How to Form Bags in Batch Steganography

Eli Dworetzky
Department of ECE
Binghamton University
edworet1@binghamton.edu

Jessica Fridrich
Department of ECE
Binghamton University
fridrich@binghamton.edu

***Abstract*—In batch steganography, the sender spreads the secret payload among multiple cover images forming a bag. The question investigated in this paper is how many and what kind of images the sender should select for her bag. We show that by forming bags with a bias towards selecting images that are more difficult to steganalyze, the sender can either lower the probability of being detected or save on bandwidth by sending a smaller bag. These improvements can be quite substantial. Our study begins with theoretical reasoning within a suitably simplified model. The findings are confirmed on experiments with real images and modern steganographic and steganalysis techniques.**

***Index Terms*—Batch steganography, pooled steganalysis, source biasing, forming bags, bias gain, bandwidth savings**

I. Introduction

As introduced by Ker [1], in batch steganography the sender spreads her secret payload among multiple covers (a bag) to decrease the chances of being detected by the Warden and / or to communicate a large message that would not fit in a single cover. The Warden pools evidence from the same bag to detect the use of steganography, a process known as pooled steganalysis.

Batch steganography and pooled steganalysis (BSPS) have been studied extensively in the past [2], [3], [4], [5], [6], [7], [8], [9], [10]. A significant portion of the work focuses on how to allocate chunks of the secret payload to covers in the bag to minimize statistical detectability [5], [1], [3], [6]. Comparatively less is available for the sender on how to form the bags themselves. The authors of [11], [9], [12] report that a sender who wishes to maintain a fixed secret rate (relative payload per bag) can decrease statistical detectability by selecting a bag whose size is neither too large nor too small. The gain in security obtained this way is called the bag gain. The authors of [6], [10] describe a method for selecting the bag size adaptively (AdaBIM) for Gaussian embedding and the image merging sender [5].

This present work is concerned with the following practical question. Given a secret message of fixed length, how should Alice form her bag of cover images for batch steganography? We emphasize that the aforementioned prior art do not directly address this question as most operate under the working assumption that Alice wants to maintain a fixed communication rate (e.g., in bits per pixel). However, maintaining a fixed rate may not be the only or the most relevant constraint for Alice. She may use the batch mode to lower the chances of being detected or simply because the payload does not fit a single image. The question then becomes a) how many images should be in the bag and b) what kind of images to include in the bag. A trivial answer to a) is to use as large a bag as possible because the statistical detectability approaches zero with increased bag size due to the square root law (SRL) [13]. However, this is not practical because of limited communication bandwidth[1] and diminishing rate of the secret communication. The sender may be satisfied with making sure the detectability is below a certain threshold while staying within a specific bandwidth. Regarding task b), as shown in [14], the sender can gain security by selecting images that are more difficult to steganalyze with some level of bias, neither too large nor too small. This gain is the so-called bias gain. The result was, however, obtained for a sender who maintains a fixed rate. Our paper builds upon this prior art but considers the bag forming problem for a fixed absolute payload.

In the next section, we introduce a previously proposed modeling framework within which a closed form expression for the ROC of Warden's optimal pooler can be derived. In Section III, we numerically analyze the trade off between detectability, bag size, and the strength of source biasing for a fixed payload size. To determine whether these findings translate to real datasets, in Section IV we carry out a series of experiments on ALASKA II with the embedding algorithm HILL and machine learning based detectors. The experimental results closely mimic the trends obtained within our simplified model. The paper is concluded in Section V.

II. Modeling framework

In this section we introduce a simple modeling framework for BSPS within which it is possible to derive a closed form for the ROC of Warden's optimal pooler. Our approach is detector-centric [11], [12], [14] in that both the impact of embedding and biased sampling are modeled through the distribution of soft outputs a steganography detector. While this makes our conclusions dependent on a detector, the advantages are substantial as the problems of detection of steganography and

WIFS'2024, December 2-5, 2024, Rome, Italy. 978-1-7281-9930-6/20/$31.00

[1]E.g., a limit on email attachment size.

biased sampling can be considered jointly through a single hypothesis test. The models are also estimable in practice, which facilitates a close match between our theoretical findings and experiments.

We denote the set of cover images by $\mathcal{X}$ and assume the number of pixels N in each image $X \in \mathcal{X}$ is fixed. For a fixed steganographic method, cover images can hold an absolute payload of C bits (e.g., $C = N \log_2 3$ for grayscale images without "wet" pixels and a ternary embedding method). We assume that the steganographer (Alice) wishes to communicate a fixed payload of αC bits, $\alpha \geq 0$.[2] When Alice does not use source biasing, she independently samples a bag of n covers from $\mathcal{X}$, $\mathbf{X} = (X_1, \ldots, X_n)$, with some choice of $n \geq \alpha$ to fit the secret message. We emphasize that α is a measure of relative payload to one image and not measured in bits per pixel. For example, $\alpha = 8$ means the message length is $8C$ bits, or 8 full images worth of payload. Alice also commits to some payload allocation strategy (or spreading strategy) to embed chunks of the secret message in each image. Denoting the absolute payload to be embedded in X_i with $\alpha_i C$, $0 \leq \alpha_i \leq 1$, the payload constraint is $\sum_{i=1}^{n} \alpha_i = \alpha$. Alice produces the ith stego image $X_i(\alpha_i)$ by embedding X_i with payload of length α_i.

