“Break Our Steganographic System” — the ins and outs of organizing BOSS

Patrick Bas¹, Tomáš Filler², and Tomáš Pevný³

¹ CNRS - LAGIS, Lille, France
patrick.bas@ec-lille.fr

² State University of New York at Binghamton, NY, USA
tomas.filler@gmail.com

³ Czech Technical University in Prague, Czech Republic
pevna@gmail.com

Abstract. This paper summarizes the first international challenge on steganalysis called BOSS (an acronym for *Break Our Steganographic System*). We explain the motivations behind the organization of the contest, its rules together with reasons for them, and the steganographic algorithm developed for the contest. Since the image databases created for the contest significantly influenced the development of the contest, they are described in a great detail. Paper also presents detailed analysis of results submitted to the challenge. One of the main difficulty the participants had to deal with was the discrepancy between training and testing source of images – the so-called cover-source mismatch, which forced the participants to design steganalyzers robust w.r.t. a specific source of images. We also point to other practical issues related to designing steganographic systems and give several suggestions for future contests in steganalysis.

1 BOSS: Break Our Steganographic System

During the years 2005 and 2007, the data-hiding community supported by the European Network of Excellence in Cryptology (ECRYPT) launched two watermarking challenges, BOWS [13] and BOWS-2 [1] (abbreviations of *Break Our Watermarking System*). The purpose of participants of both challenges was to break watermarking systems under different scenarios. The purpose of organizers was not only to assess the robustness and the security of different watermarking schemes in the environment similar to real application, but to increase the interest in watermarking and to boost the research progress within the field. Both watermarking contests showed to be popular (BOWS/BOWS2 played more than 300/150 domains and 10/15 participants respectively were ranked), and novel approaches towards breaking watermarking systems were derived during them. This, combined with a thrill associated with organization and participation, inspired us to organize the BOSS (Break Our Steganographic System) challenge.

The most important motivation for the contest was to investigate whether content-adaptive steganography improves steganographic security for empirical

covers. For the purpose of this contest, a new spatial-domain content-adaptive algorithm called HUGO (Highly Undetectable steGO) was invented [12]. The fact that in adaptive steganography the selection channel (placement of embedding changes) is publicly known, albeit in a probabilistic form, could in theory be exploited by an attacker. Adaptive schemes introduce more embedding changes than non-adaptive schemes because some pixels are almost forbidden from being modified, which causes an adaptive scheme to embed with a larger change rate than a non-adaptive one. On the other hand, the changes are driven to hard-to-model regions, because the change rate is not an appropriate measure of statistical detectability as it puts the same weight to all pixels. As compared by the state-of-the-art available in mid 2010, HUGO was largely resistant to steganalysis up to 0.4 bits per pixel in 512×512 grayscale images.

The other incentive for organizing the challenge was a hope to encourage the development of new approaches toward steganalysis, pointing to important deadlocks in steganalysis and hopefully finding solutions to them, finding weaknesses of the proposed steganographic system, and finally raising interest in steganalysis and steganography. While running the contest, we became aware of a similar contest organized within the computer vision community [7].

This paper serves as an introduction to a series of papers [4,5,6] describing the attacks on HUGO. Here, we describe the contest, image databases, and the HUGO algorithm to give the papers uniform notation and background.

1.1 Requirements and rules

In order for BOSS challenge to be attractive and fruitful for the community, we have obeyed the following conditions and limitations.

- All participants were ranked by a scalar criterion, the accuracy of detection on a database of 1,000 512×512 grayscale images called BOSSRank. Each image in the BOSSRank database was chosen to contain secret message of size 104,857 bits (0.4 bits per pixel) with probability 50% (naturally the list of stego and cover images was kept secret).
- In order to ensure that all participants start with the same degree of knowledge about the steganographic system used in the contest, we started the contest with a warm-up phase on June 28, 2010. The very same day the steganographic algorithm HUGO was presented at the International Hiding Conference 2010. For the warm-up phase, we also released the source code of the embedding algorithm. To simplify the steganalysis, a training database of 7,518 512×512 grayscale images (the BOSSBase) was released along with an implementation of the state-of-the-art feature set (the Cross Domain Features (CDF) [10]) for blind steganalysis. The motivation leading to supply this material, especially the description and implementation of the embedding algorithm, came from the Kerckhoffs' principle.

