parameter 1), indicating poor performance even after many iterations.

4.5 Independent Repetitions

As with substitution ciphers, we may wish to use several independent repeated runs of our attack, and return as our final answer whichever of the results has the largest score function (from (2) computed using the *entire* cipher text).

To consider the extent to which multiple independent repetitions might help with this problem, we compare a single long run of 50,000 iterations, with 5 repetitions of a run of 10,000 which returns the decryption key which gives the highest score. Our results are presented in Table 22. We see that multiple shorter runs are consistently better than one very long run, increasing the percentage of successful runs from 93% to 100%. This makes sense since we have already seen that the success rate with 1000 characters of cipher text for 1 run of 10,000 iterations is about 87%. So, if we use 5 independent such runs, then the probability not getting a correct answer is only $(1 - 0.87)^5 = 0.0037\%$ which is very low.

no. of repetitions x no. of iterations	accuracy	no. of successful runs
20 x 2,500	0.9972	98
10 x 5,000	1.0000	100
5 x 10,000	1.0000	100
2 x 25,000	1.0000	100
1 x 50,000	0.9856	93

Table 22: Results of attacks on transposition ciphers with key length 20, when using just 1000 characters of cipher text (still with scaling parameter 1), when dividing up the 50,000 total iterations into multiple independent repetitions.

4.6 Attack for Transposition Ciphers – Final Version

Based on the above investigations, we propose the following as our final algorithm for the attack of the transposition cipher.

- Randomly select 2000 cipher text from available cipher text
- Attack the selected cipher text with the bi-gram score function, using block slide proposal moves and parameter value 1, for 10,000 iterations (when the key length equals 20).
- Apply the decryption key to the full cipher text and calculate the log score for full decrypted text.
- Repeat the above procedure 5 times; the final result is the repetition which gives the highest log score.

To investigate how our program works, we apply the method to different cipher text and reference text. The results are presented in Table 23, which shows that the results are very good, leading to perfect runs in every case.

cipher text	reference text	accuracy	no. of successful runs
<i>Oliver Twist</i>	<i>War and Peace</i>	1.0000	100
<i>Pride and Prejudice</i>	<i>War and Peace</i>	1.0000	100
<i>Ice Hockey (Wikipedia Page)</i>	<i>War and Peace</i>	1.0000	100
<i>Pride and Prejudice</i>	<i>Oliver Twist</i>	1.0000	100
<i>War and Peace</i>	<i>Oliver Twist</i>	1.0000	100
<i>Ice Hockey (Wikipedia Page)</i>	<i>Oliver Twist</i>	1.0000	100
<i>War and Peace</i>	<i>Pride and Prejudice</i>	1.0000	100
<i>Ice Hockey (Wikipedia Page)</i>	<i>Pride and Prejudice</i>	1.0000	100
<i>Oliver Twist</i>	<i>Pride and Prejudice</i>	1.0000	100
<i>Pride and Prejudice</i>	<i>Ice Hockey (Wikipedia Page)</i>	1.0000	100
<i>War and Peace</i>	<i>Ice Hockey (Wikipedia Page)</i>	1.0000	100
<i>Oliver Twist</i>	<i>Ice Hockey (Wikipedia Page)</i>	1.0000	100

Table 23: Results of our final attack on transposition ciphers, with key length 20, when using 2000 characters of cipher text and scaling parameter 1, with 5 repetitions of 10,000 iterations each. A perfect result is obtained in every case.

5 Attacks on Substitution-Transposition Ciphers

Substitution-Transposition ciphers have 2 different keys. First the letters are switched using a substitution cipher. Then, the characters are moved around using a transposition cipher. The length of the substitution key is 26 as usual. The length of the transposition key, k , is variable as in the previous section. To break Substitution-Transposition ciphers, we shall reuse and combine the MCMC attack algorithms previously developed for the substitution cipher and the transposition cipher. In particular, we shall again attempt to maximize the same score function (2).

5.1 First Attempt

For illustrative purposes, we first try to break a substitution-transposition cipher with transposition key length 10.

As our initial algorithm, we simply combine our two previous algorithms directly, by first running a bi-gram attack as for a substitution cipher, and then running a bi-gram attack as for a transposition cipher. We use the optimal values of parameters from our previous attacks: specifically, we use 2000 randomly-chosen cipher text characters, with scaling parameter 1, and run 10,000 iterations of the bi-gram attack on a substitution cipher, followed by 2,000 iterations of the bi-gram attack on a transposition cipher.

The results are presented in Table 24. We see that we are not able to get good results, even by increasing the number of iterations. And, switching the sequence of the two attacks also does not help.

no. of iterations (substitution/transposition)	accuracy	no. of successful runs
10,000/2,000	0.0600	3
10,000/5,000	0.0587	3
10,000/10,000	0.0338	2

Table 24: Results of our first attack on substitution-transposition ciphers with transposition key length 10, using 2000 characters of cipher text, with a bi-gram substitution cipher attack followed by a bi-gram transposition cipher attack.

