

Learning to reconstruct signals from binary measurements alone

Anonymous authors

Paper under double-blind review

Abstract

Recent advances in unsupervised learning have highlighted the possibility of learning to reconstruct signals from noisy and incomplete linear measurements alone. These methods play a key role in medical and scientific imaging and sensing, where ground truth data is often scarce or difficult to obtain. However, in practice measurements are not only noisy and incomplete but also quantized. Here we explore the extreme case of learning from binary observations and provide necessary and sufficient conditions on the number of measurements required for identifying a set of signals from incomplete binary data. Our results are complementary to existing bounds on signal recovery from binary measurements. Furthermore, we introduce a novel self-supervised learning approach, which we name SSBM, that only requires binary data for training. We demonstrate in a series of experiments with real datasets that SSBM performs on par with supervised learning and outperforms sparse reconstruction methods with a fixed wavelet basis by a large margin.

1 Introduction

Continuous signals have to be quantized in order to be represented digitally with a limited number of bits in a computer. In many real-world applications, such as radar (Alberti et al., 1991), wireless sensor networks (Chen & Wu, 2015), and recommender systems (Davenport et al., 2014), the measured data is quantized with just a few bits per observation. The extreme case of quantization corresponds to observing a single bit per measurement. For example, single-photon detectors record the presence or absence of photons at each measurement cycle (Kirmani et al., 2014), and recommendation systems often observe a binary measurement of users’ preferences only (*e.g.*, via thumbs up or down).

The binary sensing problem is formalized as follows: we observe binary measurements $y \in \{-1, 1\}^m$ of a signal $x \in \mathcal{X} \subset \mathbb{S}^{n-1}$ with unit norm¹ via the following forward model

$$y = \text{sign}(Ax) \tag{1}$$

where $A \in \mathbb{R}^{m \times n}$ is a linear forward operator. Recovering the signal from the measurements is an ill-posed inverse problem since there are many signals $x \in \mathbb{S}^{n-1}$ that are consistent with a given measurement vector y . Moreover, often the measurement matrix is incomplete $m < n$, *e.g.*, as in one-bit compressed sensing (Jacques et al., 2013), which makes the signal recovery problem even more challenging.

It is possible to obtain a good estimation of x despite the binary quantization, if the set of plausible signals \mathcal{X} is low-dimensional (Bourrier et al., 2014), *i.e.*, if it occupies a small portion of the ambient space \mathbb{S}^{n-1} . A popular approach is to assume that \mathcal{X} is a single linear subspace or a union of subspaces (Jacques et al., 2013), imposing sparsity over a known dictionary. For example, the well-known total variation regularization assumes that the gradients of the signal are sparse (Rudin et al., 1992). However, in real-world settings, the set of signals \mathcal{X} is generally unknown, and sparsity assumptions on an arbitrary dictionary yield a loose description of the true set \mathcal{X} , negatively impacting the quality of reconstructions obtained under this

¹Note that the sensing model in (1) provides no information about the norm of x , so it is commonly assumed that signals verify $\|x\| = 1$.

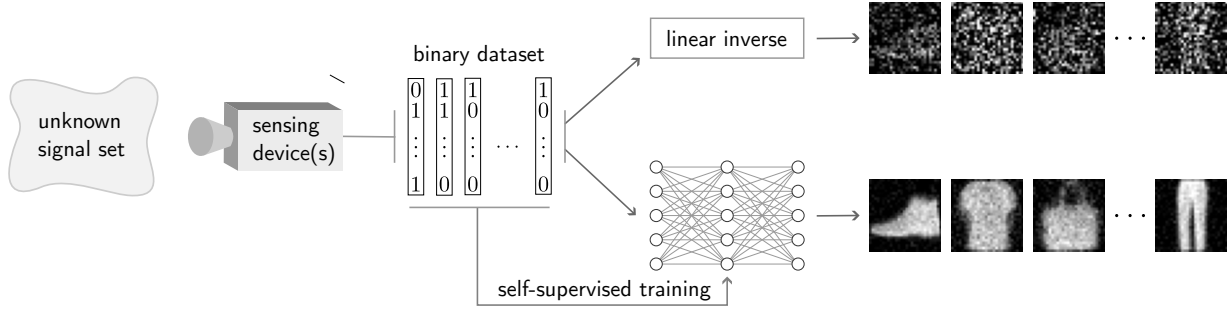


Figure 1: We propose a method for learning to reconstruct binary measurement observations, using only the binary observations themselves for training. The learned reconstruction function can discover unseen patterns in the data (in this case the clothes of fashionMNIST - see the experiments in Section 5), which cannot be recognized in the standard linear reconstructions (no learning). We also provide theoretical bounds which characterize how well we can expect to learn the set of signals from binary measurement data alone.

assumption. This limitation can be overcome by learning the reconstruction mapping $y \mapsto x$ (e.g., with a deep neural network) directly from N pairs of measurements and associated signals—i.e., a supervised learning scenario with a labeled dataset $\{(y_i, x_i)\}_{i=1}^N$ with N assumed sufficiently large. While this learning-based approach generally obtains state-of-the-art performance, it is often impractical since it can be very expensive or even impossible to obtain ground-truth signals x_i for training. For example, recommender systems generally do not have access to high-resolution user ratings on all items for training.