- We wanted all participants to have an easy access to the score of their predictions, yet prevent them to perform an oracle attack⁴ on the evaluation system. To achieve both requirements, the hosting server <http://www.agents.cz/boss> allowed to upload a prediction on BOSSRank once every three days for every IP address. Moreover, the provided score was computed from a subset of 900 randomly selected images. If the detection accuracy was above 65%, the participants could enter the the Hall of Fame.
- To impose a deadline for the participants, the challenge was divided into two phases. The warm-up phase started on June 28, 2010 and ended on September 9, 2010 by publishing the BOSSRank image database used to evaluate the participants. This was immediately followed by a four-month-long period, during which the challenge took its place. The challenge was originally scheduled to end on December 15, 2010, but it was later extended to January 10, 2011.

1.2 Source of cover images for BOSS

The BOSS webpage offered two databases of images, the BOSSBase and the BOSSRank.

BOSSBase was composed of 9,074 never-compressed cover images coming from 7 different cameras.⁵ This database was provided as the source of cover images used for the development of steganalyzers. All images were created from full-resolution color images in RAW format (CR2 or DNG). The images were first resized so that the smaller side was 512 pixels long, then they were cropped to 512×512 pixels, and finally converted to grayscale. The whole process was published in a script along with the original images in RAW format and their EXIF headers. Table 1 shows the actual number of images for each camera.

The BOSSRank database was composed of 1,000 512×512 grayscale images obtained by the same processing script. 482 of them were randomly chosen to carry the secret payload of approximately 0.4 bpp while keeping the rest without any payload. Participants did not know that 847 images were obtained by Leica M9 in RAW format and 153 images came from Panasonic Lumix DMC-FZ50 captured directly in JPEG⁶ format.

The fact that images in both databases came from slightly different sources lead to interesting consequences on steganalyzers trained purely on the BOSSBase. Although created unintentionally, this cover source mismatch forced the participants to deal with the situation, where the exact source of cover images

⁴ A method to reach 100% accuracy by learning the true classification of BOSSRank from a very large number of carefully constructed predictions.

⁵ The *BOSSBase* was released in three phases. On June 28, 2010, the version 0.90 containing 7518 images was released. When the challenge moved to its second phase, the version 0.92 was released with 9074 images. Finally, the version 1.0 containing 10000 images was released in May 2011.

⁶ Initially we wanted to use images only from one of the camera in *BOSSBase*, but because of the lack of time we had to use another camera that was not in the training database.

Camera model	# of images in <i>BOSSBase</i>	# of images in <i>BOSSRank</i>
Leica M9	2267	847
Canon EOS DIGITAL REBEL XSi	1607	0
PENTAX K20D	1398	0
Canon EOS 400D DIGITAL	1354	0
Canon EOS 7D	1354	0
NIKON D70	1033	0
Canon EOS 40D	61	0
Panasonic Lumix DMC-FZ50	0	153

Table 1. Camera models and number of images in *BOSSBase* v0.92 and *BOSSRank*.

is not fully known, a problem which surely happens in practice when detecting steganographic communication. Designing steganalyzers which are robust to the cover-source mismatch was one of the main challenges which the participants very quickly realized.

2 HUGO, the embedding algorithm for BOSS

The HUGO (Highly Undetectable steGO) algorithm used in the contest hides messages into least significant bits of grayscale images represented in the spatial domain. It was designed to follow the minimum-embedding-impact principle, where we embed a given message while minimizing a distortion calculated between cover and stego images. This strategy allows to decompose its design into two parts: the design of *image model* and the *coder*. The role of the image model is to generate a space in which the distance between points leads to a good distortion function. This function is subsequently used by the coder to determine the exact cover elements that need to be changed in order to communicate the message. In addition, the *optimal coder* minimizes the average distortion calculated over different messages of the same length. The relationship between the size of the payload (embedding rate) and the average distortion is often called the rate-distortion bound. Due to recent development in coding techniques [2,3], we believe that larger gains (in secure payload for example) can be achieved by designing distortion functions more adaptively to the image content instead of by changing the coder. From this reason, when designing HUGO we have focused on the image model.

