

# Don't Fear the Bit Flips: Optimized Coding Strategies for Binary Classification

Frederic Sala, Shahroze Kabir, Guy Van den Broeck, and Lara Dolecek  
University of California, Los Angeles

{fredsala, shkabir}@ucla.edu, guyvdb@cs.ucla.edu, dolecek@ee.ucla.edu

## Abstract

After being trained, classifiers must often operate on data that has been corrupted by noise. In this paper, we consider the impact of such noise on the features of binary classifiers. Inspired by tools for classifier robustness, we introduce the same classification probability (SCP) to measure the resulting distortion on the classifier outputs. We introduce a low-complexity estimate of the SCP based on quantization and polynomial multiplication. We also study channel coding techniques based on replication error-correcting codes. In contrast to the traditional channel coding approach, where error-correction is meant to preserve the data and is agnostic to the application, our schemes specifically aim to maximize the SCP (equivalently minimizing the distortion of the classifier output) for the same redundancy overhead.

## I. INTRODUCTION

Noise is an enduring component of all computing and communication systems. Information is corrupted when transmitted over noisy channels [1], stored in unreliable memories [2], or processed by noisy or low-quality circuits [3]. Moreover, techniques aimed at saving power or increasing efficiency can increase noise. For example, voltage scaling further increases the probability of data corruption [4]. The traditional approach to handling noisy storage media is to implement strategies to detect and correct errors. For example, many memories and drives implement error-correcting codes [5].

The typical policy of such systems is to ensure that information read back is the same as it was when written. Such a policy represents a very strong constraint; implementations tend to be expensive, with significant storage overhead dedicated to redundant data. In this paper, we ask whether we may relax this constraint in the context of feature data for classifiers. Our goal is to reduce the effect of noise on the output of the algorithm; in other words, we consider the distortion of the algorithm outputs, rather than the inputs.

The basic setup is shown in Figure 1. We have a naïve Bayes classifier with class variable  $C$  and binary features  $X_1, X_2, \dots, X_n$ . Learning is performed on noiseless training data. Afterwards, the classifier operates on noisy data; rather than the true features  $X_1, \dots, X_n$ , only the corrupted features  $X'_1, X'_2, \dots, X'_n$  are used. Feature  $X_i$ 's value is flipped from 0 to 1 or from 1 to 0 with probability  $\epsilon_i$ . A fundamental question is to determine the *same classification probability* (SCP), that is, the probability that the classifier output is the same for  $X_1, \dots, X_n$  and for  $X'_1, \dots, X'_n$ .

Features can be protected by using an error-correction strategy. A set of redundant bits is computed when the data is written; before the data is used, both the data and redundant bits are decoded, with the goal of correcting some of the errors. This procedure is known as *channel coding*. The traditional approach is to protect all bits equally, uniformly reducing the  $\epsilon_i$ s for all  $n$  features. As we will see, this is not an efficient approach. Some features are more valuable and more worthy of protection than others. In fact, given a budget of redundancy, our major goal is an optimal resulting set of  $\epsilon_1, \dots, \epsilon_n$  to minimize the classifier distortion.




Fig. 1. Noisy classifier process model. The noiseless classification uses the features  $X_1, X_2, \dots, X_n$ . Noisy feature  $X'_i$  is produced from  $X_i$  with the binary symmetric channel  $BSC(\epsilon_i)$ :  $X'_i$  is equal to  $X_i$  with probability  $1 - \epsilon_i$  and equal to  $\bar{X}_i$  with probability  $\epsilon_i$ . Classification is performed on the noisy features  $X'_1, \dots, X'_n$ , without observing the true values  $X_1, \dots, X_n$ .

TABLE I  
SCPS FOR VARIOUS REDUNDANCY ALLOCATIONS

| Parameters |            |            |            | SCP by protection strategy |       |              |           |
|------------|------------|------------|------------|----------------------------|-------|--------------|-----------|
| $p_0(x_1)$ | $p_0(x_2)$ | $p_1(x_1)$ | $p_1(x_2)$ | Uniform                    | None  | All $x_1$    | All $x_2$ |
| 0.9        | 0.89       | 0.1        | 0.11       | 0.963                      | 0.900 | <b>0.991</b> | 0.900     |
| 0.8        | 0.81       | 0.3        | 0.3        | <b>0.964</b>               | 0.905 | 0.946        | 0.946     |

We illustrate this idea with a simple example. Take a classifier with  $n = 2$  features  $X_1, X_2$ , noisy versions  $X'_1, X'_2$ , and uniform prior. The noise parameter for both binary symmetric channels (BSCs) is  $\epsilon = 0.1$  by default. Let  $p_i(x_j)$  denote the conditional probability  $p(x_j = 0|C = i)$ . We show how the SCP varies for different values of the  $p_i(x_j)$ 's and noise parameters. We also demonstrate how protecting features uniformly or non-uniformly affects the SCP. ‘‘Protection’’ refers to an error-correction scheme that effectively reduces the  $\epsilon$  noise parameter for a particular feature. Here, we allocate 4 additional redundancy bits for protection; uniform protection gives 2 such bits to each feature, while protection on  $X_i$  alone grants all 4 bits to  $X_i$ . (More precise definitions for this terminology are provided later on in the paper; the result of protection is a reduction of the  $\epsilon_i$  noise parameter for feature  $i$ .)

In the first row of Table I,  $X_1$  contains more information about  $C$  compared to  $X_2$ . Observe that allocating all bits to  $X_1$  yields a better SCP than equal protection or protection on  $X_2$ . Conversely, if we look at the last row we see that even though  $X_2$  contains more information about the class variable, the SCP is maximized by a uniform allocation of redundancy bits. Thus we cannot simply examine the conditional probabilities to decide how to allocate redundancy. Such results motivate us to seek an informed redundancy allocation strategy.

Our contributions are

- A framework for measuring the impact of noise by the *same-classification probability* (SCP),
- A low-complexity approximation for the SCP based on quantization and polynomial multiplication,
- A study of coding strategies based on allocating redundant bits in a way that minimizes the SCP (and thus the distortion on the classifier output). We show that, surprisingly, the empirical SDP is non-monotonic with respect to error protection. Afterwards, we give an optimization for optimal redundancy allocation based on a greedy approach.

## II. PRIOR WORK

Channel coding for data protection is a vast field. Error-correction has been proposed for disk drives [6], write-once memories [7], and for RAID architectures [8]. More recently, techniques have been developed to take advantage of the specific properties of modern non-volatile memories such as flash [9], [10]. Distributed storage systems can be protected through replication or erasure coding [11], [12]. Solid state drives can also benefit from sophisticated error-correction [13]. A common aspect of such research, in contrast to our work, is that the goal is to preserve the data being stored without considering the application. That is, error-correction is found in a different abstraction layer.

The research in [14] introduces an informed channel coding scheme for linear regression. Here, the error model was also bit flips applied to a finite-width binary representations of integers. The goal was to minimize the distortion on the algorithm outputs. An approximation scheme was introduced by considering the contribution to the distortion by a flip in a single bit independently of the others.

There are a number of papers that have examined issues related to learning algorithms dealing with uncertainty. Recent work on information dropout, a technique introduced in [15] to prevent neural networks from overfitting, has considered adding noise to the activations of deep neural networks [16]; unlike our work, here noise is used during learning.