In this paper, we investigate the problems of identifying the signal set and learning reconstruction mapping using a dataset of binary measurements only $\{y_i\}_{i=1}^N$. In this setting, if the measurement process is incomplete $m < n$, the matrix A has a non-trivial nullspace and there is no information in the measurement data about the set of signals \mathcal{X} in the nullspace (Chen et al., 2021). As a consequence, there is not enough information for learning the reconstruction function either. For example, the trivial pseudo-inverse reconstruction $f(y) = A^\top(AA^\top)^{-1}y$ is perfectly consistent with the binary measurements, i.e., $\text{sign}(Af(y)) = y$, but is generally far from being optimal (Boufounos et al., 2015).

Here we show that it is still possible to (approximately) identify the signal set and learn to reconstruct the binary measurements, if the measurement operator varies across observations, i.e.,

$$y_i = \text{sign}(A_{g_i}x_i) \quad (2)$$

where each signal x_i is observed via one out of G operators $g_i \in \{1, \dots, G\}$, and $i = 1, \dots, N$. This sensing assumption holds in various practical applications, where signals are observed through different operators (e.g., recommendation systems access ratings about a different set of items for each user) or through an operator which changes through time (e.g., a sensor that changes its calibration). Moreover, this assumption is also valid for the case where we obtain binary measurements via a single operator A , but the set \mathcal{X} is known to be invariant to a group of invertible transformations $\{T_g\}_{g=1}^G$, such as translations or rotations. The invariance of \mathcal{X} provides access to measurements associated with a set of (implicit) operators $\{A_g = AT_g\}_{g=1}^G$, as we have that

$$y = \text{sign}(AT_gT_g^{-1}x) = \text{sign}(AT_gx') \quad (3)$$

with $x' = T_g^{-1}x \in \mathcal{X}$ for all $g = 1, \dots, G$. This observation has been exploited to perform fully unsupervised learning on various linear inverse problems, such as magnetic resonance imaging and computed tomography (Chen et al., 2021; 2022; Tachella et al., 2023).

The problem of recovering a signal from binary measurements under the assumption of a known signal set has been extensively studied in the literature (Goyal et al., 1998; Jacques et al., 2013; Oymak & Recht, 2015). These works provide practical bounds which characterize the recovery error as a function of the number of measurements m for different classes of signal sets. However, they assume that the signal set is known (or that there is enough ground-truth training data to approximate it), which is not often the case in real-world

Assumption on $\mathcal{X} \subseteq \mathbb{S}^{n-1}$	None	None	boxdim $< k$
Assumption on $A_g \in \mathbb{R}^{m \times n}$, $g \in [G]$	rank $[A_1^\top, \dots, A_G^\top] < n$	None	Gaussian
Identification error bounds	$\delta > 1$	$\delta \gtrsim \frac{n}{mG}$	$\delta \lesssim \frac{k+n/G}{m} \log \frac{nm}{k+n/G}$
Section	Section 3.1	Section 3.1	Section 3.2

Table 1: Summary of the global model identification error δ bounds presented in this paper, as a function of the size of the signals n , the number of binary measurement operators G with m measurements and the dimension of the signal set k .

scenarios. Here we investigate the best approximation of the signal set that can be obtained from the binary observations. This approximation lets us understand how well we can learn the reconstruction function from binary data. To the best of our knowledge, the model identification problem has not been yet addressed, and we aim to provide the first answers to this problem here. The main contributions of this paper are:

- We show that for any G sensing matrices $A_1, \dots, A_G \in \mathbb{R}^{m \times n}$ and any dataset size N , there exists a signal set whose identification error (precisely defined in Section 3) from binary measurements cannot decay faster than $\mathcal{O}(\frac{n}{mG})$ when m increases.
- We prove that, if each operator A_g , $g \in \{1, \dots, G\}$, has iid Gaussian entries (a standard construction in one-bit compressed sensing), it is possible to estimate a k -dimensional² signal set up to a global error of $\mathcal{O}(\frac{k+n/G}{m} \log \frac{nm}{k+n/G})$ with high probability.
- We determine the *sample complexity* of the related unsupervised learning problem, *i.e.*, we find that the number of distinct binary observations for obtaining the best possible approximation of a k -dimensional signal set \mathcal{X} is $N = \mathcal{O}(G(\frac{m}{k})^k)$.
- We introduce a Self-Supervised learning loss for training reconstruction networks from Binary Measurement data alone (SSBM), and show experimentally that the learned reconstruction function outperforms classical binary iterative hard thresholding (Jacques et al., 2013) and performs on par with fully supervised learning on various real datasets.

A summary of the model identification bounds presented in this paper is shown in Table 1.