The image model was largely inspired by the Subtractive Pixel Adjacency Matrix (SPAM) steganalytic features [11], but steps have been taken to avoid over-fitting to a particular feature set [9]. The original publication [12] describes and analyzes several different versions of the algorithm. Here, the most powerful version used in the BOSS competition is described.

2.1 HUGO's image model

For the purpose of embedding, each image $\mathbf{X} = (x_{i,j}) \in \mathcal{X} \triangleq \{0, \dots, 255\}^{n_1 \times n_2}$ of size $n_1 \times n_2$ pixels is represented by a feature vector computed from eight three-dimensional co-occurrence matrices obtained from differences of horizontally, vertically, and diagonally neighboring pairs of pixels. The (d_1, d_2, d_3) th entry of the empirical horizontal co-occurrence matrix calculated from $\mathbf{X}$ is defined as

$$C_{d_1, d_2, d_3}^{\mathbf{X}, \rightarrow} = \frac{1}{n_1(n_2 - 2)} \left| \{(i, j) | D_{i,j}^{\rightarrow} = d_1 \wedge D_{i,j+1}^{\rightarrow} = d_2 \wedge D_{i,j+2}^{\rightarrow} = d_3\} \right|, \quad (1)$$

where $d_1, d_2, d_3 \in [-T, -T+1, \dots, T]$, $D_{i,j}^{\rightarrow} = x_{i,j} - x_{i,j+1}$ when $|x_{i,j} - x_{i,j+1}| \leq T$. Differences greater than T , $|x_{i,j} - x_{i,j+1}| > T$, are not considered in the model. Co-occurrence matrices for other directions, $k \in \{\leftarrow, \uparrow, \downarrow, \searrow, \swarrow, \nearrow, \nwarrow\}$ are defined analogically. The feature vector defining the image model is $(\mathbf{F}^{\mathbf{X}}, \mathbf{G}^{\mathbf{X}}) \in \mathbb{R}^{2(2T+1)^3}$ with

$$F_{d_1, d_2, d_3}^{\mathbf{X}} = \sum_{k \in \{\rightarrow, \leftarrow, \uparrow, \downarrow\}} C_{d_1, d_2, d_3}^{\mathbf{X}, k}, \quad G_{d_1, d_2, d_3}^{\mathbf{X}} = \sum_{k \in \{\searrow, \swarrow, \nearrow, \nwarrow\}} C_{d_1, d_2, d_3}^{\mathbf{X}, k}. \quad (2)$$

The embedding distortion between cover $\mathbf{X}$ and stego image $\mathbf{Y}$, $D(\mathbf{X}, \mathbf{Y})$, is a weighted L_1 -norm between their feature vectors:

$$D(\mathbf{X}, \mathbf{Y}) = \sum_{d_1, d_2, d_3 = -T}^T \left[w(d_1, d_2, d_3) |F_{d_1, d_2, d_3}^{\mathbf{X}} - F_{d_1, d_2, d_3}^{\mathbf{Y}}| + w(d_1, d_2, d_3) |G_{d_1, d_2, d_3}^{\mathbf{X}} - G_{d_1, d_2, d_3}^{\mathbf{Y}}| \right], \quad (3)$$

where the weights $w(d_1, d_2, d_3)$ quantify the detectability of an embedding change in the (d_1, d_2, d_3) th element of $\mathbf{F}$ and $\mathbf{G}$. The weights were heuristically chosen as

$$w(d_1, d_2, d_3) = \left(\sqrt{d_1^2 + d_2^2 + d_3^2} + \sigma \right)^{-\gamma}, \quad (4)$$

where σ and γ are scalar parameters. For the BOSS challenge, the parameters were set to $\sigma = 1$, $\gamma = 1$, and $T = 90$.

2.2 Embedding

The practical implementation of HUGO embeds the message in pixel's LSBs by using Syndrome-Trellis Code (STC), which were shown [3] to achieve near optimal rate-distortion performance. For the purpose of the challenge, only a simulator of HUGO with the STC coder replaced by a simulated optimal coder operating at the rate-distortion bound was released. This coder modifies i th pixel x_i to $y_i = \arg \min_{z \in \{x_i-1, x_i+1\}} D(\mathbf{X}, z\mathbf{X}_{\sim i})$ with probability

$$p_i = \Pr(Y_i = y_i) = \frac{1}{Z} e^{-\lambda D(\mathbf{X}, y_i \mathbf{X}_{\sim i})}, \quad (5)$$

where Z is a normalization factor and $y_i \mathbf{X}_{\sim i}$ denotes the cover image whose i th pixel has been modified to $Y_i = y_i$ and all other pixels were kept unchanged. The constant $\lambda \geq 0$ is determined by the condition

$$m = - \sum_i p_i \log_2 p_i + (1 - p_i) \log_2 (1 - p_i), \quad (6)$$

which quantifies the desire to communicate m bit long message.

