Journal of Information Hiding and Multimedia Signal Processing ©2012 ISSN 2073-4212
Ubiquitous International Volume **3**, Number **1**, January **2012**

Human- Versus Computer-Generated Text-Based Steganography: Real-World Tests of Two Algorithms

Michael Grosvald
Department of Linguistics
Department of Linguistics University of California, Davis
Davis, CA 95616
mgrosvald@ucdavis.edu

C. Orhan Orgun
Department of Linguistics
Department of Linguistics University of California, Davis
Davis, CA 95616
ocorgun@ucdavis.edu

Received April 2010; revised November 2011

Abstract. *Tests of encryption procedures using lexical (text-based) steganography are typically designed to detect statistical differences between encrypted texts and natural human language. While such a method of attack is undoubtedly useful in many contexts, it overlooks the possibility that an encrypted text might go undetected in statistical testing but nonetheless appear unnatural to a human reader. In this paper we investigate the performance of two encryption algorithms in real-world tests using a human-based method of attack. One procedure is an automated word-replacement algorithm, while the other combines an automated algorithm with human input. A quantitative and qualitative analysis of the results shows that the automated procedure produced texts that are likely to attract unwanted attention from human readers, while the procedure using human input did not.*

Keywords: steganography, lexical, text-based, Internet

1. **Introduction.** Text-based steganography has received considerable attention in the recent literature, and a wide variety of approaches has been developed and tested ([1], [3], [4], [5], [12], and [14]). Tests of text-based steganography systems generally focus on detecting the encrypted text's statistical differences from natural human language. Among the methods used are token frequency analysis and bigram (or more generally, n-gram) frequencies ([6], [7], [8], [13]). While detection under such statistical analysis is indeed a potential weakness, the statistical method of attack overlooks some important points, discussed in the following paragraphs.

Submitting a text to statistical scrutiny presupposes that there was already a suspicion of secrecy on the part of the analyst. While it is conceivable that there might be social or political contexts in which a governing body might routinely submit all texts produced by certain individuals to such analysis, most texts are probably not subject to this level of scrutiny. In many contexts, only texts that were previously found to be suspicious by a human observer will be subjected to statistical analysis. This could happen in two ways. First, the person producing the text is under suspicion and scrutiny. Second, although there is no prior suspicion directed towards the author, the text itself happens to be observed by a human, who finds it suspiciously unnatural. It is this second possibility that we find insufficiently addressed in the existing literature.

Human judgments may be of even deeper concern for another reason. It is conceivable (and in our judgment, likely) that a piece of text that passes every statistical test that has hitherto been devised (perhaps having been engineered precisely for this purpose) might nonetheless appear unnatural to a human reader. The source of unnaturalness could come from any domain of linguistics, from minute syntactic details to the grossest levels of discourse organization and pragmatics. An organization that monitors text traffic for suspicious activity need not limit itself to mass statistical analysis. It might also select a subset of texts (either randomly or based on some prior suspicion) for direct human scrutiny. Texts judged unnatural might then be submitted to more thorough cryptographic analysis.

For the foregoing reasons, we propose to test two different approaches from the current steganography literature for acceptability to human readers. The first encryption procedure, Lunabel ([2], [11]), represents a completely computer-generated approach to text-based steganography; this type of approach is shared by most systems of text-based steganography. The second procedure, Neko ([10]), is an encryption process incorporating human input. Both procedures are summarized below. It is our contention that the Neko approach is more secure from human attack than a completely computer-reliant procedure like Lunabel. To test this assertion, we have conducted two studies testing the acceptability to human readers of texts produced by both encryption procedures.

The first study is a direct quantitative investigation. It is intended to represent a situation in which texts are specifically submitted to human scrutiny for the purposes of detecting unnatural (therefore possibly illicit) texts. For this test, we created encrypted messages using Lunabel (the automated word-replacement method) and Neko (the human-generated text method). We then asked experimental subjects to read these, along with comparable pieces of natural text, and judge these texts for naturalness and clarity on a seven-point scale. The human-generated method (Neko) fared significantly better than the word-replacement method (Lunabel). In fact, there was no statistically significant difference between the acceptability scores of Neko and natural texts.