Related Work

Unsupervised learning in inverse problems. Despite providing very competitive results, most deep learning-based solvers require a supervised learning scenario, *i.e.*, they need measurements and signal pairs $\{(y_i, x_i)\}$, a labeled dataset, in order to learn the reconstruction function $y \mapsto x$. A first step to overcome this limitation is due to Noise2Noise (Lehtinen et al., 2018), where the authors show that it is possible to learn from only noisy data if two noisy realizations of the same signal $\{(x_i + n_i, x_i + n'_i)\}$ are available for training. This approach has been extended to linear inverse problems with pairs of measurements $\{(A_{g_i}x_i + n_i, A_{g'_i}x_i + n'_i)\}$ (Yaman et al., 2020; Liu et al., 2020). The equivariant imaging framework (Chen et al., 2021; 2022) shows that learning the reconstruction function from unpaired measurement data $\{Ax_i + n_i\}$ of a single incomplete linear operator A is possible if the signal model is invariant to a group of transformations. This approach can also be adapted to the case where the signal model is not invariant, but measurements are obtained via many different operators $\{A_{g_i}x_i + n_i\}$ (Tachella et al., 2022). Necessary and sufficient conditions for learning in these settings are presented in Tachella et al. (2023), however under the assumption of linear observations (no quantization). Here we extend these results to the non-linear binary sensing problem with an unsupervised dataset with multiple operators $\{\text{sign}(A_{g_i}x_i)\}$ and $g_i \in \{1, \dots, G\}$, or with a single operator and a group-invariant signal set $\{\text{sign}(Ax_i)\}$.

²The definition of dimension used in this paper is the upper box-counting dimension defined in Section 2.




Figure 2: Geometry of the 1-bit signal recovery problem with $m = 5$ and $n = 3$. **Left:** The binary sensing operator $\text{sign}(A \cdot)$ defines a tessellation of the sphere into multiple *consistency cells*, which are defined as all vectors $x \in \mathbb{S}^2$ associated with the same binary code. The consistency cell associated with a given measurement y is shown in green. Each red line is a great circle defined by all points of \mathbb{S}^2 perpendicular to one row of A . **Middle:** If the signal set consists of all vectors in the sphere, *i.e.*, $\mathcal{X} = \mathbb{S}^2$, the center of the cell is the optimal reconstruction $\hat{f}(y)$ (depicted with a blue cross) and the recovery error (denoted by δ) is given by the radius of the cell. **Right:** If the signal set (depicted in black) occupies only a small subset of \mathbb{S}^2 , *i.e.*, it has a small box-counting dimension, the optimal reconstruction corresponds to the center of the intersection between the signal set and the consistency cell, and the resulting signal recovery error is smaller.

Quantized and one-bit sensing. Reconstructing signals from one-bit compressive measurements is a well-studied problem (Goyal et al., 1998; Jacques et al., 2013; Oymak & Recht, 2015; Baraniuk et al., 2017), both in the (over)complete case $m \geq n$ (Goyal et al., 1998), and in the incomplete setting $m < n$, either under the assumption that the signals are sparse (Jacques et al., 2013), or more generally, that the signal set has small Gaussian width (Oymak & Recht, 2015). Some of these results are summarized in Section 2. The theoretical bounds presented in this paper complement those of signal recovery bounds from quantized data, as they characterize the fundamental limitations of model identification from binary measurement data.

One-bit matrix completion and dictionary learning. Matrix completion consists of inferring missing entries of a data matrix $Y = [y_1, \dots, y_N]$, whose columns can be seen as partial observations of signals x_i , *i.e.*, $y_i = \text{sign}(A_{g_i} x_i)$ where the operators A_{g_i} select a random subset of m entries of the signal x_i . In order to recover the missing entries, it is generally assumed that the signals x_i (the columns of $X = [x_1, \dots, x_N]$) belong to a k -dimensional subspace with $k \ll n$. Davenport et al. (2014) solve this learning problem via convex programming, and present theoretical bounds for the reconstruction error.

Zayyani et al. (2015) present an algorithm that learns a dictionary (*i.e.*, a union of k -dimensional subspaces) from binary data alone in the overcomplete regime $m > n$. Rencker et al. (2019) presents a similar dictionary learning algorithm with convergence guarantees. In this paper, we characterize the model identification error for the larger class of low-dimensional signal sets, which includes subspaces and the union of subspaces as special cases. Moreover, we propose a self-supervised method that learns the reconstruction mapping directly, avoiding an explicit definition (*e.g.*, a dictionary) of the signal set.

2 Signal Recovery Preliminaries

We begin with some basic definitions related to the one-bit sensing problem. The diameter of a set is defined as $\text{diam}(S) = \sup_{u, v \in S} \|u - v\|$, and the radius is defined as half the diameter. Each row $a_i \in \mathbb{R}^n$ in the operator A divides the unit sphere \mathbb{S}^{n-1} into two hemispheres, *i.e.*, $\{x \in \mathbb{S}^{n-1} \mid a_i^\top x \geq 0\}$ and $\{x \in \mathbb{S}^{n-1} \mid a_i^\top x < 0\}$. Considering all rows, the operator $\text{sign}(A \cdot)$ defines a *tessellation* of \mathbb{S}^{n-1} into *consistency cells*, where each cell is composed of all the signals that are associated with a binary code y , *i.e.* $\{x \in \mathbb{S}^{n-1} \mid \text{sign}(Ax) = y\}$. The radius and number of consistency cells play an important role in the

analysis of signal recovery and model identification. Figure 2 illustrates the geometry of the problem for $n = 3$ and $m = 5$.