The robustness of various algorithms under noise exposure has been considered in a series of papers. [17] performed an experimental study of algorithms when data sets were corrupted by synthetic noise. Deleted features were tackled in [18] and [19]. [18], which is perhaps the most closely related work to the present paper, also considers corrupted features, and introduces two techniques to tackle noise, one based on linear programming, and the other using an online batching scheme. However, unlike our work, this paper considers only adversarial (not probabilistic) noise and does not propose optimized channel coding. [19] uses a game-theoretic approach to avoid over-reliance on a feature that could be deleted. Another work tackling the case of missing data in Bayes classifiers is [20].

Noise and coding are often important issues in distributed learning. For example, [21] introduced optimality guarantees for distributed estimators in a setting where nodes can be isolated, but limited communication between them is allowed. [22] considered problems including hypothesis testing and parameter estimation in the case of multiterminal algorithms where each terminal has a data compression constraint.

A related area is known as value of information [23], [24]. The idea is to maximize an expected reward on some observed features. One example is to observe those features that maximize the information gain [25].

Decision robustness in the case of hidden variables can be measured by the *same decision probability* (SDP) [26], [27]. SDP has been applied to the evaluation of adaptive testing [28] and of mammography-based diagnosis [29]. The same classification probability we employ in this paper bears some similarity to the SDP. The problem of handling label noise is considered in

[30]. Label noise in crowdsourcing is tackled through coding-theoretic means in [31]. Additionally, our work can be viewed as a form of feature selection [32], although we do not remove features, but rather allow them to be noisier than certain more critical features.

### III. NOISE IN BINARY CLASSIFICATION

We use the standard notation of upper-case letters for random variables and lower-case letters for their instantiations. Boldface denotes vectors of variables.

Throughout this paper we consider binary classification problems with binary features. For some test point  $\mathbf{x} = (x_1, x_2, \dots, x_n)$  with  $n$  binary features, we denote the conditional probabilities by  $\alpha_i = p(x_i|C=0)$  and  $\beta_i = p(x_i|C=1)$  for  $1 \leq i \leq n$ . Throughout this paper, we employ a naïve Bayes classifier such that  $\mathbf{x}$  is classified to

$$c(\mathbf{x}) = \arg \max \left\{ p(C=0) \prod_{i=1}^n \alpha_i, p(C=1) \prod_{i=1}^n \beta_i \right\}$$

Taking logarithms,  $\mathbf{x}$  is classified according to the value of

$$c(\mathbf{x}) = \text{sgn} \left( \log \frac{p(C=0)}{p(C=1)} + \sum_{i=1}^n \log(\alpha_i/\beta_i) \right). \quad (1)$$

#### A. Impact of Noise

In this work, we employ a simple noise model: we define an *noise parameter vector*  $\boldsymbol{\epsilon} = (\epsilon_1, \epsilon_2, \dots, \epsilon_n)$  with  $0 \leq \epsilon_i < 1/2$  for  $1 \leq i \leq n$ . The binary feature  $X_i$  is flipped to its opposite value  $\bar{X}_i$  with probability  $\epsilon_i$  and stays unchanged with probability  $1 - \epsilon_i$ . Using information theory terminology, we view this operation as placing a binary symmetric channel on each of the features, as shown in Figure 1. We express the resulting *error vector* as  $\mathbf{E} = (E_1, E_2, \dots, E_n)$ , where  $E_i = 1$  if an error has occurred for feature  $i$  and 0 otherwise. As a special case, we consider the vector  $\boldsymbol{\epsilon} = (\epsilon, \epsilon, \dots, \epsilon)$ , where the error probability is identical on all features. In this setting, the probability of some error instantiation  $\mathbf{e}$  is a function of its Hamming weight  $wt(\mathbf{e})$ :

$$Pr(\mathbf{e}) = \epsilon^{wt(\mathbf{e})} (1 - \epsilon)^{n - wt(\mathbf{e})}.$$

#### B. Same Classification Probability

Now we examine the effect of the feature bit errors on the output of the classification algorithm. Our goal is to determine when noise does not impact the algorithm output. That is, we wish to compute the probability that the noiseless point  $\mathbf{X}$  and the noisy version  $\mathbf{X}(\mathbf{E}) = (x'_1, \dots, x'_n)$  have the same classification:

$$SCP(\mathbf{X}, \mathbf{E}) = Pr(c(\mathbf{X}) = c(\mathbf{X}(\mathbf{E}))). \quad (2)$$

We call this quantity the (real) *same classification probability*. We are mainly concerned with an empirical version of the SCP for a feature realization  $\mathbf{x}$ ,

$$SCP(\mathbf{x}, \mathbf{E}) = \sum_{\mathbf{e} \in \{0,1\}^n} Pr(c(\mathbf{x}) = c(\mathbf{x}(\mathbf{e}))) Pr(\mathbf{e}). \quad (3)$$

The SCP resembles the same decision probability (SDP) introduced in [26], defined as the probability of producing a decision (using evidence  $\mathbf{e}$ ) confirmed by  $Pr(d|\mathbf{e} \geq T)$  as when some hidden variable  $\mathbf{H}$  is revealed,  $\sum_{\mathbf{h}} Pr(d|\mathbf{e}, \mathbf{h}) Pr(\mathbf{h}|\mathbf{e})$ . A key difference from SDP is that SCP considers two separate distributions: the augmented distribution that captures the data generation process (Figure 1), and the original naïve Bayes network where classification is performed. Instead, SDP is calculated a single Bayesian classifier that is assumed to capture both processes.

If  $e_j = 1$  and thus there is a feature bit flip at position  $j$ , the corresponding term in (1) changes as

$$\log \left( \frac{\alpha_j}{\beta_j} \right) \rightarrow \log \left( \frac{1 - \alpha_j}{1 - \beta_j} \right).$$

We denote  $A_j := \log \left( \frac{\alpha_j}{\beta_j} \right)$  and  $B_j := \log \left( \frac{1 - \alpha_j}{1 - \beta_j} \right)$ . Let us write  $D_j$  for the difference

$$D_j = B_j - A_j = \log \left( \frac{1 - \alpha_j}{1 - \beta_j} \right) - \log \left( \frac{\alpha_j}{\beta_j} \right).$$

An error  $e_j = 1$  replaces  $A_j$  in (1) with  $B_j$ ; equivalently,  $D_j$  is added to this sum. Therefore, for some error vector  $\mathbf{e}$  with  $\ell$  errors given by the components  $e_{i_1}, e_{i_2}, \dots, e_{i_\ell} = 1$ , the noisy classification is given by

$$c(\mathbf{x}(\mathbf{e})) = \text{sgn} \left( \log \frac{p(c=0)}{p(c=1)} + \sum_{i=1}^n A_i + \sum_{j=1}^{\ell} D_{i_j} \right). \quad (4)$$