We will assume that the Warden has a single-image steganography detector (SID), which is a mapping $d : \mathcal{X} \to \mathbb{R}$ that assigns to each image a scalar referred to as the soft output (or response) of the detector. Given an intercepted bag of n images $\mathbf{Y} = (Y_1, \ldots, Y_n)$, the Warden infers whether Alice uses steganography by computing $d(Y_i)$ for all $i = 1, \ldots, n$ and comparing $\pi(d(Y_1), \ldots, d(Y_n))$ against a threshold. Here, $\pi : \mathbb{R}^n \to \mathbb{R}$ is Warden's *pooler*.

A. Modeling detector outputs

For modeling purposes, we assume that sampling from the cover source $\mathcal{X}$ is a two stage process. Alice first selects a scene and then acquires it with a digital camera. For a fixed scene, the output of the detector on all possible acquisitions of the same scene with the same camera and settings is assumed to follow a Gaussian distribution. After embedding a secret message, the mean of this distribution is assumed to change linearly with payload:

$$d(X_i(\alpha)) \sim \mathcal{N}(\mu_i + b_i \alpha, \sigma_i^2), \quad 0 \leq \alpha \leq 1, \tag{1}$$

where $b_i \geq 0$ is the *slope* of the detector's response. The Gaussianity of this conditional distribution is heuristically justified by the independent heteroscedastic acquisition noise model and the fact that d can be linearized on a small neighborhood of the noise-free scene. Even though the soft output of typical detectors built with machine learning is not generally linear w.r.t. payload, it is approximately true for small payloads (see, e.g., Figure 3 in [9]).

[2]Note that α is not necessarily bounded as we allow the message to be arbitrarily large.

B. Source model and source biasing

From existing prior art on source modeling and biasing for steganography [15], [16], [17], we build upon the work of Dworetzky et al. [14] because this study is already framed within the framework of BSPS. Within the adopted model (1), we assume that the slopes b_i follow a two-valued distribution controlled by parameter $p \in (0, 1)$:

$$b_i \sim \mathcal{B}(p) = \begin{cases} \varepsilon & \text{with probability } p \\ 1 & \text{with probability } 1 - p. \end{cases} \tag{2}$$

In other words, our cover source contains only two types of images – those where steganography is easily detectable (slope $b_i = 1$) and difficult-to-steganalyze images with $b_i = \varepsilon \ll 1$. Given the biasing parameter $q \in (0, 1)$, source biasing simply involves sampling images so that the distribution of slopes follows the biased distribution $b_i \sim \mathcal{B}(q)$. Clearly, Alice should choose $q \geq p$ so that she is more likely to embed difficult-to-steganalyze images.

Within this bivalued model of the cover source, we assume Alice uses a bivalued spreading strategy as introduced in [14]. First, we let $K = |\{i : b_i = \varepsilon\}|$ be the number of images whose slopes are ε, so K is a binomial random variable. Alice assigns $\alpha_i = \alpha_\varepsilon$ to all K images with $b_i = \varepsilon$ and $\alpha_i = \alpha_1$ to all $(n - K)$ images with $b_i = 1$. The pair $\{\alpha_\varepsilon, \alpha_1\}$ is determined to satisfy the absolute payload constraint $\alpha = K\alpha_\varepsilon + (n - K)\alpha_1$.

C. Optimal pooler

In order to obtain a simple hypothesis test for Warden's detection problem, we will assume that the parameters p and ε are known both to Alice and the Warden and the Warden knows her spreading strategy, biasing parameter q, and the payload α. Moreover, we assume that the embedding does not change the slopes b_i.

Given a bag of images $\mathbf{Y} = (Y_1, \ldots, Y_n)$, we further simplify by assuming that $\sigma_i^2 = 1$ for all i and that μ_i is known to the Warden, which makes Warden's hypothesis test simple

$$\begin{aligned} \mathcal{H}_0 &: \quad b_i \sim \mathcal{B}(p), \quad y_i \sim \mathcal{N}(0, 1) && \text{for all } i \\ \mathcal{H}_1 &: \quad b_i \sim \mathcal{B}(q), \quad y_i \sim \mathcal{N}(b_i \alpha_i(K), 1) && \text{for all } i \end{aligned} \tag{3}$$

where $y_i = d(Y_i)$ and K is the number of images with slope ε in the bag. Observe that Alice's assignment of α_i depends on the realization of K, but regardless the y_i are conditionally Gaussian given the b_i. Hence, the most powerful detector for this problem is the likelihood ratio test (LRT)

$$\begin{aligned} L(\mathbf{b}, \mathbf{y}) = &\sum_{i=1}^{n} y_i b_i \alpha_i(K) - \frac{1}{2} \sum_{i=1}^{n} b_i^2 \alpha_i^2(K) \\ &+ K \log \frac{q}{p} + (n - K) \log \frac{1 - q}{1 - p}. \end{aligned} \tag{4}$$