During embedding, whenever a pixel's LSB needs to be changed, the sender has a freedom to choose between a change by $+1$ or -1 (with the exception of boundaries of the dynamic range). The sender first chooses the direction that leads to a smaller distortion (3), embeds the message and then perform the embedding changes. Moreover, in strategy S2 (the most secure version of the algorithm), the embedding changes are performed sequentially and the sender recomputes the distortion at each pixel that is to be modified because some of the neighboring pixels might have already been changed. This step does not change the communicated message and enables HUGO to consider mutual interaction of embedding changes and thus further minimize the statistical detectability.

To illustrate the adaptivity of the algorithm, Figure 1 shows the average probability of changing each pixel in the Lena image⁷ estimated by embedding 500 different messages of the same length using the simulated coding algorithm.

Fig. 1. Probabilities of pixel being changed during embedding in the Lena image. Probabilities were estimated by embedding 500 different pseudo-random messages with sizes 0.25/0.5 bits per pixel (bpp).

⁷ Obtained from <http://en.wikipedia.org/wiki/File:Lenna.png>.

3 Final results and analysis of the submissions

From a large number of received submissions, only 3 participant teams have entered the Hall of Fame, namely A. Westfeld, the team of J. Fridrich called Hugobreakers and finally the team of G. Gül & F. Kurugollu. Final competition results and scores: (1) Hugobreakers 80.3%, (2) Gül & Kurugollu 76.8%, and (3) A. Westfeld 67%. As can be seen from the number of unique IP addresses from which the BOSSRank image database was downloaded, many other researchers tried to play BOSS. Figure 2 shows the distribution of 142 unique IP addresses among different countries.

Fig. 2. Number of unique IP addresses from which the BOSSRank image database was downloaded during the contest. Total 142 IP addresses were recorded.

3.1 Cover-source mismatch

The cover-source-mismatch problem refers to a scenario, where images used for training the steganalyzer do not come from the same source as images w.r.t. which the steganalyzer is tested. If the source of images is very different and the steganalyzer is not robust with respect to this discrepancy, this can lead to decrease of the detection accuracy. By accident, the addition of pictures coming from a camera which was not used in BOSSBase has caused the cover-source mismatch problem.. Figure 3 shows the accuracy of submissions entered to the hall of fame according to the camera model. It can clearly be seen that all submissions are more accurate on images coming from the Leica M9 than on images captured by the Panasonic DMC-FZ50. The cover-source mismatch can be used to partly explain this phenomenon, the other reason might be that images coming from the DMC-FZ50 are more difficult to classify because of their contents.

The loss of accuracy is higher for steganalyzers developed by Hugobreakers than by other groups. It is also interesting to observe that on the beginning of the challenge, the accuracy of the first submission of Hugobreakers was nearly random on images coming from the Panasonic camera. From this analysis, it also appears that Gül & Kurugollu’s steganalyzers were more immune to the problem of model mismatch than the classifier proposed by Hugobreakers.

Fig. 3. Scores for each cameras for the different submissions in the Hall of Fame.

To learn from this analysis more, it would be interesting to know the design of Hugobreakers’ steganalyzers which scored at 71% and 75%, because between these two submissions, the cover-source mismatch was significantly reduced. Did this improvement come from training on a more diverse set of images, or it is due to different features or machine learning algorithm? Moreover, it should be also investigated, why steganalyzers of A. Westfeld and Gül & Kurugollu were more robust. Answers to these questions are important for building more robust and thus practically usable steganalyzers.