Producing the same classification requires  $c(\mathbf{x}(\mathbf{e})) = c(\mathbf{x})$ , or, equivalently,

$$\begin{aligned} & \text{sgn} \left( \log \frac{p(c=0)}{p(c=1)} + \sum_{i=1}^n \log(\alpha_i/\beta_i) \right) \\ &= \text{sgn} \left( \log \frac{p(c=0)}{p(c=1)} + \sum_{i=1}^n A_i + \sum_{j=1}^{\ell} D_{i_j} \right). \end{aligned} \quad (5)$$

Without loss of generality, fix the noiseless classifier output to  $c = 0$ , so that the sign function value is positive<sup>1</sup>. Next, write

$$T := - \left( \log \frac{p(c=0)}{p(c=1)} + \sum_{i=1}^n A_i \right).$$

Here,  $T$  represents a target value. Then, the equality (5) is equivalent to

$$\sum_{j=1}^{\ell} D_{i_j} > T \quad (6)$$

Recall that the  $D_{i_j}$  correspond to all the components in  $\mathbf{e}$  that are equal to 1. Then, to sum over all  $\mathbf{e}$  in (3), we must determine all the subsets of  $\{D_1, D_2, \dots, D_n\}$  with sum greater than the target  $T$ . A subset of size  $\ell$  corresponds to an error vector with  $\ell$  feature bit errors that does not change the classification relative to the noiseless version. Such a subset contributes a term (given by (3)) to the same classification probability.

**Example:** Consider a model with  $n = 3$  binary features and equal error probability  $\epsilon = (\epsilon, \epsilon, \epsilon)$ . Take  $D_1 + D_2 > T$ . This implies that the error vector  $\mathbf{e} = (1, 1, 0)$  does not change the classification output, contributing a probability term of  $\epsilon^2(1 - \epsilon)$  to the SCP.

#### IV. SCP APPROXIMATION

In this section, we consider issues surrounding the computation of the SCP. Our main result is an efficient SCP approximation.

##### A. Exact Empirical SCP

The problem described by (6) is a variant of the subset sum problem with real numbers. The naïve approach is to enumerate the  $2^n$  subsets of  $\{D_1, D_2, \dots, D_n\}$  and compute each of their sums. This task is not tractable for large  $n$ .

A compact representation of the SCP is the function

$$G(z) = \prod_{i=1}^n ((1 - \epsilon_i) + \epsilon_i z^{D_i})$$

We define  $[> T]G(z)$  to be the sum of the coefficients of terms in  $G(z)$  with exponent larger than  $T$ . It is easy to see that

$$\text{SCP}(\mathbf{x}, \mathbf{E}) = [> T]G(z).$$

Any subset  $\{D_{i_1}, D_{i_2}, \dots, D_{i_j}\}$  with sum greater than  $T$  produces a term in  $G(z)$  with exponent  $D_{i_1} + \dots + D_{i_j} > T$  and coefficient

$$\epsilon_{i_1} \epsilon_{i_2} \dots \epsilon_{i_j} \prod_{i \neq i_1, \dots, i_j} (1 - \epsilon_i).$$

The sum of all such coefficients is indeed the SCP. This may lead us to attempt to expand the function  $G(z)$  and examine the resulting coefficients. However, the fact that the exponents are real-valued prevents us from doing so efficiently.

##### B. Quantization

Observe that if the exponents of terms in  $G(z)$  were non-negative integers,  $G(z)$  would be a polynomial, allowing for fast multiplication. In fact,  $G(z)$  would become a *generating function* [33]. Inspired by this notion, we introduce the following quantization scheme.

The key idea is to quantize  $D_i$  into  $k$  buckets, for  $k$  a constant. By performing the quantization in a clever way, we induce a structure that enables us to approximate the SCP in no more than  $O(n^2 \log n)$  operations. The approximation is described in Algorithm 1; the concept is detailed below.

Our quantization scheme only relies on the minimal and maximal values of  $\{D_1, \dots, D_n\}$ . Let  $D_{\min} = \min\{D_1, \dots, D_n\}$  and  $D_{\max} = \max\{D_1, \dots, D_n\}$ . Let  $D_I = D_{\max} - D_{\min}$ . Consider the family of  $k$  intervals

<sup>1</sup>If the noiseless classifier output is  $c = 1$ , the outcome is that we seek subsets whose sum is less than the target value  $T$  rather than greater. All of the arguments we present in this paper are unchanged. For simplicity of notation, we do not add the case to our setup.

---

**Algorithm 1** SCP Approximation
 

---

**Input:** Difference terms  $D_1, D_2, \dots, D_n$ , Error probabilities  $\epsilon_1, \epsilon_2, \dots, \epsilon_n$ , Target  $T$ , Number of buckets in quantization scheme  $k$

**Output:** SCP approximation  $SCP_{app}$

**Initialize**  $S_1, \dots, S_k$  to 0,  $SCP_{app}$  to 0

**for**  $j = 1$  **to**  $n$  **do**

**if**  $D_j \in [E_i^{\text{begin}}, E_i^{\text{end}})$  **then**

$S_i \leftarrow S_i + 1$

$\epsilon_{iS_i} \leftarrow \epsilon_j$

**end if**

**end for**

**expand**  $G[i, j] = \prod_{i=1}^k \prod_{j=1}^{S_i} ((1 - \epsilon_{ij}) + \epsilon_{ij}yz^i)$

**for**  $i = 0$  **to**  $n$  **do**

$T' \leftarrow \frac{k-1}{D_I} \left( T - i \left( D_{\min} - \frac{D_I}{k-1} \right) \right)$

**for**  $j \geq 0$  **do**

**if**  $G[i, j] > T'$  **then**

$SCP_{app} \leftarrow SCP_{app} + G[i, j]$

**end if**

**end for**

**end for**

---

$$\left[ D_{\min} - \frac{D_I}{2(k-1)}, D_{\min} + \frac{D_I}{2(k-1)} \right), \left[ D_{\min} + \frac{D_I}{2(k-1)}, D_{\min} + \frac{3D_I}{2(k-1)} \right), \dots, \left[ D_{\min} + \frac{(2k-5)D_I}{2(k-1)}, D_{\min} + \frac{(2k-3)D_I}{2(k-1)} \right), \left[ D_{\min} + \frac{(2k-3)D_I}{2(k-1)}, D_{\max} + \frac{(2k-1)D_I}{2(k-1)} \right). \quad (7)$$

Note that each bucket has width  $D_I/(k-1)$ . For compactness, we also write the  $i$ th interval as  $[E_i^{\text{begin}}, E_i^{\text{end}})$ .

We quantize any value of  $D_j$  that falls into a bucket with the midpoint of the bucket interval. We compute  $S_i$ , the number of  $D_j$  in each bucket:

$$S_i = \left| \left\{ D_j : D_j \in \left[ D_{\min} + \frac{(2i-3)D_I}{2(k-1)}, D_{\min} + \frac{(2i-1)D_I}{2(k-1)} \right), 1 \leq j \leq n \right\} \right|,$$

for  $1 \leq i \leq k$ . In addition, we re-label the noise parameters such that  $\epsilon_{i1}, \epsilon_{i2}, \dots, \epsilon_{iS_i}$  are the noise parameters corresponding to the features  $D_j$  that fall into the  $i$ th bucket.