The second study is a qualitative real-world test of the two encryption systems. The aim of this test is to see how well these two methods do in terms of not arousing suspicion under human scrutiny in a situation where no prior suspicion of the author exists. For this test, we posted messages to high-volume Internet discussion forums and observed other forum participants' reactions to those messages. The human-generated method did not elicit any responses indicating suspicion, while the word-replacement method was noticed several times for unnaturalness of linguistic structure (though never directly suspected of encrypting messages).

2. **Overview of steganography methods used.**

In this section, we present brief summaries of the methods of encryption used by Lunabel and Neko. Detailed algorithms can be found in [2], [11] and [10]. We discuss each encryption procedure separately.

2.1. **Lunabel.**

Lunabel is a word replacement method. It starts with a cover text and replaces words in that cover text with other words according to a particular word substitution scheme. The cover text can be any piece of text. Substitution sets of words are compiled in advance according to certain linguistic criteria such as part of speech, semantic compatibility, frequency of occurrence, and so on. Because no substitution scheme is linguistically perfect, texts generated by Lunabel do not always appear to be natural human language. While different word-replacement methods may have varying degrees of success in choosing mutually substitutable sets of words to target for replacement, the type of unnaturalness we observe here is inherent in the very notion of word replacement.

The following plain text (from [9]) and corresponding Lunabel-generated encryption illustrate this problem of unnaturalness. The Lunabel text differs from the cover text only where word list elements have been replaced in order to code the message. Those word list elements have been highlighted in both the cover text and the Lunabel encryption.

2.1.1. *Natural (Cover) text.* They reigned over the earth for more than 100 million **years** and suddenly, mysteriously disappeared. What caused the demise of this ubiquitous **group** of reptiles which included some of the largest animals to ever walk the planet? One of the great mysteries in science is the extinction of the dinosaurs at the end of the Mesozoic Era some 65 million **years** ago. Who (or more likely what) caused it is unknown and a **subject** of great debate. Dinosaurs appeared at the beginning of the Mesozoic Era and were the dominant **form** of **life** until the end of that era. They lived almost everywhere there was land including Antarctica. We can see their bones in the geological record. The lower stratum of rock contains the earliest and most primitive species of dinosaur, and the upper stratum contains the newer species. Then, suddenly, at a geological strata **line** called the Cretaceous-Tertiary Boundary (often referred to as the K-T Boundary), the dinosaurs disappear. . .

2.1.2. *Lunabel text.* They reigned over the earth for more than 100 million **times** and suddenly, mysteriously disappeared. What caused the demise of this ubiq- uitous **universe** of reptiles which included some of the largest animals to ever walk the planet? One of the great mysteries in science is the extinction of the dinosaurs at the end of the Mesozoic Era some 65 million **things** ago. Who (or more likely what) caused it is unknown and a **characteristic** of great debate. Dinosaurs appeared at the beginning of the Mesozoic Era and were the dominant **sense** of **question** until the end of that era. They lived almost everywhere there was land including Antarctica. We can see their bones in the geological record. The lower stratum of rock contains the earliest and most primitive species of dinosaur, and the upper stratum contains the newer species. Then, suddenly, at a geological strata **context** called the Cretaceous-Tertiary Boundary (often referred to as the K-T Boundary), the dinosaurs disappear. . .

It can be seen that the unnaturalness of expressions can be local, for example *65 million things ago*, where only a time expression in place of things would be locally natural (and anything smaller than a year would still be unnatural in the larger pragmatic contexts: *??over 65 million weeks ago*). On the other hand, the expression *a characteristic of great debate* is perfectly natural in its local structure, but in the larger discourse context, one expects a word like *topic* or *subject*. While further refinement of bigram, trigram, etc., frequencies in word lists may leave some limited room for improvement in local sources of unnaturalness, the more global types of problem are currently beyond reach, since they require not only a grasp of high-level discourse organization, but also vast amounts of world knowledge. From our current vantage point, this would appear to encompass no less than the full problem of artificial intelligence.