Given a decision threshold $x \in \mathbb{R}$, the ROC of Warden's optimal pooler has the following parametric form owing to the law of total probability

$$P_{\rm FA}(x) = \sum_{k=0}^{n} \binom{n}{k} p^k (1-p)^{n-k} Q\left(\frac{x - \ell(k) + \Delta^2(k)}{\sqrt{2\Delta^2(k)}}\right) \tag{5}$$

$$P_{\rm D}(x) = \sum_{k=0}^{n} \binom{n}{k} q^k (1-q)^{n-k} Q\left(\frac{x - \ell(k) - \Delta^2(k)}{\sqrt{2\Delta^2(k)}}\right), \tag{6}$$

where $P_{\rm FA}$ is the false alarm rate, $P_{\rm D}$ is the true positive rate, Q is the $\mathcal{N}(0,1)$ tail function,

$$\Delta^2(k) = \frac{1}{2}\sum_{i=1}^{n} b_i^2 \alpha_i^2(k) \tag{7}$$

is the steganographic deflection coefficient of a bag conditioned on the event $K = k$, and

$$\ell(k) = k \log \frac{q}{p} + (n-k) \log \frac{1-q}{1-p}. \tag{8}$$

The closed form ROC (5)–(6) allows us to compute $P_{\rm D}$ at a desired $P_{\rm FA}$ via a root-finding algorithm and thus study the effect of the bag size n and the biasing parameter q on security numerically in the next section.

III. Model analysis

As reported in our previous work [14], biasing has a complicated effect on Warden's ROC. For fixed bag size n, payload size α, and small $P_{\rm FA}$, the true positive rate $P_{\rm D}$ decreases as q increases until reaching a local minimum where the trend reverses. The value of q at which the minimum occurs depends on $P_{\rm FA}$. For $P_{\rm FA}$ sufficiently large, biasing only increases $P_{\rm D}$. Thus, in the rest of the paper we measure security in terms of $P_{\rm D}$ at a fixed $P_{\rm FA}$ rather than a global measure of the ROC such wAUC.

In this paper, we will assume that Alice has a maximal detectability requirement $\widetilde{P}_{\rm D}$ and hence she wants to choose bag size n and bias q so that $P_{\rm D} \leq \widetilde{P}_{\rm D}$. As will be seen below, she might additionally be interested to meet this requirement with the smallest bag size n to avoid raising suspicion from high bandwidth consumption or to comply with a hard restriction $n \leq n_{\max}$ if her overt communication channel is bandwidth limited. This reasoning motivated us to introduce the following two concepts – the bias gain and bandwidth savings.

A. Bias gain

For fixed payload α and bag size n, let us denote the true positive rate at fixed $P_{\rm FA}$ when biasing with q as $P_{\rm D}(q, P_{\rm FA})$. Alice's bias gain is the difference

$$\gamma_{\rm bias}(p, P_{\rm FA}) = P_{\rm D}(p, P_{\rm FA}) - P_{\rm D}(q_{\rm opt}, P_{\rm FA}), \tag{9}$$

where

$$q_{\rm opt}(p, P_{\rm FA}) \coloneqq \operatorname{argmin}_q P_{\rm D}(q, P_{\rm FA}), \tag{10}$$

is the optimal biasing parameter. In the rest of this section, whenever the value of $P_{\rm FA}$ is clear from the context and since p will be fixed, for brevity we will simply write $\gamma_{\rm bias}$,$q_{\rm opt}$, and $P_{\rm D}(q)$ instead of (9)–(10).

B. Bandwidth savings

Saving on bandwidth is the other benefit Alice can enjoy due to biasing. Formally, given Warden's pooler with her decision threshold set to achieve false alarm rate $P_{\rm FA}$, $q \geq p$, and a target $\widetilde{P}_{\rm D} \geq P_{\rm FA}$, let $n(q, \widetilde{P}_{\rm D}, P_{\rm FA})$ be the smallest bag size for which Warden's true positive rate satisfies $P_{\rm D} \leq \widetilde{P}_{\rm D}$.[3] The bandwidth savings is the difference between the bag sizes that meet the target maximum detectability without and with biasing:

$$w(p, P_{\rm D}, P_{\rm FA}) = n(p, \widetilde{P}_{\rm D}, P_{\rm FA}) - n(q_{\rm opt}, \widetilde{P}_{\rm D}, P_{\rm FA}). \tag{11}$$

The relative bandwidth savings is the ratio $w(p, \widetilde{P}_{\rm D}, P_{\rm FA})/n(p, \widetilde{P}_{\rm D}, P_{\rm FA})$.