Recall that we seek to determine the subsets of  $\{D_1, D_2, \dots, D_n\}$  that have sum greater than  $T$ . We approximate any  $D_j$  that falls into the  $i$ th interval with the midpoint of that interval:

$$D_{\min} + \frac{i-1}{k-1} D_I. \quad (8)$$

Additionally, observe that  $D_{\min}$  and  $D_{\max}$  falls are represented by their own values.

We approximate the SCP by computing the SCP on the quantized versions of the  $D_i$ ; we show that this can be performed efficiently by Algorithm 1.

Consider the multivariate generating function

$$G(y, z) = \prod_{i=1}^k \prod_{j=1}^{S_i} ((1 - \epsilon_{ij}) + \epsilon_{ij}yz^i). \quad (9)$$

We write  $[\ell, > R]G(y, z)$  for the sum of all coefficients of terms with exponent  $\ell$  in  $y$  and exponent greater than  $R$  in  $z$ . Note that in the case where  $\epsilon = (\epsilon, \dots, \epsilon)$ ,  $G(y, z)$  reduces to

$$G(y, z) = \prod_{i=1}^k ((1 - \epsilon) + \epsilon yz^i)^{S_i}.$$

We show that the appropriate coefficients in the generating function yield the SCP for the quantized versions of the  $D_i$ :

**Theorem 1.** Under the quantization scheme given by Algorithm 1, the resulting SCP is given by

$$\text{SCP}_{app}(\mathbf{x}, \mathbf{E}) = \sum_{\ell=0}^n [\ell, > T'(\ell)] G(y, z),$$

where

$$T'(\ell) := \frac{k-1}{D_I} \left( T - \ell \left( D_{min} - \frac{D_I}{k-1} \right) \right).$$

*Proof:* Consider some subset with size  $\ell$  and sum greater than the target  $T$  in the quantization scheme given by (7). We distribute our  $\ell$  choices of  $D_j$  into the  $k$  buckets, where we may take up to  $S_i$  values from bucket  $i$ . Let  $(a_1, a_2, \dots, a_k)$  be such a choice;  $a_i$ , for  $0 \leq a_i \leq S_i$  represents the number of  $D_j$  falling into the  $i$ th bucket. A subset of size  $\ell$  requires  $a_1 + a_2 + \dots + a_k = \ell$ . The total sum for the subset is given by

$$a_1 D_{min} + a_2 \left( D_{min} + \frac{D_I}{k-1} \right) + \dots + a_k D_{max}.$$

Here we used the fact that terms in each bucket are quantized to the midpoint of that bucket (i.e., expression (8)).

Since the sum of the subset is greater than  $T$ , we may write

$$\begin{aligned} (a_1 + a_2 + \dots + a_k) \left( D_{min} - \frac{D_I}{k-1} \right) + \left( a_1 \frac{D_I}{k-1} + a_2 \frac{2D_I}{k-1} + \dots + a_k \frac{kD_I}{k-1} \right) &> T \\ \implies \ell \left( D_{min} - \frac{D_I}{k-1} \right) + \frac{D_I}{k-1} (a_1 + 2a_2 + \dots + ka_k) &> T \\ \implies a_1 + 2a_2 + \dots + ka_k > \frac{k-1}{D_I} \left( T - \ell \left( D_{min} - \frac{D_I}{k-1} \right) \right) &:= T'. \end{aligned}$$

Now we relate the generating function given in (9) to the equation  $a_1 + 2a_2 + \dots + ka_k > T'$ . The generating function is the product of  $n$  binomials. To generate a term in the expansion, we select from each binomial either  $1 - \epsilon_{ij}$  or  $\epsilon_{ij} y z^i$ . The subsets of size  $\ell$  satisfying  $a_1 + 2a_2 + \dots + ka_k > T'$  correspond to choosing  $a_1$  terms  $\epsilon_{1j} y z$ ,  $a_2$  terms  $\epsilon_{2j} y z^2$ , and so on. The products of such terms have degree  $\ell$  in  $y$  and degree  $a_1 + 2a_2 + \dots + ka_k > T'$  in  $z$ . Considering the subsets of size  $0, 1, \dots, n$ , we indeed have that  $\text{SCP}_{app}(\mathbf{x}, \mathbf{E}) = \sum_{\ell=0}^n [\ell, > T'] G(y, z)$ . ■

### C. Computation

To compute  $\text{SCP}_{app}(\mathbf{x}, \mathbf{E})$ , we must perform the expansion of the generating function  $G(y, z) = \prod_{i=1}^k \prod_{j=1}^{S_i} ((1 - \epsilon_{ij}) + \epsilon_{ij} y z^i)$  and examine the coefficients. Polynomial multiplication is equivalent to convolution (of vectors or arrays, depending on whether the polynomials are univariate or multivariate); this effort can be sped up by using the Fast Fourier Transform (fft) [34]. In the Fourier domain, convolution is equivalent to multiplication. That is, for two polynomials  $q(x), r(x)$ ,

$$q(x)r(x) = \text{ifft}(\text{fft}(q) \times \text{fft}(r)),$$

where the multiplication on the right side is performed term by term and  $\text{ifft}$  denotes the inverse transform. If the maximum degree of  $q(x)$  and  $r(x)$  is  $d$ , the number of operations required is  $O(d \log d)$ .

First, we expand the terms  $\prod_{j=1}^{S_i} ((1 - \epsilon_{ij}) + \epsilon_{ij} y z^i)$ . Replace  $yz^i$  by  $v$ ; then, we are multiplying  $S_i$  polynomials of degree 1. This can be performed in  $O(S_i \log \log S_i)$  operations. This fact is not difficult to check: our  $S_i$  monomials can be paired up and the pairs multiplied; the  $S_i/2$  resulting degree-2 polynomials are paired up, and so on. The total time is  $O(S_1 \log \log S_1) + O(S_2 \log \log S_2) + \dots + O(S_k \log \log S_k) \leq O(n \log \log n)$ . We note that each expansion can be performed in parallel.

Note that in the case of identical  $\epsilon$ , the polynomial to be expanded is  $((1 - \epsilon) + \epsilon y z^i)^{S_i}$ , and the expansion can be performed in  $O(S_i)$  operations with the binomial theorem.

Next, each of the resulting expansions must be multiplied. The expansion of  $\prod_{j=1}^{S_i} ((1 - \epsilon_{ij}) + \epsilon_{ij} y z^i)$  can be represented as a two-dimensional array, where rows represent the exponent of  $y$  and columns the exponent of  $z$ . In our case, we have  $k$  arrays. The maximal degree for  $y$  is  $n$ , while the maximal degree for  $z$  is  $kn$ , so that the product of all of the arrays (and thus any intermediate multiplication) has size at most  $n \times kn$ . Each multiplication has cost at most  $O(kn^2 \log(kn))$ ; the total cost is thus no more than  $O(k^2 n^2 \log(kn))$ . (Again, it is possible to speed to perform the multiplications in parallel for speedup.)