The problem of recovering a signal from one-bit compressed measurements with a known signal set has been well studied (Goyal et al., 1998; Jacques et al., 2013; Oymak & Recht, 2015; Baraniuk et al., 2017). These works characterize the maximum estimation error across all signals obtained by an optimal reconstruction function \hat{f} , *i.e.*,

$$\delta = \max_{x \in \mathcal{X}} \|x - \hat{f}(\text{sign}(Ax))\| \quad (4)$$

as a function of the number of measurements and complexity of the signal model. From a geometric viewpoint (see Figure 2), the optimal reconstruction function with respect to the norm $\|\cdot\|$ is given by the centroid (with respect to the same norm $\|\cdot\|$) of the intersection between the consistency cell associated with the measurement $y = y(x) = \text{sign}(Ax)$, *i.e.*, $S_y := \{u \in \mathbb{S}^{n-1} : y = \text{sign}(Au)\}$, and the signal set \mathcal{X} , *i.e.*,

$$\hat{f}(y) = \text{centroid}(S_y \cap \mathcal{X}). \quad (5)$$

while the maximum reconstruction error is given by the intersection with maximal radius, that is

$$\delta = \max_{x \in \mathcal{X}} \text{radius}(S_{y(x)} \cap \mathcal{X}). \quad (6)$$

In the overcomplete case $m > n$, assuming that all unit vectors are plausible signals, *i.e.*, $\mathcal{X} = \mathbb{S}^{n-1}$, the mean reconstruction error δ is given by the consistency cell with maximal radius, which scales as $\frac{n}{m}$ (see Proposition 4). The optimal rate is achieved by measurement consistent reconstruction functions, *i.e.*, those verifying $y = \text{sign}(Af(y))$ (Goyal et al., 1998).

In the incomplete case $m < n$, non-trivial signal recovery is only possible if the set of signals occupies a low-dimensional subset of the unit sphere \mathbb{S}^{n-1} (Oymak & Recht, 2015). For example, a common assumption is that \mathcal{X} is the set of k -sparse vectors (Jacques et al., 2013). In this paper, we characterize the class of low-dimensional sets using a single intuitive descriptor, the box-counting dimension. The upper box-counting dimension (Falconer, 2004, Chapter 2) is defined for a compact subset $S \subset \mathbb{R}^n$ as

$$\text{boxdim}(S) = \limsup_{\epsilon \rightarrow 0^+} \frac{\log \mathfrak{N}(S, \epsilon)}{\log 1/\epsilon} \quad (7)$$

where $\mathfrak{N}(S, \epsilon)$ is the minimum number of closed balls of radius ϵ with respect to the norm $\|\cdot\|$ that are required to cover S . This descriptor has been widely adopted in the inverse problems literature (Puy et al., 2017; Tachella et al., 2023), and it captures the complexity of various popular models, such as smooth manifolds (Baraniuk & Wakin, 2009) and union of subspaces (Blumensath & Davies, 2009; Baraniuk et al., 2017). For example, the set of $(k+1)$ -sparse vectors with unit norm has a box-counting dimension equal to k . The upper box-counting dimension is particularly useful to obtain an upper bound on the covering number of a set: if $\text{boxdim}(\mathcal{X}) < k$, there exists a set-dependent constant $\epsilon_0 \in (0, \frac{1}{2})$ for which

$$\mathfrak{N}(\mathcal{X}, \epsilon) \leq \epsilon^{-k} \quad (8)$$

holds for all $\epsilon \leq \epsilon_0$ (Puy et al., 2017). The following theorem exploits this fact to provide a bound on the number of measurements needed for recovering a signal with an error smaller than δ from generic binary observations:

Theorem 1. *Let A be a matrix with iid entries sampled from a standard Gaussian distribution and assume that $\text{boxdim}(\mathcal{X}) < k$, such that $\mathfrak{N}(\mathcal{X}, \epsilon) \leq \epsilon^{-k}$ for all $\epsilon < \epsilon_0$ with $\epsilon_0 \in (0, \frac{1}{2})$. If the number of measurements verifies*

$$m \geq \frac{2}{\delta} \left(2k \log \frac{4\sqrt{n}}{\delta} + \log \frac{1}{\xi} \right) \quad (9)$$

then for all $x, s \in \mathcal{X}$ and $\delta \leq \min\{4\sqrt{n}\epsilon_0, \frac{1}{2}\}$, we have that

$$\text{sign}(Ax) \neq \text{sign}(As) \implies \|x - s\| < \delta \quad (10)$$

with probability greater than $1 - \xi$.




Figure 3: Illustration of the model identification problem from binary measurements with $n = 3$, $m = 4$, and $G = 3$. A signal set with box-counting dimension 1 is depicted in black. The red lines define the frontiers of the consistency cells associated with operators A_1, \dots, A_3 . **From left to right:** The signal set, the estimation of the signal set associated with A_1, \dots, A_3 and the overall estimate $\hat{\mathcal{X}}$.