C. Numerical study

To simplify matters in our numerical study we opted to fix the cover source parameters to $p = 0.3$ and $\varepsilon = 0.01$ which best reflect the dataset "Binarized Alaska II" in [14].[4] The payload allocation strategy we study is the bivalued Greedy sender [14] – it first spreads the payload uniformly across all images in the bag with $b_i = \varepsilon$ and, if necessary, spreads the rest of the payload uniformly across images with $b_i = 1$.

First, in Figures 1, 2, and 3 we present a few plots that juxtapose Alice's security without biasing, $P_{\rm D}(p)$, and with optimal biasing, $P_{\rm D}(q_{\rm opt})$ as functions of bag size n. Each plot fixes Alice's payload size α and the Warden's $P_{\rm FA}$. The black curves and multicolored curves correspond to $P_{\rm D}(p)$ and $P_{\rm D}(q_{\rm opt})$, respectively. The marker color depicts the value of the optimal biasing parameter, $q_{\rm opt}$. Some of the plots are shown on a log-linear scale since Alice is interested in achieving small $P_{\rm D}$. Observe that as $n \to \infty$, $P_{\rm D}$ tends to $P_{\rm FA}$ (not 0) since the Warden will be randomly guessing asymptotically.

Figure 1 depicts an example of Alice choosing the smallest bag size that achieves her desired maximum detectability $P_{\rm D} \leq \widetilde{P}_{\rm D} = 10^{-2}$. If she does not bias, she should select $n = 46$ but if she does bias she should select $n = 34$, saving 12 images on the size of her bag—her bandwidth savings (11) is 12. Figure 3 shows an example of Alice choosing a bag size to fit within a bandwidth limit $n_{\max} = 10$. Regardless of whether she biases, she should select the largest bag size $n = 10$ to minimize $P_{\rm D}$. However, by biasing she decreases the probability she is caught by 0.1.

Figure 4 shows the relative bandwidth savings as a function of $\widetilde{P}_{\rm D}$ for $P_{\rm FA} = 10^{-3}$ for four payloads. Note that the savings can be significant, up to 35%, depending on the target $\widetilde{P}_{\rm D}$ (c.f. Figure 2 right).

[3] Such bag size must exist since $\lim_{n\to\infty} P_{\rm D}(p, P_{\rm FA}) = P_{\rm FA}$ due to the SRL.

[4] The relative number of hard-to-steganalyze images was about 0.3 and the ratio between the average b_i of hard images and easy images was 0.01.

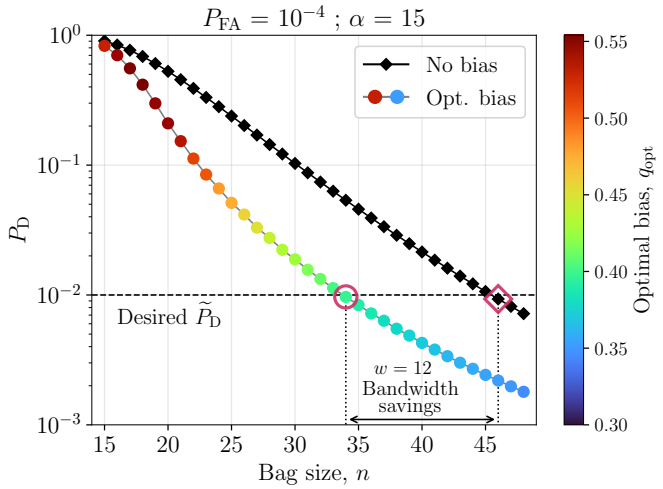

Figure 1. The Warden's $P_{\rm D}$ (bivalued model) vs. bag size n under the model. The bandwidth savings (11) for target $\widetilde{P}_{\rm D} = 10^{-2}$ is shown graphically. Observe $P_{\rm D}$ is shown on a log scale while n is on a linear scale.

IV. Experiments on a real dataset

All experiments were executed on ALASKA II developed as in [18] without the final JPEG compression step. This dataset contains 75,000 images, which we randomly divide into four disjoint partitions for our experiments. One partition of 23k images was used to train Alice's SID (to spread her message), and another partition of 23k images was used by the Warden to train her SID for detection. Once the Warden's SID is trained, another partition of 10k images was used to form bags for training the Warden's pooler.[5] The remaining 19k images were used to form bags for evaluation. Since we are interested in small $P_{\rm FA}$, we form 10k cover and 10k stego bags for each setting.

Given SID d and cover image X, the slope of X with capacity C bpp was estimated from 100 embeddings with random stego keys $b = (\mathbb{E}[d(X(C))] - d(X))/C$ with the expectation taken over the embeddings.

A. Greedy batch sender

Since the slopes in real datasets follow a continuous distribution, we study the following modification of the greedy sender. Given a set of n images, the sender uses her own SID to estimate the slopes, orders them from the smallest slope to the largest, and then using the embedding algorithm HILL [19], she embeds the images one by one at their capacity until the required payload is embedded. The last image may be embedded only partially.