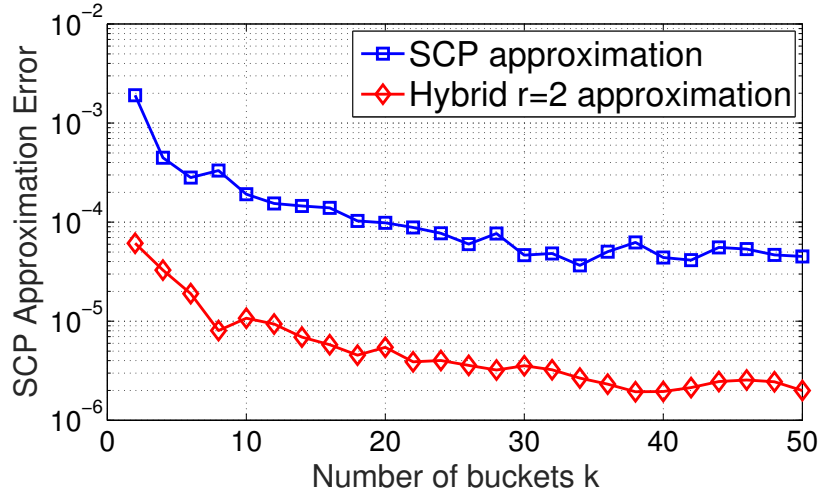


Fig. 2. SCP approximation error (against the real SCP) for the two approximation schemes. The parameters here are  $\epsilon = 10^{-2}$  and  $n = 16$ . 50 data points were selected randomly; the true SCP was computed and the average absolute error was measured for various  $k$  in the two schemes.

#### D. Approximation Quality

Next we comment on the quality of the approximation. Clearly, the quality is a function of the number of buckets  $k$ ; the larger the  $k$ , the finer the approximation. We note, however, that the approximation error is not a monotonic function, since there are edge effects for different  $k$ , as can be seen in Figure 2.

We introduce two improvements to optimize the worst-case error of the SCP approximation. We describe these errors as a function of the noise parameter  $\epsilon$ , and for simplicity take the case where all the components equal to  $\epsilon$ . The first improvement comes at no extra cost and reduces the worst-case error to  $O(\epsilon^2)$ , while the second improvement requires  $O(n^r)$  time but reduces the worst-case error to  $O(\epsilon^{r+1})$ . We adopt the first improvement for our experiments.

**Improvement 1.** For the first improvement, consider the intervals defined in (7). If our target value  $T$  lines somewhere in  $[D_{min}, D_{max}]$ , we can translate (shift) all of our intervals over by some constant  $T_{shift}$  that is no larger than half an interval width,  $D_I/2(k-1)$ , such that  $T$  is now precisely the edge between two intervals. Then, any  $D_i < T$  is mapped to a bucket center smaller than  $T$ , while any  $D_i \geq T$  is mapped to a bucket center larger than  $T$ . We conclude that any single-element subset of  $\{D_1, \dots, D_n\}$  that contributes to the SCP also contributes to the SCP approximation; thus, all SCP approximation errors must be for two- or more element subsets, reducing this error to  $O(\epsilon^2)$ .

**Improvement 2.** We refer to the second improvement as a *hybrid SCP approximation*. The idea here is simple: we use the approximation only for the subsets of size greater than  $r$  for some constant  $2 \leq r < n$ . For those subsets that are of size  $r$  or smaller, we check the sums using the actual values of  $D_i$ , requiring at most  $O(n^r)$  time. This strategy is most suitable for very small values of  $r$ , such as  $r = 2$ .

**Example.** An evaluation was performed on the small data set house-votes-84 from the UCI repository [35]. This set has  $n = 16$  binary features corresponding to ‘yes’ or ‘no’ votes on various congressional proposals; the binary class represents democrat or republican. The small size of  $n$  enables us to compute the real SCP value (though much slower than the approximate SCP). Figure 2 shows the error versus the true SCP for the two approximation schemes and various  $k$  averaged over 50 random test points; the noise parameter was  $\epsilon = 10^{-2}$ .

## V. OPTIMIZED CHANNEL CODING

Thus far we have only considered the impact of noise without taking any action to protect the algorithm output from distortion. In this section, we consider a simple but effective form of error correction. Our goal is to tailor this strategy to the characteristics of the features in order to minimize the classifier distortion. This is a departure from traditional error protection, which is typically present in a different abstraction layer and is agnostic to the algorithm.

Consider storing multiple copies of a certain feature bit; these copies are produced when the data is written and before the addition of noise. Afterwards, when the bit is ready to be used, the repeated values of the bit (some of which may now be corrupted) are subject to a majority vote (equivalent to maximum-likelihood decoding) and the decoded value is placed back into the algorithm.

What is the effect of repetition on the error probability  $\epsilon$ ? If the  $j$ th feature  $x_j$ , corrupted by a bit flip with probability  $\epsilon$ , is repeated  $2r + 1$  times for some  $r \geq 0$ , the probability  $\epsilon^{(r)}$  that the feature is decoded incorrectly is equal to the probability

**Algorithm 2** Channel Code Optimization

---

**Input:** Test points  $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_t$ , Noise vector  $\boldsymbol{\epsilon} = (\epsilon_1, \dots, \epsilon_n)$ , redundancy budget  $r$   
**Output:** Optimized redundancy allocation vector  $\mathbf{r}^r$   
**Initialize**  $\mathbf{r}(0)$  to 0  
**for**  $j = 0$  **to**  $r - 1$  **do**  
     $i \leftarrow \arg \min_i \frac{1}{i} \sum_{i=1}^t \text{SCP}_{app}(\mathbf{x}_t, \mathbf{E}^{\mathbf{r}(j)+\mathbf{1}(i)})$   
     $\mathbf{r}(j+1) \leftarrow \mathbf{r}(j) + \mathbf{1}(i)$   
**end for**

---

that the majority of votes are corrupted:

$$\epsilon^{(r)} = \sum_{i=r+1}^{2r+1} \binom{2r+1}{i} \epsilon^i (1-\epsilon)^{2r+1-i}.$$

For example, if  $r = 1$ , the smallest meaningful repetition,  $\epsilon^{(1)} = \epsilon^3 + 3\epsilon^2(1-\epsilon)$ . Note that we always add pairs of repetitions, since an even number of repeated bits can produce ties. In general, employing  $r$  pairs of repeated bits drives the error probability  $\epsilon$  to  $O(\epsilon^{1+r})$ .

### A. Monotonicity

Intuitively, we expect that reducing the error probability on features will increase the SCP. Clearly, reducing all of the  $\epsilon_i$  to 0 yields an SCP of 1, as the only non-zero term in the SCP is  $(1-\epsilon_1)(1-\epsilon_2) \times \dots \times (1-\epsilon_n) = 1$ .

Nevertheless, the SCP does not necessarily increase when we reduce noise parameters  $\epsilon$ . The following result is applicable to both the SCP and the SCP approximation:

**Theorem 2.** Let  $\mathbf{x}$  be a feature vector and  $\boldsymbol{\epsilon} = (\epsilon_1, \dots, \epsilon_{i-1}, \epsilon_i, \epsilon_{i+1}, \dots, \epsilon_n)$  and  $\boldsymbol{\epsilon}' = (\epsilon_1, \dots, \epsilon_{i-1}, \epsilon'_i, \epsilon_{i+1}, \dots, \epsilon_n)$  be noise parameter vectors with  $\epsilon_i \geq \epsilon'_i$ . Let the associated error vectors be  $\mathbf{E}, \mathbf{E}'$ , respectively. If  $D_i > 0$ , then  $\text{SCP}(\mathbf{x}, \mathbf{E}) \geq \text{SCP}(\mathbf{x}, \mathbf{E}')$ , while if  $D_i \leq 0$ , then  $\text{SCP}(\mathbf{x}, \mathbf{E}) \leq \text{SCP}(\mathbf{x}, \mathbf{E}')$ .