2.2. **Neko.** Neko does not use a cover text prior to information hiding. Instead, it uses word lists that are compiled solely based on frequency of occurrence in a given context. For information hiding, Neko generates a sequence of words. Human input is then required to create a text that includes those words in the given order. Because the final text is human-generated, it is (by definition) natural language and is free from algorithmic artifacts. The following example shows a Neko encryption based on the same key words as the previously shown Lunabel text.

2.2.1. *Neko text.* One of the greatest scientific achievements of modern **times** is the great leap in our understanding of the large-scale temporal and spatial structure of the **universe**. Human beings are by nature very curious, and it has often been said that this curiosity is one of the **things** that defines us. This trait that is so **characteristic** of our species has allowed us, by extension, to gain a greater **sense** not only of where we live, but who we are as well. However, one key **question** that science has yet to answer is whether we are alone in this great cosmos. If we find that other beings have also attained consciousness and intelligence as we have, can we still consider ourselves unique? It is in the **context** of such a discussion that one is forced to examine possibilities which at first glance may make us uncomfortable, but which ultimately must be investigated if we are to continue to pursue such scientific endeavors. . .

The only limiting factor in the naturalness of Neko texts is the skill of the person writing them. However, as this factor is also equally present in writing plain text, there is no principled reason for which Neko texts should appear any less natural than ordinary plain text. Nonetheless, the studies we will present are designed to test this expectation.

3. Real-world comparisons of the two encryption methods.

3.1. Real-world Test 1: Questionnaire response data.

3.1.1. *Description of test materials.* In our first study, we conducted a direct assessment of the naturalness of Lunabel and Neko texts as judged by human readers. This simulates a situation in which a piece of text is submitted to human scrutiny because of suspected manipulation. For this test, we used three pieces of natural text:

i) A newspaper article on life events
ii) A discussion of the extinction of dinosaurs
iii) A geological history of Yellowstone National Park

We created Lunabel encryptions using these three pieces as cover text. We also created Neko encryptions based on the same words that were targeted for replacement in the Lunabel encryptions. We thus had three families of texts:

i) Natural texts
ii) Lunabel encryptions of those natural texts
iii) Neko texts

We then collected linguistic judgments from 31 native speakers of English. The speakers were undergraduate students at the University of California, Davis, who were not previously informed about the purpose of the tests. We presented each subject with a piece of natural text, a Lunabel encryption, and a Neko encryption. The order of presentation varied randomly across subjects. Each subject was asked to respond to two questions after reading each piece of text. The questions concerned the *naturalness* and the *clarity* of the texts; these terms were defined to subjects as follows.

In the present context, a text is *natural* if it seems like something a competent user of English might write. On the other hand, a text is *clear* if it is apparent what the intended meaning of the author was. Note that a text can be natural but not clear (e.g. text written by a competent user of English which is nonetheless very dense and technical), or clear but not natural (e.g. a simple story written by someone who is obviously a non-native user of English, but whose errors do not interfere with the reader's understanding of the story).

Keeping the foregoing in mind, subjects were asked to reply to the following two questions about each text that they read, with each question to be answered on a 7-point scale, from 1=least to 7=most:

Q1: How natural did you find this text?
Q2: How clear did you find this text?

In this way, we are able to separate, and quantify, two independent dimensions of acceptability to human readers of each of our three categories of text (plain, Lunabel, and Neko).

3.1.2. *Questionnaire statistical results.* For each question, paired t-tests were conducted to determine significant differences in the mean responses for the three kinds of text (plain, Lunabel, Neko). For Q1, the respective mean scores [and standard deviations] for plain, Lunabel, Neko texts were 5.55[1.34], 4.58[1.43], and 5.42[1.29]. The Lunabel text was scored as significantly less natural than the plain text ($t(30) = 3.32; p < 0.01$), while the Neko and plain text scores were not significantly different ($t(30) = 0.37; n.s.$). Therefore, respondents' scores indicate that they found the human-generated Neko text to be very similar to the plain text in terms of naturalness, while the completely automated Lunabel procedure resulted in text that respondents found significantly less natural.