B. Single-image detectors and pooler

Both Alice and the Warden train an SRNet [20] on their splits. Both were pre-trained on ImageNet as described in [21] and refined to detect HILL on stego images embedded with relative payloads randomly drawn from $\mathcal{P}$ ={0.05, 0.1, 0.2, ..., 1.5}. Each 23k image partition was further randomly split into disjoint subsets of 22k and 1k images for training and validation, respectively. The CNNs logit was used as the detector's soft response d.

For fixed bag size n, the Warden's pooling function was implemented as a random forest (RF) (Python's package `scikit-learn`) on a $2n + 1$-dimensional feature vector extracted from all images in the bag $(d(X_1), \ldots, d(X_n), b_1, \ldots b_n, \pi_{\rm CORR})$, where b_i is Warden's detector response slope of the ith image and $\pi_{\rm CORR} = \sum_{i=1}^{n} \alpha_i d(X_i)$ is the correlation of Warden's SID soft outputs with payloads that might reside in the images.

C. Biasing

Since the distribution of slopes in real datasets is continuous, we use a version of compounding to bias the sampling of cover images as described in detail in [14]. Denoting the CDF of slopes b with F, biasing the sampling of F was implemented by a modified inverse transform sampling algorithm with the beta distribution $\mathrm{Beta}(1/q, 1)$, where $q > 0$ is the biasing parameter. Selecting cover images (their slopes) with bias q is achieved by sampling $F^{-1}(G_q^{-1}(U))$, where G_q is the CDF of $\mathrm{Beta}(1/q, 1)$ and U is uniformly distributed on $[0, 1]$. Note that since $\mathrm{Beta}(1, 1) = U$, selecting $q = 1$ corresponds to unbiased sampling. The optimal value of q was found experimentally by a grid search $1/q \in \{0.05, 0.1, 0.15, \ldots, 0.95, 1.0\}$.

D. Results

In Figure 5, we report Warden's $P_{\rm D}$ at $P_{\rm FA} = 10^{-2}$ (top) and $P_{\rm FA} = 10^{-3}$ (bottom) versus bag size n for four payloads when Alice biases with $q_{\rm opt}$. Overall, the graphs show similar trends to those observed for the simplified model in the previous section. Alice's optimal bias changes with bag size and the effect of biasing diminishes towards larger bag sizes. Alice can also enjoy both bias gain and bandwidth savings. The relative bandwidth savings for the smaller payloads $\alpha \in \{2, 4\}$ at fixed $P_{\rm FA} = 10^{-2}$ ranges between 20-30% if the desired $\widetilde{P}_{\rm D}$ is approximately between 0.03 and 0.4. For the larger payloads $\alpha \in \{8, 16\}$ at fixed $P_{\rm FA} = 10^{-3}$, the relative bandwidth savings can reach between 30-50% for desired $\widetilde{P}_{\rm D} > 0.1$.

In summary, the experiments on a real dataset demonstrate a close qualitative match between the trends observed within our model in Section III. Selecting bags with a bias can benefit Alice in two ways. She can either lower her chances of being caught or save on the number of images she needs to communicate. She can also opt for a trade off between both of these benefits.

V. Conclusions

This paper studies the problem of forming bags in batch steganography – how many and what kind of images Alice should use. When Alice is allowed to sample the images for her bag in a biased fashion to prefer sending harder-to-steganalyze images, she can enjoy a gain in terms of a lower detection probability $P_{\rm D}$ or she can choose to send a smaller bag size at the same $P_{\rm D}$. She can also trade off between both. We provide theoretical insight by working within a suitably

[5] For each setting, Warden independently forms 5k cover and 5k stego training bags.