*Proof:* Every summand in  $\text{SCP}(\mathbf{x}, \mathbf{E})$  is either a multiple of  $\epsilon_i$  (if it corresponds to a subset of the  $D_j$ s that includes  $D_i$ ) or a multiple of  $(1-\epsilon_i)$  (if it corresponds to a subset of the  $D_j$  that does not include  $D_i$ ). We write

$$\text{SCP}(\mathbf{x}, \mathbf{E}) = \epsilon_i C_1 + (1-\epsilon_i) C_2,$$

where  $C_1, C_2$  are non-negative.

*Case 1)  $D_i > 0$ .* Consider the subsets of  $\{D_1, \dots, D_n\} \setminus \{D_i\}$  corresponding to  $C_1$ . These subsets must have sum larger than  $T - D_i$ , since including  $D_i$ , their sum must exceed  $T$ . In the case of  $C_2$ , these subsets have sum larger than  $T$ . Since  $D_i$  is positive, any subset with sum larger than  $T$  has sum larger than  $T - D_i$  as well, so all subsets corresponding to  $C_2$  also correspond to terms in  $C_1$ . We may write  $C_1 = C_2 + C_3$  for some  $C_3 \geq 0$ . Then,

$$\begin{aligned} \text{SCP}(\mathbf{x}, \mathbf{E}) &= \epsilon_i C_1 + (1-\epsilon_i) C_2 \\ &= \epsilon_i (C_2 + C_3) + (1-\epsilon_i) C_2 \\ &= C_2 + \epsilon_i C_3. \end{aligned}$$

Since  $C_3 \geq 0$ , reducing  $\epsilon_i$  to  $\epsilon'_i$  reduces the SCP as well.

*Case 2)  $D_i \leq 0$ .* Then,  $T - D_i \geq T$ , so all subsets corresponding to  $C_1$  also correspond to  $C_2$ . Thus, we can write

$$\begin{aligned} \text{SCP}(\mathbf{x}, \mathbf{E}) &= \epsilon_i C_1 + (1-\epsilon_i) C_2 \\ &= \epsilon_i C_1 + (1-\epsilon_i) (C_1 + C_3) \\ &= C_1 + (1-\epsilon_i) C_3. \end{aligned}$$

Reducing  $\epsilon_i$  to  $\epsilon'_i$  increases  $(1-\epsilon_i) C_3$ , so the SCP is increased. ■

In words, noise helps the SCP when applied to those features that disagree with the classification, since flipping these bits increases confidence (and can cover up for other bit flips that reduce it). However, noise hurts the SCP when applied to features that agree with the classification; the sign of  $D_i$  is a direct consequence of this idea.

The fact that reducing the error probability (by adding redundancy, etc.) is not always helpful leads us to seek an optimized solution. Clearly a uniform allocation of protection for all features is not always a good idea.



TABLE II  
EXPERIMENTAL RESULTS FOR CLASSIFICATION CHANGE PROBABILITY (1-SCP)

| R  | Movie dataset, feature set 1 |          |        | Movie dataset, feature set 2 |          |        | Voting dataset |         |        | NLTCs dataset |         |               |
|----|------------------------------|----------|--------|------------------------------|----------|--------|----------------|---------|--------|---------------|---------|---------------|
|    | Uniform                      | Optimal  | Ratio  | Uniform                      | Optimal  | Ratio  | Uniform        | Optimal | Ratio  | $R_b$         | Optimal | No protection |
| 1  | 0.10112                      | 0.08664  | 1.1670 | 0.1263                       | 0.10245  | 1.2328 | 0.05635        | 0.04135 | 1.3628 | 6             | 0.0131  | 0.0218        |
| 2  | 0.03484                      | 0.03024  | 1.1524 | 0.04419                      | 0.03028  | 1.4595 | 0.03182        | 0.02377 | 1.3390 | 7             | 0.0120  | 0.0218        |
| 3  | 0.01154                      | 0.01002  | 1.1513 | 0.01469                      | 0.00879  | 1.6711 | 0.01903        | 0.01429 | 1.3319 | 8             | 0.0108  | 0.0218        |
| 4  | 0.00382                      | 0.00331  | 1.1535 | 0.00486                      | 0.00256  | 1.9028 | 0.01154        | 0.0085  | 1.3555 | 9             | 0.0097  | 0.0218        |
| 5  | 0.00127                      | 0.00109  | 1.1576 | 0.00162                      | 0.00077  | 2.1019 | 0.00702        | 0.00504 | 1.3911 | 10            | 0.0086  | 0.0218        |
| 6  | 0.00043                      | 0.00037  | 1.1602 | 0.00055                      | 0.00024  | 2.2549 | 0.00423        | 0.00298 | 1.4328 | 11            | 0.0074  | 0.0218        |
| 7  | 0.00015                      | 0.00013  | 1.1622 | 0.00019                      | 7.70E-05 | 2.3983 | 0.00261        | 0.00177 | 1.4742 | 12            | 0.0062  | 0.0218        |
| 8  | 4.94E-05                     | 4.24E-05 | 1.1636 | 6.30E-05                     | 2.48E-05 | 2.539  | 0.0016         | 0.00105 | 1.5182 | 13            | 0.0052  | 0.0218        |
| 9  | 1.69E-05                     | 1.45E-05 | 1.1646 | 2.16E-05                     | 8.08E-06 | 2.6720 | 0.00098        | 0.00063 | 1.5632 | 14            | 0.0044  | 0.0218        |
| 10 | 5.83E-06                     | 5.00E-06 | 1.1654 | 7.43E-06                     | 2.64E-06 | 2.8109 | 0.00060        | 0.00038 | 1.6107 | 15            | 0.0037  | 0.0218        |

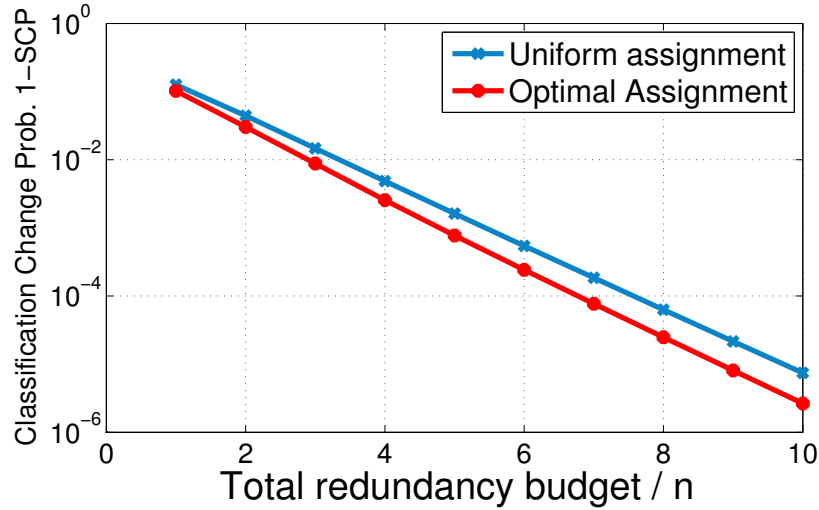


Fig. 3. Classification change probability (1 – SCP) for a movie reviews data set with  $n = 20$  features and a test set of size of  $t = 250$ . Comparison of the results of optimized versus uniform redundancy allocation.