This result extends Theorem 2 in Jacques et al. (2013), which holds for k -sparse sets only, to general low-dimensional sets and is included in Appendix B. For example, if \mathcal{X} is the intersection of L $(s+1)$ -dimensional subspaces with the unit sphere, Theorem 1 holds with constant $\epsilon_0 = (3^s L)^{-\frac{1}{k-s}}$ and $k > s$ (Vershynin, 2018, Chapter 4.2). This theorem tells us that we can recover sparse signals from binary measurements up to an error of

$$\mathcal{O}\left(\frac{k}{m} \log \frac{nm}{k}\right)$$

which is sharp, up to the logarithmic factor (Jacques et al., 2013). Oymak and Recht (Oymak & Recht, 2015) present a similar result, stated in terms of the Gaussian width³ of the signal set instead of the box-counting dimension.

3 Model Identification from Binary Observations

In this section, we study how well we can identify the signal set from binary measurement data associated with G different measurement operators $A_1, \dots, A_G \in \mathbb{R}^{m \times n}$. We focus on the problem of identifying the set \mathcal{X} from the binary sets $\{\text{sign}(A_g \mathcal{X})\}_{g=1}^G$. In practice, we observe a subset of each binary set $\text{sign}(A_g \mathcal{X})$, however, in Section 3.4 we show that the number of elements in each of these sets is controlled by the box-counting dimension of \mathcal{X} , which is typically low in real-world settings (Hein & Audibert, 2005).

We start by analyzing how the different operators provide us with information about \mathcal{X} . Each forward operator A_g constrains the signal space by the following set

$$\hat{\mathcal{X}}_g = \{v \in \mathbb{S}^{n-1} \mid \exists x_g \in \mathcal{X}, \text{sign}(A_g v) = \text{sign}(A_g x_g)\}. \quad (11)$$

Each set $\hat{\mathcal{X}}_g$ is thus composed of all unit vectors v that are *consistent* with at least one point x_g of \mathcal{X} according to the binary mapping $\text{sign}(A_g \cdot)$. We thus conclude that $\hat{\mathcal{X}}_g$ is essentially a *dilation* of \mathcal{X} —and we clearly have $\mathcal{X} \subset \hat{\mathcal{X}}_g$ —whose extension is locally determined by specific cells of $\text{sign}(A_g \cdot)$. A three-dimensional example with $m = 4$ measurements and $G = 3$ operators is presented in Figure 3. Note that, for a given binary mapping $\text{sign}(A_g \cdot)$, each cell is characterized by one binary vector in the range of this mapping, so that, as shown in this figure, all cells provide a different tessellation of \mathbb{S}^{n-1} whose size and dimension will play an important role in our analysis.

Since each $\hat{\mathcal{X}}_g$ is a dilation of \mathcal{X} , we can infer the signal set from the following intersection

$$\hat{\mathcal{X}} := \bigcap_{g=1}^G \hat{\mathcal{X}}_g, \quad (12)$$

which can be expressed concisely as

$$\hat{\mathcal{X}} = \{v \in \mathbb{S}^{n-1} \mid \exists x_1, \dots, x_G \in \mathcal{X}, \text{sign}(A_g v) = \text{sign}(A_g x_g), \forall g = 1, \dots, G\}. \quad (13)$$

³The Gaussian width of a set S is defined as $\mathbb{E}_s \{\sup_{x \in S} x^\top s\}$ where s is distributed as a standard Gaussian vector.




Figure 4: Illustration of the oracle argument in the example of Figure 3. **Left:** The signal set $\mathcal{X} \subset \mathbb{S}^2$ is depicted in black. **Middle:** Cells intersected by the oracle system are indicated in green. **Right:** The identified set $\hat{\mathcal{X}}$ is indicated in green, and is larger than the oracle counterpart.

Due to the binary quantization, the inferred set will be larger than the true set, *i.e.*, $\mathcal{X} \subset \hat{\mathcal{X}}$. However, we will show that it is possible to learn a slightly *larger* signal set, defined in terms of a global identification error $\delta > 0$, *i.e.*, the open δ -tube

$$\mathcal{X}_\delta = \{v \in \mathbb{S}^{n-1} \mid \|x - v\| < \delta, x \in \mathcal{X}\} \quad (14)$$

such that the inferred set is contained in it, *i.e.*, $\hat{\mathcal{X}} \subset \mathcal{X}_\delta$. For our developments to be valid, we will further assume that \mathcal{X} is not too dense over \mathbb{S}^{n-1} so that two tubes of \mathcal{X} with two distinct radii are distinct.

Assumption 1. The set \mathcal{X} is closed and there exists a maximal radius $0 < \delta_0 < 2$ for which $\mathcal{X}_\delta \subsetneq \mathcal{X}_{\delta_0}$ for any $0 < \delta < \delta_0$.