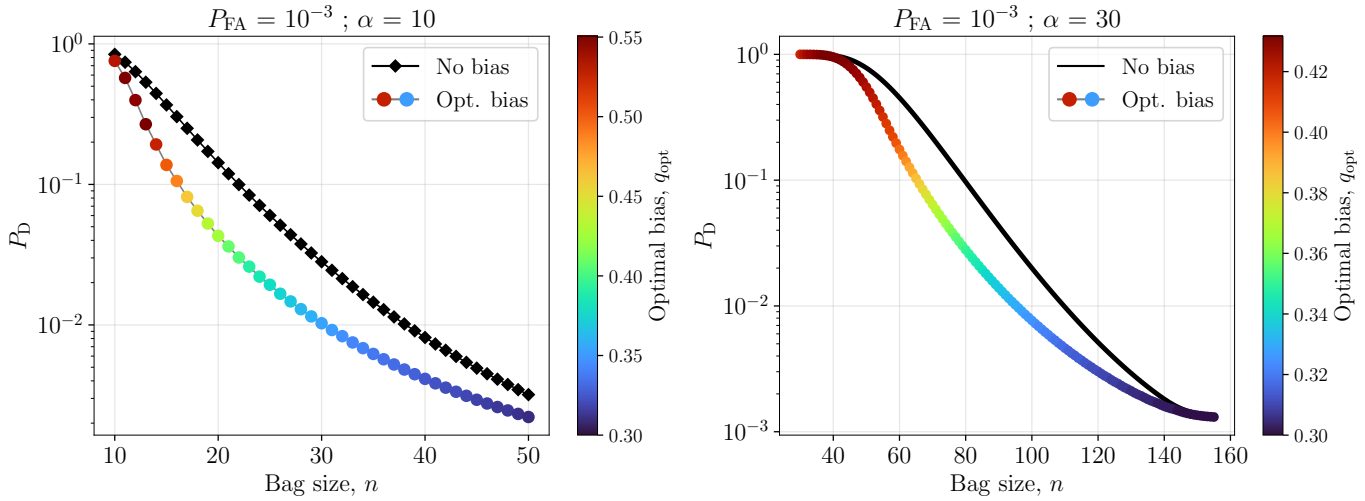

Figure 2. The Warden's $P_{\rm D}$ (bivalued model) vs. bag size n for two payloads $\alpha = 10$ and 30 with fixed $P_{\rm FA} = 10^{-3}$. Both plots show that biasing can decrease $P_{\rm D}$ by half an order of magnitude. Observe $P_{\rm D}$ tends to 10^{-3} as $n \to \infty$ since $P_{\rm FA} = 10^{-3}$. Both plots are on a log-linear scale.

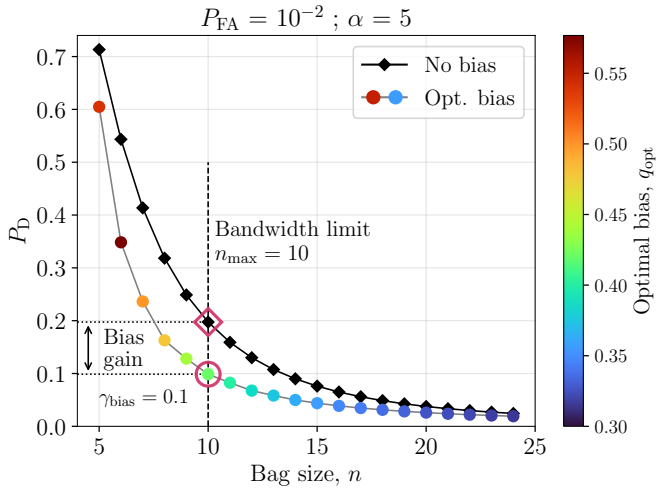

Figure 3. The Warden's $P_{\rm D}$ (bivalued model) vs. bag size n for $\alpha = 5$ and $P_{\rm FA} = 10^{-2}$. The bias gain (9) for bandwidth limit $n = 10$ is shown graphically. Note that both axes are on a linear scale since the dynamic range of $P_{\rm D}$ is only two orders of magnitude.

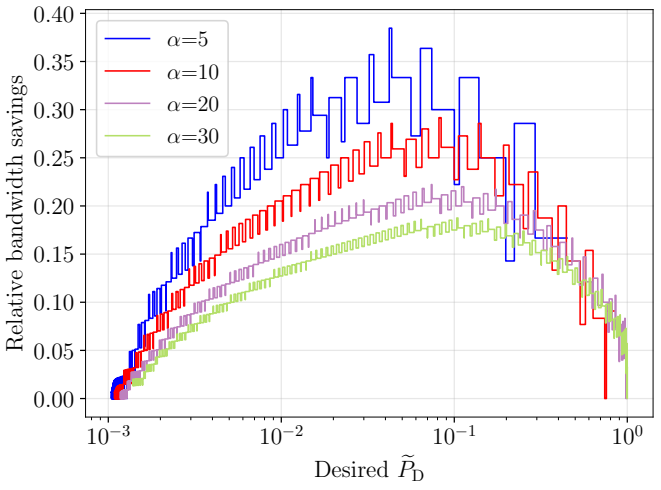

Figure 4. Relative bandwidth savings vs. Alice's desired maximum detectability $\widetilde{P}_{\rm D}$ for four payloads, $P_{\rm FA} = 10^{-3}$. The oscillations are due to the fact that the bag size is an integer.

simplified modeling framework. Our findings are confirmed in a practical setup on the ALASKA II dataset when Alice embeds with HILL and Warden steganalyzes with a machine learning based pooler. The bias gain and bandwidth savings stay relatively consistent across different Warden's detection thresholds (fixed false alarms) and payloads.

Acknowledgment

The work on this paper was supported by NSF grant No. 2028119.

References

[1] A. D. Ker, "Batch steganography and pooled steganalysis," in *Information Hiding, 8th International Workshop* (J. L. Camenisch, C. S. Collberg, N. F. Johnson, and P. Sallee, eds.), vol. 4437 of Lecture Notes in Computer Science, (Alexandria, VA), pp. 265–281, Springer-Verlag, New York, July 10–12, 2006.

[2] A. D. Ker, "Batch steganography and the threshold game," in *Proceedings SPIE, Electronic Imaging, Security, Steganography, and Watermarking of Multimedia Contents IX* (E. J. Delp and P. W. Wong, eds.), vol. 6505, (San Jose, CA), pp. 04 1–13, January 29–February 1, 2007.