### B. Greedy Optimization

In the remainder of this work, we introduce an algorithm to optimize the allocation of a redundancy budget for coded feature protection. We use a budget of  $2r$  redundancy bits, which will be used to protect our  $n$  features. Feature  $i$  is then represented by  $i + 2r_i$  copies, and has error probability  $\epsilon^{(r_i)}$ , for  $1 \leq i \leq n$ . The noise vector for all  $n$  features  $\epsilon^{(\mathbf{r})} = (\epsilon^{(r_1)}, \epsilon^{(r_2)}, \dots, \epsilon^{(r_n)})$ . The corresponding error vector is written  $\mathbf{E}^{(\mathbf{r})}$ . The values  $\mathbf{r} = (r_1, r_2, \dots, r_n)$  are constrained by  $r_1 + r_2 + \dots + r_n = r$ . To minimize the distortion due to noise on point  $\mathbf{x}$ , we must maximize the SCP with respect to  $\mathbf{r}$  and the resulting  $\epsilon$  vector. We have the following optimization

$$\arg \max_{\mathbf{r}} \text{SCP}(\mathbf{x}, \mathbf{E}^{(\mathbf{r})}) \quad \text{s.t.} \quad \sum_{i=1}^n r_i = r, \quad r_i \geq 0.$$

Our first step is to perform a tractable version of this optimization using our SCP approximation:

$$\arg \max_{\mathbf{r}} \text{SCP}_{app}(\mathbf{x}, \mathbf{E}^{(\mathbf{r})}) \quad \text{s.t.} \quad \sum_{i=1}^n r_i = r, \quad r_i \geq 0.$$

Although we can now efficiently check each value of  $\mathbf{r}$ , doing so is still computationally expensive. The number of solutions to  $r_1 + r_2 + \dots + r_n = r$  with  $r_i \geq 0$  is given by  $\binom{n+r-1}{r}$ . We turn to a greedy (myopic) optimization approach to further reduce the complexity. This approach enables us to perform the optimization one redundancy unit (two repeated bits) at a time. After the  $j$ th step, we write the redundancy vector as  $\mathbf{r}(j) = (r_1(j), r_2(j), \dots, r_n(j))$ . In the  $(j+1)$ st step, one of the  $r_i(j)$  terms is selected and increased by 1. For ease of notation, let us write  $\mathbf{1}(i)$  for the vector  $(0, 0, \dots, 0, 1, 0, \dots, 0)$  with a 1 in the  $i$ th position and 0s elsewhere. Then, the  $(j+1)$ st step is given by

$$\arg \max_i \text{SCP}_{app}(\mathbf{x}, \mathbf{E}^{\mathbf{r}(j)+\mathbf{1}(i)}) \quad \text{s.t.} \quad 1 \leq i \leq n.$$

This reduces our complexity to that of performing  $n$  SCP computations per each of the  $r$  steps. Of course, we can also perform the optimization over  $t$  test data points. The procedure is given in Algorithm 2

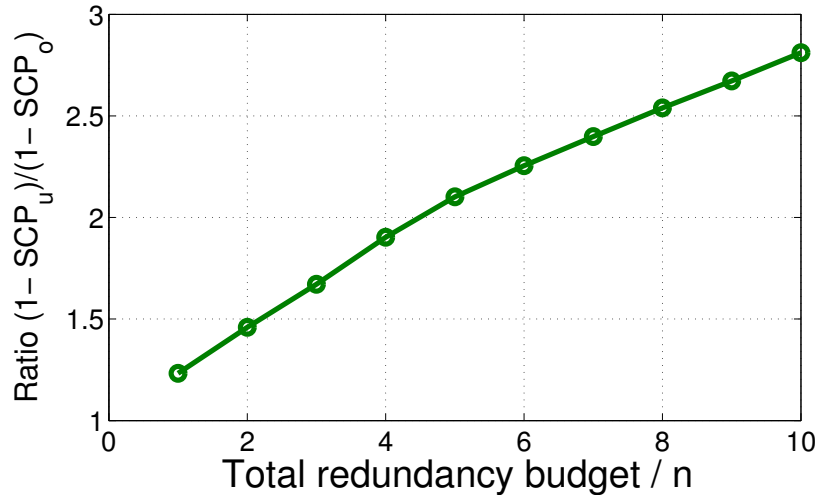


Fig. 4. The classification change probability (uniform to optimized allocation) ratio reaches  $3\times$ . The setup is the same as that of Figure 2.

We demonstrate our approach with SCP experiments using the voting dataset [35], the movie review dataset from [36], and the NLTCs dataset in [37]. In all cases, we used  $k = 50$  for the number of buckets in our SCP approximations. For the movie reviews set we used two sets of  $n = 20$  features and 250 test data points. The  $\epsilon$  parameter was 0.2 for the first 2 datasets (a very high amount of noise that can nevertheless be handled through channel coding). Table II shows the classification change  $(1 - \text{SCP})$  probability for uniform and optimized assignment of redundancy bits given a budget of  $2nR$  bits. We report the ratio of the classification change probability (uniform versus optimized); this ratio reaches up to  $3\times$ .

In the movie review dataset, the first feature set included features with similar dependence on the class variable. The second feature set included a mixture of features (some more informative of the class variable than others). In this case the optimized assignment has a clear benefit over the uniform assignment. The results on the voting dataset (using  $n = 16$  features and  $t = 454$  test data points) are in between those of the two movie review sets. A plot of the SCP and the improvement ratio for the more beneficial assignment is shown in Figures 3 and 4.

For the NLTCs dataset, we tested the case of a small amount of redundancy  $R_b \leq 15$ , so that a uniform allocation is impossible, and we must rely on the optimization. Here,  $n = 15$  features and  $t = 3236$  test data points were used. The results are compared to the unprotected version.

## VI. CONCLUSION

We studied binary naïve Bayes classifiers operating on noisy test data. First, we characterized the impact of bit flips due to noise on the classifier output with the same classification probability (SCP). We introduced a low-complexity approximation for the SCP based on quantization and polynomial multiplication. Next, we considered minimizing the classifier distortion by allocating redundant bits among the features. Our informed approach for redundancy allocation is among the first principled methods combining coding theory with machine learning. We demonstrated the results of this idea with experiments.