This assumption amounts to saying that there exists at least one open ball in \mathbb{S}^{n-1} that does not belong to $\mathcal{X} \subset \mathbb{S}^{n-1}$. For instance, $\mathcal{X} = \mathbb{S}^{n-1}$ does not verify this assumption, and $\mathcal{X} = \mathbb{S}^{n-1} \cap \{x \in \mathbb{R}^n : x_1 \geq 0\}$ verifies it for $\delta_0 \leq \sqrt{2}$ since $\mathcal{X}_\delta = \mathbb{S}^{n-1}$ for any $\delta \geq \sqrt{2}$. The next subsections provide lower and upper bounds for δ .

3.1 A Lower Bound on the Identification Error

We first aim to find a lower bound on the best δ achievable via the following oracle argument: if we had oracle access to G measurements of each point x in \mathcal{X} through each of the G different operators, we could stack them together to obtain a larger measurement operator, defined as

$$\begin{bmatrix} y_1 \\ \vdots \\ y_G \end{bmatrix} = \text{sign}(\bar{A}x) \text{ with } \bar{A} = \begin{bmatrix} A_1 \\ \vdots \\ A_G \end{bmatrix} \in \mathbb{R}^{mG \times n}. \quad (15)$$

This oracle measurement operator provides a refined approximation of the signal set, specified as

$$\hat{\mathcal{X}}_{\text{oracle}} = \{v \in \mathbb{S}^{n-1} \mid \exists x \in \mathcal{X}, \text{sign}(\bar{A}v) = \text{sign}(\bar{A}x)\}, \quad (16)$$

which is again a dilation of \mathcal{X} .

Figure 4 shows an example with the oracle set $\hat{\mathcal{X}}_{\text{oracle}}$, which provides a better (or equal) approximation of the signal set than (13), due to the fact that $\mathcal{X} \subset \hat{\mathcal{X}}_{\text{oracle}} \subseteq \hat{\mathcal{X}}$ by the construction of these sets. As the oracle estimate is composed of the cells associated with $\text{sign}(\bar{A}\cdot)$ which are intersected by the signal set, the oracle approximation error depends on the diameter of the intersected cells. Given a certain oracle tessellation of \mathbb{S}^{n-1} , the worst estimate of \mathcal{X} is obtained when it intersects the largest cells in the tessellation. The following proposition formalizes the intuition that the maximum consistency cell diameter—*i.e.*, the greatest distance separating two binary consistent vectors of \mathcal{X} according to \bar{A} —serves as a lower bound on the model identification error δ .

Proposition 2. *Given $\bar{A} \in \mathbb{R}^{mG \times n}$, for any set $\mathcal{X} \subset \mathbb{S}^{n-1}$ respecting Assumption 1 with $0 < \delta_0 < 2$, there exists a rotation matrix $R \in SO(n)$ such that the rotated set*

$$\mathcal{X}' = \{v \in \mathbb{S}^{n-1} | v = Rx, x \in \mathcal{X}\} = R\mathcal{X} \quad (17)$$

verifies $\hat{\mathcal{X}}'_{\text{oracle}} \not\subset \mathcal{X}'_\delta$ for any $\delta < \min\{d, \delta_0\}$ where $0 < d < 2$ is the largest cell diameter of the tessellation induced by $\text{sign}(\bar{A} \cdot)$.

Proof. Given $\delta < \delta_0$, the proof consists in choosing an appropriate rotation matrix, such that we can find a point v which belongs to the oracle estimate $\hat{\mathcal{X}}'_{\text{oracle}}$ of the rotated set \mathcal{X}' , but doesn't belong to the δ -tube \mathcal{X}'_δ of this set. From Assumption 1 and since the δ -tube \mathcal{X}_δ is open, there exists $x \in \mathcal{X}$ and $v \notin \mathcal{X}_\delta$ such that $\|x - v\| = \delta$. Let S denote the largest cell in the tessellation of \mathbb{S}^{n-1} induced by $\text{sign}(\bar{A} \cdot)$, such that $d = \text{diam}(S)$. If $\delta < d$, we can always pick a rotation $R \in SO(n)$ such that both $x' = Rx$ and $v' = Rv$ belong to S . As $x' \in S$, \mathcal{X}' intersects S and we have that $S \subseteq \hat{\mathcal{X}}'_{\text{oracle}}$, and thus that $v' \in \hat{\mathcal{X}}'_{\text{oracle}}$. \square

In the rest of this subsection, we focus on bounding the maximum cell diameter, as it is directly related to the model identification error through Proposition 2. We start with the following proposition which shows that, if the stacked matrix is rank-deficient, there exist cells having the maximum possible diameter.

Proposition 3. *Consider the tessellation defined by $\text{sign}(\bar{A} \cdot)$ with $\bar{A} \in \mathbb{R}^{mG \times n}$. If*

$$\text{rank}(\bar{A}) < n \quad (18)$$

there are cells in the tessellation with a diameter equal to 2.