The results for Q2 were similar. The respective mean scores [and standard deviations] for plain, Lunabel, and Neko texts were 5.81[1.01], 4.90[1.47], and 5.29[1.49]. The Lunabel text was scored as significantly less clear than the plain text ($t(30) = 3.10; p < 0.01$), while the Neko and plain text scores were not significantly different ($t(30) = 1.63; n.s.$).

In this case, the Neko texts do have a numerical score that is noticeably less than that of the plain texts, though as already noted, this difference does not reach statistical significance. Furthermore, as we have defined the terms here, *clarity* is not always a hallmark of human-created text, while we have taken *naturalness* to be so. Therefore, we consider it particularly important that the Neko-produced texts were rated as very similar to the plain texts in terms of naturalness.

3.2. **Real-world Test 2: Internet forum postings.**

3.2.1. *Fictitious contributors and their posts.* The second test is a real-world trial of the feasibility of using steganography for information exchange. In contrast to our first test, where we subjected texts to direct scrutiny for the purposes of a *quantitative* analysis, here our focus was on seeing whether human readers would suspect steganographic texts to be unnatural without any prior expectation that manipulated texts might be involved. Our aim was to determine whether our different categories of text (plain, Lunabel, Neko) might provoke *qualitatively* different reactions on the part of such readers.

For this qualitative evaluation, we chose to post encrypted messages on public discussion forums on the Internet. We chose two Internet newsgroups for this purpose. Both of these newsgroups have high volumes of daily contribution and heated discussions between people who hold widely varying opinions on the newsgroups'topics of discussion. We selected these newsgroups for several reasons. First, the fact that they are very active meant that our posts would be visible to a large number of people, and therefore open to plenty of scrutiny. On the other hand, since each newsgroup had a wide variety of individual styles and language backgrounds, it was likely that departures from what is considered standard English might be less surprising to forum participants. On the other hand, some contributors specifically address others' perceived lack of command of English, and this opened up the possibility that any unexpected linguistic structure in our posts might be noticed and commented on.

Some of our posts elicited responses from a large number of individuals; furthermore, all of our posts and all responses to them were public and visible not only to newsgroup participants, but to anyone on the Internet. Hence, any quotes taken from such postings would be searchable, and could yield other information about participants up to and possibly including their real identities. For this reason, we are not including specific information about the newsgroups or actual quotes from our posts and the responses they elicited, in order to protect the privacy of newsgroup participants who might prefer not to have attention drawn to themselves as a result of being discussed or quoted directly in this paper.

As described earlier, both encryption methods tested here require the preparation of suitable word lists (for substitution in a cover text in the case of Lunabel, and as the seed for human-generated text in the case of Neko). To prepare the lists for this part of our study, we analyzed words appearing in previous newsgroup postings for frequency, so that word lists contained items seen at comparably frequent rates in newsgroup postings. For each of the two newsgroups, 40 word lists (of 16 words each) were prepared, and Lunabel and Neko encryptions used the same lists. The word lists were of the following categories: nouns, verbs, adjectives and adverbs. To insure agreement with any co-occurring articles "a" or "an," each noun or adjective list contained items agreeing in either beginning with a consonant or vowel sound.

For each newsgroup, we created a pair of fictitious users. One user (henceforth "L") used Lunabel encryptions, while the other (henceforth "N") used Neko encryptions. In one newsgroup, user L was representative of the point of view stated in the newsgroup's offcial mission, while N was antagonistic to this point of view. In the other newsgroup, L was antagonistic and N was friendly. This was done in order to (a) elicit attention from regular newsgroup participants and (b) to balance the roles of Lunabel and Neko across both groups.