## REFERENCES

- [1] C. E. Shannon, "Communication in the presence of noise," *Proceedings of the IEEE*, vol. 86, no. 2, pp. 447–457, 1998.
- [2] L. Dilillo, P. Girard, S. Pravossoudovitch, A. Virazel, and M. B. Hage-Hassan, "Data retention fault in SRAM memories: analysis and detection procedures," in *Proceedings of the 23rd IEEE VLSI Test Symposium (VTS)*, 2005.
- [3] S. Mittal, "A survey of techniques for approximate computing," *ACM Computing Surveys*, vol. 48, no. 4, 2016.
- [4] V. Chandra and R. Aitken, "Impact of technology and voltage scaling on the soft error susceptibility in nanoscale CMOS," in *Proc. Defect and Fault Tolerance of VLSI Systems (DFTVS '08)*, (Boston, MA), 2008.
- [5] V. Sridharan and D. Liberty, "A field study of DRAM errors," in *Proc. International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*, (Salt Lake City, UT), 2012.
- [6] C. M. Riggle and S. G. McCarthy, "Design of error correction systems for disk drives," *IEEE Transactions on Magnetics*, vol. 34, no. 4, pp. 2362–2371, 1998.
- [7] R. L. Rivest and A. Shamir, "How to reuse a "write-once memory"," *Information and Control*, vol. 55, no. 1-3, pp. 1–19, 1982.
- [8] M. Blaum, J. Brady, J. Bruck, and J. Menon, "EVENODD: an efficient scheme for tolerating double disk failures in RAID architectures," *IEEE Transactions on Computers*, vol. 44, no. 2, pp. 192–202, 1995.
- [9] Y. Cassuto, M. Schwartz, V. Bohossian, and J. Bruck, "Codes for asymmetric limited-magnitude errors with application to multilevel flash memories," *IEEE Transactions on Information Theory*, vol. 56, no. 4, pp. 1582–1595, 2010.
- [10] L. Dolecek and F. Sala, "Channel coding methods for non-volatile memories," *Foundations and Trends in Communications and Information Theory*, vol. 13, no. 1, pp. 1–136, 2016.
- [11] H. Weatherspoon and J. Kubiatowicz, "Erasure coding vs. replication: A quantitative comparison," in *Proceedings of First International Workshop on Peer-to-Peer Systems (IPTPS '01)*, 2002.

- [12] A. G. Dimakis, P. B. Godfrey, Y. Wu, M. J. Wainwright, and K. Ramchandran, "Network coding for distributed storage systems," *IEEE Transactions on Information Theory*, vol. 56, no. 9, pp. 4539–4551, 2010.
- [13] K. Zhao, W. Zhao, H. Sun, T. Zhang, X. Zhang, and N. Zheng, "LDPC-in-SSD: Making advanced error correction codes work effectively in solid state drives," in *Proceedings of Conference on File and Storage Technologies (FAST '13)*, (San Jose, CA), 2013.
- [14] K. Mazooji, F. Sala, G. Van den Broeck, and L. Dolecek, "Robust channel coding strategies for machine learning data," in *Proc. Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pp. 609–616, 2016.
- [15] N. Srivastava, G. E. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: a simple way to prevent neural networks from overfitting," *Journal of Machine Learning Research*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [16] A. Achille and S. Soatto, "Information dropout: Learning optimal representations through noisy computation," tech. rep., University of California, Los Angeles, 2016.
- [17] E. Kalapanidas, N. Avouris, M. Craciun, and D. Neagu, "Machine learning algorithms: a study on noise sensitivity," in *Proc. 1st Balcan Conference in Informatics*, (Thessaloniki, Greece), pp. 356–365, 2003.
- [18] O. Dekel and O. Shamir, "Learning to classify with missing and corrupted features," in *Proc. 25th International Conference on Machine Learning (ICML 2008)*, (Helsinki, Finland), 2008.
- [19] A. Globerson and S. Roweis, "Nightmare at test time: robust learning by feature deletion," in *Proc. 25th International Conference on Machine Learning (ICML 2006)*, (Pittsburgh, PA), 2006.
- [20] M. Ramoni and P. Sebastiani, "Robust bayes classifiers," *Artificial Intelligence*, vol. 125, no. 1-2, pp. 209–226, 2001.
- [21] J. C. Duchi, M. I. Jordan, M. J. Wainwright, and Y. Zhang, "Optimality guarantees for distributed statistical estimation," tech. rep., University of California, Berkeley, 2014.
- [22] T. S. Han and S. Amari, "Statistical inference under multiterminal data compression," *IEEE Transactions on Information Theory*, vol. 44, no. 6, pp. 2300–2324, 1998.
- [23] D. Heckerman, E. Horvitz, and B. Middleton, "An approximate nonmyopic computation for value of information," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 15, no. 3, pp. 292–298, 1993.
- [24] M. Bilgic and L. Getoor, "Value of information lattice: Exploiting probabilistic independence for effective feature subset acquisition," *Journal of Artificial Intelligence Research (JAIR)*, vol. 41, pp. 69–95, 2011.
- [25] A. Krause and C. Guestrin, "Optimal value of information in graphical models," *Journal of Artificial Intelligence Research (JAIR)*, vol. 35, pp. 557–591, 2009.
- [26] A. Choi, Y. Xue, and A. Darwiche, "Same-Decision Probability: A confidence measure for threshold-based decisions," *International Journal of Approximate Reasoning (IJAR)*, vol. 2, pp. 1415–1428, 2012.
- [27] S. Chen, A. Choi, and A. Darwiche, "Algorithms and applications for the Same-Decision Probability," *Journal of Artificial Intelligence Research (JAIR)*, vol. 49, pp. 601–633, 2014.
- [28] S. Chen, A. Choi, and A. Darwiche, "Computer adaptive testing using the same-decision probability," in *Proceedings of the Twelfth UAI Conference on Bayesian Modeling Applications Workshop*, pp. 34–43, 2015.
- [29] F. J. Gimenez, Y. Wu, E. S. Burnside, and D. L. Rubin, "A novel method to assess incompleteness of mammography reports," in *AMIA Annual Symposium Proceedings*, vol. 2014, p. 1758, American Medical Informatics Association, 2014.
- [30] B. Frenay and A. Kaban, "A comprehensive introduction to label noise," in *Proc. European Symposium on Artificial Neural Networks, Comp. Intelligence and Machine Learning (ESANN)*, (Bruges, Belgium), 2014.
- [31] A. Vempaty, L. R. Varshney, and P. K. Varshney, "Reliable crowdsourcing for multi-class labeling using coding theory," *IEEE Journal of Selected Topics in Signal Processing*, vol. 8, no. 4, pp. 667–679, 2014.
- [32] I. Guyon and A. Elisseeff, "An introduction to variable and feature selection," *Journal of Machine Learning Research*, vol. 3, pp. 1157–1182, 2003.
- [33] H. S. Wilf, *Generatingfunctionology*. A K Peters/CRC Press, 3rd ed., 2005.
- [34] K. R. Rao, D. N. Kim, and J.-J. Hwang, *Fast Fourier Transform - Algorithms and Applications*. Springer Publishing Company, 1st ed., 2010.
- [35] M. Lichman, "UCI machine learning repository," 2013.
- [36] J. Bekker, J. Davis, A. Choi, A. Darwiche, and G. Van den Broeck, "Tractable learning for complex probability queries," in *Proceedings of Conference on Neural Information Processing Systems (NIPS)*, (Montreal, Canada), 2015.
- [37] D. Lowd and J. Davis, "Learning markov network structure with decision trees," in *Proc. IEEE International Conference on Data Mining (ICDM)*, (Sydney, Australia), 2010.