Proof. If \bar{A} has a rank smaller than n , it has a non-trivial nullspace. Let $v \in \mathbb{S}^{n-1}$ be a generator of the nullspace with unit norm. Consider a cell associated with the code $\text{sign}(\bar{A}x)$ for some $x \in \mathbb{R}^n$ inside the range of \bar{A}^\top . The points $\frac{x+v}{\|x+v\|}, \frac{x-v}{\|x-v\|} \in \mathbb{S}^{n-1}$ belong to this cell since they share the same code. As $\|x \pm v\| = \sqrt{\|v\|^2 + \|x\|^2}$ due to orthogonality, the distance between these two points is

$$\frac{2\|v\|}{\sqrt{\|v\|^2 + \|x\|^2}} = \frac{2}{\sqrt{1 + \|x\|^2}} \quad (19)$$

which tends to 2 as $\|x\|$ goes to zero, without modifying the cell code $\text{sign}(\bar{A}x)$. \square

This proposition tells us that n/G measurements are necessary in order to obtain non-trivial cell diameters, and thus to obtain a non-trivial estimation of \mathcal{X} . It provides a practical necessary condition for model identification, *i.e.*, that we have at least

$$m > \frac{n}{G}$$

measurements per operator. Moreover, in practice, it is possible to compute the rank of the stacked matrix \bar{A} via numerical approximations. The following theorem provides a more refined characterization of the oracle error for $m \geq n/G$:

Proposition 4. *Consider the tessellation defined by $\text{sign}(\bar{A} \cdot)$ with $\bar{A} \in \mathbb{R}^{mG \times n}$. The largest cell in the tessellation has a diameter of size at least $\frac{2}{3} \frac{n}{mG}$.*

Proof. According to Thao & Vetterli (1996, Theorem A.7), the maximum number of cells $C_{\bar{A}}$ induced by a tessellation defined by $\text{sign}(\bar{A} \cdot)$ with $\bar{A} \in \mathbb{R}^{mG \times n}$ can be upper bounded as

$$C_{\bar{A}} \leq \binom{mG}{n} 2^n.$$

As $\binom{mG}{n} \leq (\frac{emG}{n})^n$, we have that $C_{\bar{A}} \leq (\frac{2emG}{n})^n$. We can inscribe all cells into spherical caps S_i^4 of radius $\delta/2$, where δ is the maximum cell diameter. As shown in (Ball et al., 1997, Lemma 2.3), a spherical cap

⁴A spherical cap of radius r around a point $v \in \mathbb{S}^{n-1}$ is defined as $\{x \in \mathbb{S}^{n-1} | \|x - v\| < r\}$.

of radius $\delta/2$ has measure bounded by $\sigma(S_i) \leq (\frac{\delta}{4})^{n-1} \sigma_{n-1}$ where σ_{n-1} is the measure of \mathbb{S}^{n-1} . Since the tessellation covers the unit sphere \mathbb{S}^{n-1} , we have that $\mathbb{S}^{n-1} \subseteq \cup_{i=1}^{C_{\bar{A}}} S_i$ and thus

$$\sum_{i=1}^{C_{\bar{A}}} \sigma(S_i) \geq \sigma_{n-1} \Rightarrow \left(\frac{2emG}{n}\right)^n \left(\frac{\delta}{4}\right)^n \sigma_{n-1} \geq \sigma_{n-1} \Rightarrow \delta \geq \frac{2}{3} \frac{n}{mG}.$$

□

As stated in the following corollary, Proposition 4 shows that the model identification error cannot decrease faster with the number of measurements and operators than $\mathcal{O}(\frac{n}{mG})$, since the largest cell in any oracle tessellation has a diameter of at least $\frac{2}{3} \frac{n}{mG}$.

Corollary 5. *Given the G operators $A_1, \dots, A_G \in \mathbb{R}^{m \times n}$, and any set $\mathcal{X} \subset \mathbb{S}^{n-1}$ respecting Assumption 1 with $0 < \delta_0 < 2$, for any $0 < \delta < \min(\delta_0, \frac{2}{3} \frac{n}{mG})$, there exists a rotation R such that the inferred signal set $\hat{\mathcal{X}}'$ of $\mathcal{X}' = R\mathcal{X}$ is not included in \mathcal{X}'_δ , i.e., $\hat{\mathcal{X}}' \not\subseteq \mathcal{X}'_\delta$.*

Proof. If $\mathcal{X} \subset \mathbb{S}^{n-1}$ respects Assumption 1 with $0 < \delta_0 < 2$, and $\hat{\mathcal{X}}_{\text{oracle}}$ and $\hat{\mathcal{X}}$ are the oracle set associated with \bar{A} and the inferred set of \mathcal{X}' , respectively, then, as derived previously, we know that $\mathcal{X} \subset \hat{\mathcal{X}}_{\text{oracle}} \subset \hat{\mathcal{X}}$. According to Proposition 2 and Proposition 4, there exists a rotation R such that $\hat{\mathcal{X}}'_{\text{oracle}} = R\hat{\mathcal{X}}_{\text{oracle}} \not\subseteq \mathcal{X}'_\delta = R\mathcal{X}_\delta$ for $0 < \delta < \min(\delta_0, \frac{2}{3} \frac{n}{mG})$. Therefore, from the inclusion above, we thus see that there exists $\hat{\mathcal{X}}' = R\hat{\mathcal{X}} \not\subseteq \mathcal{X}'_\delta$. □

3.2 A Sufficient Condition for Model Identification

We now seek a sufficient condition on the number of measurements per operator that guarantees the identification of \mathcal{X} up to a global error of δ . As with the sufficient conditions ensuring signal recovery (see Section 2), we assume that \mathcal{X} is low-dimensional to provide a bound that holds with high probability if the entries of the operators are sampled from a Gaussian distribution.