Of course, it is not surprising that these results are poor. The basic problem is that our first attack is attempting to break a substitution cipher, but it is working with text which also had an unknown transposition applied. Thus, there is no particular reason that the pair frequencies of the transposed text should in any way match those of the reference text. So, the first attack is doomed from the start. And, of course, if the first attack fails completely, then the second attack is similarly handicapped.

5.2 Multiple Cycles

We have previously seen that with MCMC cipher attacks, sometimes several shorter runs are better than one longer run. Inspired by this, we let our algorithm run for several cycles (Table 25). Each cycle consists of a bi-gram attack on substitution cipher, followed by a bi-gram attack on transposition cipher. We use the result of one cycle as the starting point for next one. Table 25 shows that our results do improve upon increasing the number of cycles. However, the improvement

does not continue much beyond 3 cycles: we get just 63% accuracy on average and 62 success out of 100 runs even after 10 cycles.

cycles	accuracy	no. of successful runs	duration (in seconds)
1	0.0600	3	0.7446
2	0.3713	35	1.4919
3	0.5462	52	2.5222
5	0.5250	52	3.6850
10	0.6300	62	7.4400

Table 25: Results of attacks on substitution-transposition ciphers with transposition key length 10, for various numbers of cycles. Each cycle uses 2000 characters of cipher text, and consists of a 10,000-iteration bi-gram substitution cipher attack followed by a 2,000-iteration bi-gram transposition cipher attack.

It seems that, even with multiple cycles, the problem remains that if the substitution attack fails massively, then the transposition attack has little chance of success, and vice-versa. This is a sort of “chicken and the egg” problem: if *one* of the attacks were successful, or even nearly successful, then the other attack would perform well and the problem would quickly be solved. The question remains, how can we get initial near-success? We consider that next.

5.3 Using a Uni-gram Attack for Initialization

Recall that the standard uni-gram attack on substitution ciphers is quick and simple, but it is not terribly accurate, i.e. it tends to only partially reveal the original text. We can make use of this in the attack for breaking the substitution-transposition cipher. Even though the uni-gram attack does not reveal all the letters, it can quickly provide a very good starting point.

Inspired by this, we modify our algorithm to *first* run the uni-gram attack for substitution ciphers, and *then* run multiple cycles of bi-gram attacks. The results are presented in Table 26 (for transposition key length 10, up to 2 cycles) and Table 27 (for transposition key length 20, up to 3 cycles), and indicate very high success rates in both cases.

cycles	accuracy	no. of successful runs	duration (in seconds)
1	0.7937	68	0.8369
2	1.0000	100	1.5800

Table 26: Results of attacks on substitution-transposition ciphers with transposition key length 10, after initializing with a uni-gram attack, for various numbers of cycles. Each cycle uses 2000 characters of cipher text, and consists of a 10,000-iteration bi-gram substitution cipher attack followed by a 2,000-iteration bi-gram transposition cipher attack.

cycles	accuracy	no. of successful runs	duration (in seconds)
1	0.2394	0	1.8440
2	0.9133	82	3.6105
3	0.9906	98	5.3200

Table 27: Results of attacks on substitution-transposition ciphers with transposition key length 20, after initializing with a uni-gram attack, for various numbers of cycles. Each cycle uses 2000 characters of cipher text, and consists of a 10,000-iteration bi-gram substitution cipher attack followed by a 2,000-iteration bi-gram transposition cipher attack.

5.4 Attack for Substitution-Transposition Ciphers – Final Version

Putting the above together, we propose the following MCMC algorithm to attack the substitution-transposition cipher.

- First, run the uni-gram attack for substitution cipher on the original cipher text.
- Then, for several cycles (3, for key length 20):
 - Run the bi-gram attack for transposition cipher on the resulting text, for an appropriate number of iterations (10,000, for key length 20).
 - Run the bi-gram attack for substitution cipher on the resulting text, for an appropriate number of iterations (10,000, for key length 20).

We ran this final algorithm on randomly-generated substitution-transposition ciphers with transposition key length 20, with different combinations of cipher text and reference. The overall results are presented in Table 28. Once again, we find that runs using the short and modern text *Ice Hockey (Wikipedia Page)* are worse than using other text. However, for experiments using the classic novels, the accuracy is always above 80% and the number of successful runs is always above 70, indicating quite good performance for this challenging problem.