On each of the two newsgroups, the "antagonistic" user initiated a total of three discussion threads over the period of one to two weeks; these six threads are the subject of this analysis. Both fictitious users replied to the other user's posts as well as posting occasional replies in independent discussion threads initiated by regular forum users. Encrypted texts were only posted in discussions initiated by our fictitious users; in each of the six threads, the first posting (by the "antagonistic" user) and a later reply to that posting (by the "friendly" user) consisted of encrypted text. As intended, our encrypted posts sometimes elicited replies from regular forum users. We examined these replies for any indications of suspicion of text manipulation. The results are summarized in the next section.

3.2.2. *Summary of forum user reactions.* For this analysis, we compiled reactions from regular forum users to the posts made by our fictitious users. Most regular users replied to the content of the posts, but a smaller number made comments concerning language use in L's posts. Notably, none of the posts by N elicited any reaction concerning the language used. This was so for N's postings on both newsgroups (i.e. in both the "friendly" and "antagonistic" roles). The hidden message in each case was a confirmation of the day and time for a meeting between the two authors, followed by a brief instruction to email us if anyone were to break the code. A typical encrypted text might read as follows:

```
T11:30?
email [email address] for prize.
END
```

We categorized these reactions along a dimension of "level of suspicion" as follows:

I: Reponses with reference solely to the content of one of our postings, without any remarks on linguistic structure

II: Reponses noting the use of non-standard English in one of our postings

III: Reponses expressing suspicion that one of our postings might be computer-generated or computer-modified

IV: Reponses expressing suspicion that one of our postings might contain an encrypted message

Table 1 summarizes reactions to the threads initiated by both fictitious users. These numbers conflate responses to the initial posts and to other posts within the same thread. All replies of types II, III and IV were given to L postings and never to the N responses to L within the L-initiated threads.

The last category, IV, was of particular interest for this study, as any such response would show clearly that encrypted texts were vulnerable to suspicion that steganography was being employed by one of our fictitious users. While this level of arousal was never expressed by any regular user, there were two instances of regular users guessing that L might be using computer-modified text in various ways (e.g., the output of a computerized word game, or a machine translation). As shown in Table 1, posts made by N never elicited any evidence of such suspicion on the part of regular users.

Table 1. Categorization of replies to Lunabel and Neko Internet postings

Lunabel-initiated thread	Unique authors (regular users)	I	II	III	IV
1	14	25	0	1	0
2	24	90	5	1	0
3	13	11	10	0	0

Neko-initiated thread	Unique authors (regular users)	I	II	III	IV
1	3	10	0	0	0
2	0	0	0	0	0
3	1	4	0	0	0

4. **Discussion.** The outcomes of both tests show that the human-generated (Neko) texts fared better than those produced by a fully automatic procedure (Lunabel). In the first test, human respondents reading plain text as well as Lunabel-generated and Neko-generated texts found the Lunabel texts to be significantly less natural and less clear than plain text, while there was no significant difference in scoring for either naturalness or clarity between Neko text and plain text. The results of the second test bolster those findings, and they also provide a qualitative complement to the quantitative outcomes observed in the first test.

The data summarized in Table 1 show clearly that Lunabel generated plenty of direct replies but Neko did not. In some cases, the text generated by Lunabel was noticeably non-standard, and resulted in a number of replies such as “Do you speak English?” or “Did you even read what you wrote?” Most importantly, Neko postings drew no comments on language use. We can conclude from this that Neko passes an important real-world test: Neko texts that hide a message are readily accepted as natural language by people who read them closely enough to reply to their content.

Lunabel had a more mixed rate of success. It will be noted that the majority of responses to Lunabel posts were on the content of the messages. That is, most replies seemed to be predicated on the acceptance of Lunabel posts as real language. However, there were a considerable number of posts pointing out what appeared to be non-native or non-standard English in Lunabel posts.