3.2 False positives, False negatives

We now extend the analysis from the previous subsection to false positive and false negative rates defined here as probability of cover image classified as stego and stego image classified as cover, respectively. Figure 4 shows these rates on BOSSRank together with rates on each camera separately for two best submissions of Hugobreakers and Gül & Kurugollu. We have noticed that Hugobreakers’ steganalyzer suffered from very high false positive rate on images captured by the Panasonic camera. Their best submission had an almost 47% false positive rate, but only 8% false negative rate. Surprisingly, the final steganalyzer of Gül & Kurugollu did not exhibit such an imbalance between the false positive and false negative rates. Although the score used during the challenge evaluated overall accuracy of the steganalyzers, for the practical application, it is very important to limit the false positive rate. According to the results, the cover-source mismatch can make these errors even worse.

3.3 Clustering analysis

Clustering analysis provides an interesting insight, how diverse were participants’ submissions and how they evolved in time. Figure 5 shows an MDS plot of Ham-

Fig. 4. False positive and false negative rates according to the camera for the four best submissions.

ming distances between submission vectors from the Hall of fame [8]⁸. The MDS plot reveals that the initial detector of Hugobreakers (H 68%) was similar to the detector of A. Westfeld. Later, as the challenge progressed, Hugobreakers improved their detector and departed from the initial solution. Towards the end of the contest, Hugobreakers were merely tuning their detector, but no significant change has been introduced. This can be recognized by many submissions forming a tiny cluster. On the other hand, the detector developed by Gül & Kurugollu was from the very beginning different from detectors of other participants, as their submissions form a small compact cluster within the space.

It is interesting to see that Hugobreakers and Gül & Kurugollu have developed detectors with similar accuracy but independent errors. This is supported by the fact that only two images out of 1000 were always incorrectly classified (both images, image no. 174 and image no. 353, were false positives). In other words for 99.8% of the images there has been at least one submission in which the image was classified correctly. These suggest that the accuracy can be improved by fusing the classifiers developed in the contest as is shown in the next section.

⁸ Multi-Dimensional Scaling (MDS) plot tries to map points from high-dimensional space to low-dimensional space such that distances between individual points are preserved.

Fig. 5. MDS plot of submissions entered to Hall of fame. Legend: A — Andreas Westfeld, G — Gül & Kurugollu, and H — Hugobreakers. Each submission is labeled by the score as calculated on 900 random images measured at the time of submission. Final solutions are labeled by the score calculated w.r.t. the whole BOSSRank database.

4 Mixing strategies

From the informed analysis done in the previous section, we noticed that the submission $\mathbf{h} = (h_1, \dots, h_{1000}) \in \{0, 1\}^{1000}$ ⁹ of Hugobreakers scoring 79% is more immune to cover-source mismatch and false positive errors than their final submission $\mathbf{h}' = (h'_1, \dots, h'_{1000}) \in \{0, 1\}^{1000}$ scoring 80.3%. In order to decrease the false positive errors of the final solution we fuse the two submissions and define new vector $\mathbf{c} = (c_1, \dots, c_{1000}) \in \{0, 1\}^{1000}$ as:

$$c_i = \begin{cases} 1 & \text{if } h_i = 1 \text{ and } h'_i = 1 \text{ (both submissions call } i\text{th image stego)} \\ 0 & \text{otherwise.} \end{cases}$$

⁹ Element 0 (1) in the of the submission vector corresponds to cover (stego) prediction.

Fig. 6. Comparisons between the results of the collusion and the winner of the challenge.

Figure 6 compares the performances of the collusion vector c with the best submissions of BOSS. This vector c achieves 81.3%, which is 1% more than the final score of Hugobreakers. The fused vector is also less sensitive to false positive errors and cover-source mismatch. Note however that this is an a posteriori submission using results from the test set and consequently it should be evaluated on other test sets in order to consider the comparison fair.

5 Conclusion and perspectives

As can be seen from [4,5,6], BOSS challenge has stimulated research and forced the participants to deal with many challenging problems in steganalysis. The accuracy of detection of the HUGO algorithm, developed for the challenge, has increased from 65% to 81% for an embedding capacity of 0.4bpp and further improvement is to be expected. Moreover, according to the clustering analysis presented in this report, at least two different steganalyzers with similar performance have been developed which can lead to better results after the players exchange their ideas.

In possible extensions of HUGO, authors should consider avoiding the payload-limited sender regime, where the same amount of payload is embedded in every image. Instead, the stegosystem should try to embed different amount of payload based on the image content and possibly spread the payload among multiple cover objects, i.e., use batch steganography.