cipher text	reference text	accuracy	no. of successful runs
<i>Oliver Twist</i>	<i>War and Peace</i>	0.8894	86
<i>Pride and Prejudice</i>	<i>War and Peace</i>	0.8433	80
<i>Ice Hockey (Wikipedia Page)</i>	<i>War and Peace</i>	0.6811	58
<i>Pride and Prejudice</i>	<i>Oliver Twist</i>	0.9100	89
<i>War and Peace</i>	<i>Oliver Twist</i>	0.8367	78
<i>Ice Hockey (Wikipedia Page)</i>	<i>Oliver Twist</i>	0.7211	62
<i>War and Peace</i>	<i>Pride and Prejudice</i>	0.8189	74
<i>Ice Hockey (Wikipedia Page)</i>	<i>Pride and Prejudice</i>	0.6811	56
<i>Oliver Twist</i>	<i>Pride and Prejudice</i>	0.8244	81
<i>Pride and Prejudice</i>	<i>Ice Hockey (Wikipedia Page)</i>	0.7961	74
<i>War and Peace</i>	<i>Ice Hockey (Wikipedia Page)</i>	0.6761	64
<i>Oliver Twist</i>	<i>Ice Hockey (Wikipedia Page)</i>	0.7778	71

Table 28: Results of our final attacks on substitution-transposition ciphers with transposition key length 20, after initializing with a uni-gram attack, for various texts. Each cycle uses 2000 characters of cipher text, and consists of a 10,000-iteration bi-gram substitution cipher attack followed by a 10,000-iteration bi-gram transposition cipher attack.

Further experimentation using the classic novels indicates that with enough iterations and cycles, the accuracy and success rates remain quite high even with transposition keys up to size 40 (Table 29).

transposition key length	no. of iterations (subst./trans.)	cycles	accuracy	no. of successful runs	duration
10	10,000/2,000	3	1.0000	100	3.93
20	10,000/10,000	3	0.8894	86	5.32
30	10,000/50,000	5	0.8618	85	34.08
40	10,000/100,000	5	0.7645	73	65.51

Table 29: Results of our final attacks on substitution-transposition ciphers with various transposition key lengths, for *Oliver Twist*, using *War and Peace* as the reference text. Each attack first initializes with a uni-gram attack, and then repeats the specified number of cycles. Each cycle uses 2000 characters of cipher text, and consists of a 10,000-iteration bi-gram substitution cipher attack followed by a 10,000-iteration bi-gram transposition cipher attack.

Overall this indicates quite good performance, even for the difficult substitution-transposition cipher, achieving accuracies and success rates above 70% even with key length 40.

6 Summary

In this paper, we successfully applied MCMC algorithms to break substitution ciphers, transposition ciphers, and even substitution-transposition ciphers. The attacks are based on the frequency analysis of the cipher text together with a reference text, and primarily consist of bi-gram attacks.

We have experimented significantly with such issues as number of MCMC iterations, scaling (inverse temperature) parameter, amount of cipher text to use, number of independent repetitions, swap moves versus slide moves versus block-slide moves, etc., in an attempt to optimize our choices. For substitution-transposition ciphers, we required additional innovations such as repeatedly cycling between substitution-type and transposition-type attacks, and using a simple uni-gram substitution attack as an initialization point.

Overall, our simulations indicate good success of our algorithms. In particular, we are able to break the simple substitution-transposition cipher with accuracy and success rates above 70%, even with transposition key length up to 40. This indicates the potential for MCMC algorithms to provide significant help in deciphering challenging encryptions.

References

- [1] S. Conner (2003), *Simulation and solving substitution codes*, Master’s thesis, Department of Statistics, University of Warwick.
- [2] P. Diaconis (2008), *The Markov Chain Monte Carlo Revolution*. Bull. Amer. Math. Soc., Nov. 2008.

- [3] A. Dimovski and D. Gligoroski (2003), *Attacks on the Transposition Ciphers Using Optimization Heuristics*. Proceedings of the XXXVIII International Scientific Conference on Information, Communication & Energy Systems & Technologies, Heron Press, Birmingham, U.K.
- [4] P. Garg (2009), *Cryptanalysis of SDES via Evolutionary Computation Techniques*, IJCSIS **1(1)**, May 2009.
- [5] W.R. Gilks, S. Richardson, and D.J. Spiegelhalter, ed. (1996), *Markov chain Monte Carlo in practice*. Chapman and Hall, London.
- [6] Ice Hockey (Wikipedia Page). http://en.wikipedia.org/wiki/Ice_hockey
- [7] N. Metropolis, A. Rosenbluth, M. Rosenbluth, A. Teller, and E. Teller (1953), *Equations of state calculations by fast computing machines*. J. Chem. Phys. **21**, 1087–1091.
- [8] K.S. Ooi and B.C. Vito (2002), *Cryptanalysis of S-DES*. Available at: <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.79.6617>
- [9] Project Gutenberg. <http://www.gutenberg.org/>
- [10] G.O. Roberts and J.S. Rosenthal (2004), *General state space Markov chains and MCMC algorithms*. Prob. Surv. **1**, 20–71.
- [11] B. Schneier (1996), *Applied Cryptography, Second Edition*. John Wiley & Sons, New York.
- [12] C. E. Shannon (1949), *Communication Theory of Secrecy Systems*. Bell System Technical Journal **28(4)**, 656–715.