As can be seen in column IV in Table 1, there were no comments made by any regular user that either Lunabel or Neko messages might contain encrypted messages. However, the two replies reported in column III are of particular importance. Those were made by individuals who suspected that a computer was somehow used in generating Lunabel posts. One of these suggested that Lunabel might be using a web-based translation service. The other is more troublesome. This reply suggested that Lunabel might be using something akin to Mad Libs®. Mad Libs is a language game that provides a text with certain words missing from it. Parts of speech are specified for the missing words and the player supplies random words in the appropriate part of speech. The missing words are then filled in accordingly and the resulting text is read for humorous effect. It will be recognized that the procedure is very similar to that used in Lunabel's substitution scheme. It is therefore fair to consider the unnaturalness of Lunabel's output a potential liability in the real world.

This in turn points to an additional advantage of encryption algorithms incorporating human input. In the preceding tests, we used the same word lists to generate texts using both algorithms in order to evaluate Lunabel and Neko under comparable conditions. However, it is only the automated word substitution method that requires that relatively infrequent words be used, as the use of highly frequent lexical items like function words (e.g. "the," "a," "from") would result in texts that would be much too suspicious. So the use of content words in these tests may if anything underestimate Neko's advantages, as a human generating a text can very easily work common words like "the" into a text in a natural fashion. The resulting texts can also be correspondingly shorter.

In fact, with Neko one option is to use only high-frequency word lists. Then, the user composing Neko text is free to write about essentially any topic, using content that is appropriate to any forum where such messages will be placed. Taking the idea even further, using Neko one can use just one word list containing very high-frequency words like "from," "is" or "there." To demonstrate this advantage of Neko, we are using this entire paragraph to code a short message ("NekoWins") with such a word list. One can see how this was accomplished, as follows.

To convert from plain text to a Neko message coding that text, follow these steps in sequence, in accordance with the procedure described in [10]. Note that the sequence is reversible, in the case where one starts with a Neko-encrypted text and wishes to decrypt it to obtain the intended message (i.e. the plain text).

1) Start with plain text: NekoWins
2) Convert into ASCII codes and extract two hexadecimal digits from each:
   N = ASCII code 78 = 16 * **4** + **14**
   e = ASCII code 101 = 16 * **6** + **5**
   k = ASCII code 107 = 16 * **6** + **11**
   o = ASCII code 111 = 16 * **6** + **15**
   W = ASCII code 87 = 16 * **5** + **7**
   i = ASCII code 105 = 16 * **6** + **9**

n = ASCII code 110 = 16 * **6 + 14**
s = ASCII code 115 = 16 * **7 + 3**

3) Put the resulting numbers in sequence (note that each character in the plain text corresponds to two digits in the resulting sequence):
4, 14; 6, 5; 6, 11; 6, 15; 5, 7; 6, 9; 6, 14; 7, 3

4) Use the Neko word lists to replace each digit with a word. For this example, we have used just one word list: [**over, they, here, as, in, the, to, can, those, of, on, that, go, much, with, be**]. Note that the indexing goes from 0 to 15 rather than 1 to 16. Note also that in the present example, the words indexed by 0, 1, 2, 8, 10, 12 and 13 happen not to be used in generating the encrypted text. Retrieving the words in the just-given word list that are indexed by the sequence of numbers obtained in step (3), we obtain the corresponding sequence of words:
in, with; to, the; to, that; to, be; the, can; to, of; to, with; can, as

5) Finally, compose a suitable text around the obtained word sequence, making sure to avoid inserting any additional Neko word list words (the use of an editor that highlights word list words is convenient for this). In this example, the result is the paragraph that introduced this topic while simultaneously carrying the coded message.

> “**In** fact, **with** Neko one option is **to** use only high-frequency word lists. Then, **the** user composing Neko text is free **to** write about essentially any topic, using content **that** is appropriate **to** any forum where such messages will **be** placed. Taking **the** idea even further, using Neko one **can** use just one word list containing very high-frequency words like “from,” “is” or “there.” **To** demonstrate this advantage **of** Neko, we are using this entire paragraph **to** code a short message (“NekoWins”) **with** such a word list. One **can** see how this was accomplished, **as** follows.”

5. **Conclusion.** We tested a cover text substitution method of steganography (Lunabel) and an approach that generates a sequence of words and calls for a human to create a custom text that includes them (Neko). We used two assessment techniques, both of which used a human-centered method of attack.