[3] A. D. Ker and T. Pevný, "Batch steganography in the real world," in *Proceedings of the 14th ACM Multimedia & Security Workshop* (J. Dittmann, S. Craver, and S. Katzenbeisser, eds.), (Coventry, UK), pp. 1–10, September 6–7, 2012.

[4] L. Li, W. Zhang, C. Qin, K. Chen, W. Zhou, and N. Yu, "Adversarial batch image steganography against CNN-based pooled steganalysis," *Signal Processing*, vol. 181, pp. 107920–107920, 2021.

[5] V. Sedighi, R. Cogranne, and J. Fridrich, "Practical strategies for content-adaptive batch steganography and pooled steganalysis," in *Proceedings IEEE, International Conference on Acoustics, Speech, and Signal Processing*, March 5–9, 2017.

[6] M. Sharifzadeh, M. Aloraini, and D. Schonfeld, "Adaptive batch size image merging steganography and quantized Gaussian image steganography," *IEEE Transactions on Information Forensics and Security*, vol. 15, pp. 867–879, 2020.

[7] T. Pevný and I. Nikolaev, "Optimizing pooling function for pooled steganalysis," in *IEEE International Workshop on Information Forensics*

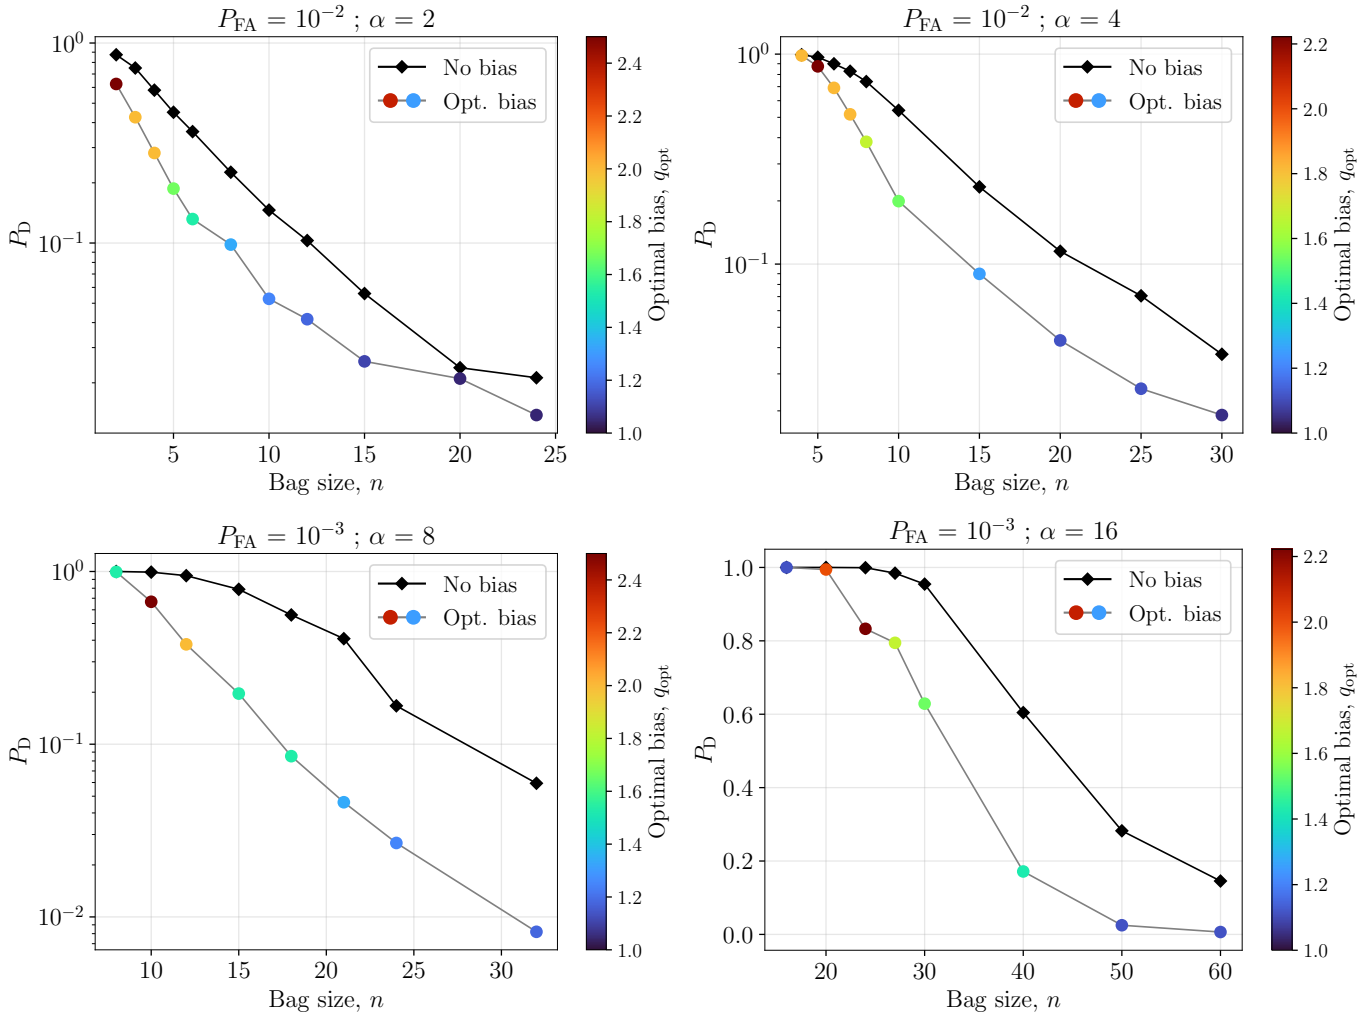

Figure 5. The Warden's $P_{\rm D}$ vs. bag size n on ALASKA II images. Four payloads $\alpha \in \{2, 4, 8, 16\}$ using Greedy sender & HILL for embedding. $P_{\rm FA} = 10^{-2}$ (top) and $P_{\rm FA} = 10^{-3}$ (bottom). Since $P_{\rm D}(p)$ (no bias) for $\alpha = 16$ has a smaller dynamic range, we show the bottom right plot on a linear scale. The other plots are on a log-linear scale. When $P_{\rm D}(q_{\rm opt}) \approx 1$, the monotonic trend in $q_{\rm opt}$ breaks due to statistical noise.

and Security, (Rome, Italy), pp. 1–6, November 16–19, 2015.

[8] A. Zakaria, M. Chaumont, and G. Subsol, "Pooled steganalysis in JPEG: how to deal with the spreading strategy?," in *IEEE International Workshop on Information Forensics and Security (WIFS)*, (Delft, The Netherlands), pp. 1–6, December 9–12 2019.

[9] Y. Yousfi, E. Dworetzky, and J. Fridrich, "Detector-informed batch steganography and pooled steganalysis," in *The 10th ACM Workshop on Information Hiding and Multimedia Security* (J. Butora, C. Veilhauer, and B. Tondi, eds.), (Santa Barbara, CA), ACM Press, 2022.

[10] M. Aloraini, M. Sharifzadeh, and D. Schonfeld, "Quantized gaussian JPEG steganography and pool steganalysis," *IEEE Access*, vol. 10, pp. 38031–38044, 2022.