Besides that, BOSS challenge pointed out that cover-source mismatch is a significant problem for practical applications of steganalyzers based on a combination of steganalytic features and machine learning algorithms. We believe that the future research should focus to mitigate the cover source mismatch to-

gether with a problem of excessively high false positive rates. These findings also underline the need to develop a methodology to compare steganalyzers in a fair manner.

One of the incentives to organize BOSS was to investigate if steganalysis can exploit the knowledge of probability of pixel changes. For adaptive schemes, which represent current state-of-the-art in steganography, this probability is not uniform and can be well estimated from the stego image. Whether this fact presents any weakness has not been proved yet, but according to our knowledge, none of the successful participants of BOSS contest was able to utilize such information.

References

1. Bas, P., Furon, T.: BOWS-2. <http://bows2.gipsa-lab.inpg.fr> (July 2007)
2. Filler, T., Fridrich, J.: Gibbs construction in steganography. *Information Forensics and Security*, IEEE Transactions on 5(4), 705–720 (dec 2010)
3. Filler, T., Judas, J., Fridrich, J.: Minimizing additive distortion in steganography using syndrome-trellis codes. *IEEE Transactions on Information Forensics and Security* (2010), under review
4. Fridrich, J., Goljan, M., Kodovský, J., Holub, V.: Steganalysis of spatially-adaptive steganography. In: Filler, T., Pevný, T., Ker, A., Craver, S. (eds.) *Information Hiding*, 13th International Workshop. Lecture Notes in Computer Science, Prague, Czech Republic (May 18–20, 2011)
5. Fridrich, J., Kodovský, J., Goljan, M., Holub, V.: Breaking hugo - the process discovery. In: Filler, T., Pevný, T., Ker, A., Craver, S. (eds.) *Information Hiding*, 13th International Workshop. Lecture Notes in Computer Science, Prague, Czech Republic (May 18–20, 2011)
6. G. GÄEL, F.K.: A new methodology in steganalysis : Breaking highly undetectable steganography (hugo). In: Filler, T., Pevný, T., Ker, A., Craver, S. (eds.) *Information Hiding*, 13th International Workshop. Lecture Notes in Computer Science, Prague, Czech Republic (May 18–20, 2011)
7. Goldenstein, S., Boulton, T.: The first IEEE workitorial on vision of the unseen. <http://www.liv.ic.unicamp.br/wvu/> (2008)
8. Gower, J.: Some distance properties of latent root and vector methods used in multivariate analysis. *Biometrika* 53(3-4), 325 (1966)
9. Kodovský, J., Fridrich, J.: On completeness of feature spaces in blind steganalysis. In: Ker, A.D., Dittmann, J., Fridrich, J. (eds.) *Proceedings of the 10th ACM Multimedia & Security Workshop*. pp. 123–132. Oxford, UK (September 22–23, 2008)
10. Kodovský, J., Pevný, T., Fridrich, J.: Modern steganalysis can detect YASS. In: Memon, N.D., Delp, E.J., Wong, P.W., Dittmann, J. (eds.) *Proceedings SPIE, Electronic Imaging, Security and Forensics of Multimedia XII*. vol. 7541, pp. 02–01–02–11. San Jose, CA (January 17–21, 2010)
11. Pevný, T., Bas, P., Fridrich, J.: Steganalysis by subtractive pixel adjacency matrix. In: Dittmann, J., Craver, S., Fridrich, J. (eds.) *Proceedings of the 11th ACM Multimedia & Security Workshop*. pp. 75–84. Princeton, NJ (September 7–8, 2009)
12. Pevný, T., Filler, T., Bas, P.: Using high-dimensional image models to perform highly undetectable steganography. In: Fong, P.W.L., Böhme, R., Safavi-Naini,

- R. (eds.) Information Hiding, 12th International Workshop. pp. 161–177. Lecture Notes in Computer Science, Calgary, Canada (June 28–30, 2010)
13. Piva, A., Barni, M.: The first BOWS contest: Break our watermarking system. In: Delp, E.J., Wong, P.W. (eds.) Proceedings SPIE, Electronic Imaging, Security, Steganography, and Watermarking of Multimedia Contents IX. vol. 6505. San Jose, CA (January 29–February 1, 2007)