In the first test, human readers were given plain, Lunabel and Neko texts and asked to rate each for naturalness and clarity. While the Neko texts scored comparably to the plain texts on both measures, the Lunabel texts scored significantly lower than plain texts for both naturalness and clarity.

In the second test, we posted messages generated by these two encryption methods to Internet discussion groups and tallied the responses of regular users. Messages generated by both encryption methods were largely accepted, and both elicited replies by other users to the surface content of the postings rather than the hidden messages. However, a considerable number of users pointed out non-standard language use in Lunabel-generated texts. In addition, two users suspected that a computer might be used in generating them. By contrast, Neko-generated texts were uniformly accepted as ordinary language.

The fact that Neko uses human input to generate its variant of the cover text is of course an important factor in its success. However, the main advantage of the human-generated encryption approach is that cover text generation is carried out after the information hiding stage. This allows the final text to remain natural.

Acknowledgements. Department of Linguistics, University of California, Davis, Davis, CA 95616, mgrosvald@ucdavis.edu, ocorgun@ucdavis.edu

REFERENCES

[1] R. Bergmair, *Towards Linguistic Steganography: A Systematic Investigation of Approaches, Systems, and Issues*, Technical Reports, University of Derby, 2004.

[2] V. Chand and C. O. Orgun, Exploiting linguistic features in lexical steganography: design and proof-of-concept implementation, *Proc. of the 39th IEEE Annual Hawaii International Conference on System Sciences*, vol. 6, pp. 126b-136b, 2006.

[3] M. Chapman and G. Davida, Hiding the hidden: A software system for concealing ciphertext in innocuous text, *Proc. of the International Conference on Information and Communications Security*, vol. 133, pp. 335–345, 1997.

[4] M. Chapman and G. Davida, Plausible deniability using automated linguistic stegonagraphy, *Proc. of International Conference on Infrastructure Security*, pp. 276-287, 2002.

[5] M. Chapman, et al., A practical and effective approach to large-scale automated linguistic steganography, *Proc. of the Information Security Conference*, pp. 156-165, 2001.

[6] Z. L. Chen, L. S. Huang, Z. S. Yu, L. J. Li, and W. Yang, A statistical algorithm for linguistic steganography detection based on distribution of words, *Proc. of the 3rd International Conference on Availability Reliability and Security*, pp. 558-563, 2008.

[7] Z. L. Chen, L. S. Huang, Z. S. Yu, L. J. Li, and W. Yang, Effective linguistic steganography detection, *Proc. of the IEEE 8th International Conference on Computer and Information Technology*, pp. 224-229, 2008.

[8] Z. L. Chen, L. S. Huang, Z. S. Yu, L. J. Li, and W. Yang, Linguistic dteganography detection using statistical characteristics of correlations between words, *Proc. of the 10th International Conference on Information Hiding*, pp. 224-235, 2008.

[9] Krystek Lee, Who dunnit to the dinosaurs, *The Museum of Unnatural History*, http://www.unmuseum.org/deaddino.htm.

[10] M. Grosvald and C. O. Orgun, Free from the cover text: a human-generated natural language approach to text-based steganography, *Journal of Information Hiding and Multimedia Signal Processing*, vol. 2, no. 2, pp. 133-141, 2011.

[11] C. O. Orgun and V. Chand, The human attack in linguistic steganography, *Handbook of Research on Social and Organizational Liabilities in Information Security*, IGI Global, New York, USA, pp. 380-397, 2009.

[12] M. Shirali-Shahreza, et al., Text steganography in SMS, *Proc. of International Conference on Convergence Information Technology*, vol. 21, pp. 2260-2265, 2007.

[13] C. M. Taskiran, U. Topkara, M. Topkara and E. Delp, Attacks on lexical natural language steganography systems, *Proc. of the SPIE International Conference on Security, Steganography, and Watermarking of Multimedia Contents*, pp. 97-105, 2006.

[14] Keith Winstein, The tyrannosaurus Lex system, http://alumni.imsa.edu/ keithw/tlex/.