[11] E. Dworetzky and J. Fridrich, "Explaining the bag gain in batch steganography," *IEEE Transactions on Information Forensics and Security*, vol. 18, pp. 3031–3043, 2023.

[12] E. Dworetzky, E. Kaziakhmedov, and J. Fridrich, "Observing bag gain in JPEG batch steganography," in *IEEE International Workshop on Information Forensics and Security (WIFS)*, (Nürnberg, Germany), December 4–7 2023.

[13] A. D. Ker, "The square root law of steganography," in *The 5th ACM Workshop on Information Hiding and Multimedia Security* (M. Stamm, M. Kirchner, and S. Voloshynovskiy, eds.), (Philadelphia, PA), ACM Press, June 20–22, 2017.

[14] E. Dworetzky, E. Kaziakhmedov, and J. Fridrich, "Improving steganographic security with source biasing," in *The 12th ACM Workshop on Information Hiding and Multimedia Security* (F. P. Gonzales, P. Comesana, C. Krätzer, and H. V. Zhao, eds.), (Baiona, Spain), ACM Press, 2024.

[15] Z. Wang and X. Zhang, "Secure cover selection for steganography," *IEEE Access*, vol. 7, pp. 57857–57867, May 2019.

[16] Z. Wang, X. Zhang, and Z. Yin, "Joint cover-selection and payload-allocation by steganographic distortion optimization," *IEEE Signal Processing Letters*, vol. 25, pp. 1530–1534, October 2018.

[17] Q. Giboulot, T. PevnÃœ, and A. D. Ker, "The non-zero-sum game of steganography in heterogeneous environments," *IEEE Transactions on Information Forensics and Security*, vol. 18, pp. 4436–4448, 2023.

[18] R. Cogranne, Q. Giboulot, and P. Bas, "ALASKA–2: Challenging academic research on steganalysis with realistic images," in *IEEE International Workshop on Information Forensics and Security*, (New York, NY), December 6–11, 2020.

[19] B. Li, M. Wang, and J. Huang, "A new cost function for spatial image steganography," in *Proceedings IEEE, International Conference on Image Processing, ICIP*, (Paris, France), October 27–30, 2014.

[20] M. Boroumand, M. Chen, and J. Fridrich, "Deep residual network for steganalysis of digital images," *IEEE Transactions on Information Forensics and Security*, vol. 14, pp. 1181–1193, May 2019.

[21] J. Butora, Y. Yousfi, and J. Fridrich, "How to pretrain for steganalysis," in *The 9th ACM Workshop on Information Hiding and Multimedia Security* (D. Borghys and P. Bas, eds.), (Brussels, Belgium), ACM Press, June 22–25, 2021.