Theorem 6. *Given the operators $A_1, \dots, A_G \in \mathbb{R}^{m \times n}$ with entries i.i.d. as a standard Gaussian distribution, a low-dimensional signal set \mathcal{X} , with $\text{boxdim}(\mathcal{X}) < k$, such that $\mathfrak{N}(\mathcal{X}, \epsilon) \leq \epsilon^{-k}$ for all $\epsilon < \epsilon_0$ with $\epsilon_0 \in (0, \frac{1}{2})$, and some failure probability $0 < \xi < 1$, if the number of measurements per operator verifies*

$$m \geq \frac{4}{\delta} \left[\left(k + \frac{n}{G}\right) \log \frac{5\sqrt{n}}{\delta} + \frac{1}{G} \log \frac{1}{\xi} + \frac{n}{G} \log 3 \right] \quad (20)$$

for $0 < \delta \leq \min\{4\sqrt{n}\epsilon_0, \frac{1}{2}\}$, then with probability at least $1 - \xi$, we have that $\hat{\mathcal{X}} \subseteq \mathcal{X}_\delta$.

The proof is included in Appendix C. Theorem 6 provides a bound on δ , i.e., how precisely we can characterise the signal set \mathcal{X} , which we can compare with the lower bound in Proposition 4. From (20) we have that (see Appendix C for a detailed derivation),

$$\delta = \mathcal{O}\left(\frac{k+n/G}{m} \log \frac{nm}{k+n/G}\right) \quad (21)$$

The bound in (21) is consistent with existing model identification bounds in the linear setting (Tachella et al., 2023), which require $m > k + n/G$ measurements per operator for uniquely identifying the signal set.

3.3 Learning to Reconstruct

The best reconstruction function \hat{f} that can be learned from binary measurements alone can be defined as a function of the identified set $\hat{\mathcal{X}}$, as defined in (5)

$$\hat{f}(y) = \text{centroid}(S_y \cap \hat{\mathcal{X}}) \quad (22)$$

for a binary input y with associated consistency cell $S_y = \{v \in \mathbb{S}^{n-1} \mid \text{sign}(Av) = y\}$ and the error is given by the radius of $S_y \cap \hat{\mathcal{X}}$. Unfortunately, Theorems 1 and 6 do not automatically translate into a bound on the optimal reconstruction error of the reconstruction function defined in (22). Theorem 6 implies that the optimal unsupervised reconstruction $\hat{f}(\text{sign}(Ax))$ is at most $\mathcal{O}(\frac{k+n/G}{m} \log \frac{nm}{k+n/G})$ away from the signal set \mathcal{X} , but does not guarantee that it is close to x . Nonetheless, we conjecture that this rate holds with high probability if the operators follow a Gaussian distribution:

Conjecture 7. *Given binary measurements from the operators $A_1, \dots, A_G \in \mathbb{R}^{m \times n}$ with entries i.i.d. from a standard Gaussian distribution, the optimal reconstruction function defined in (22) has a maximal reconstruction error that is upper bounded as $\mathcal{O}(\frac{k+n/G}{m} \log \frac{nm}{k+n/G})$ with high probability.*

Conjecture 7 hypothesizes that the optimal unsupervised reconstruction function should obtain a similar performance than the supervised one, i.e., $\mathcal{O}(\frac{k}{m} \log \frac{nm}{k})$ shown in Theorem 1, if the number of operators is sufficiently large, i.e., $G > n/k$. In the experiments in Section 5, we provide empirical evidence that supports this hypothesis.

3.4 Sample complexity

We end our theoretical analysis of the unsupervised learning problem by bounding its *sample complexity*, i.e., we bound the number N of *distinct* binary measurement vectors $\{y_i\}_{i=1}^N$ that must be acquired for obtaining the best approximation of the signal set \mathcal{X} from binary data.

Since we observe binary vectors $y \in \{\pm 1\}^m$, there is a limited number of different binary observations. We could naively expect to observe up to 2^m different vectors per measurement operator (i.e., all possible binary codes with m bits), requiring at most $N \leq G2^m$ samples to fully characterize the best approximation of the signal set \mathcal{X} defined in (13). Fortunately, as already exploited in the proof of Proposition 4, this upper bound can be significantly reduced if the signal set has a low box-counting dimension, as not all cells in the tessellation will be intersected by the signal set (see fig. 3). We can thus obtain a better upper bound by counting the number of intersected cells, denoted as $|\text{sign}(A\mathcal{X})|$.

If \mathcal{X} is the intersection of a single k -dimensional subspace with the unit sphere, (Thao & Vetterli, 1996, Theorem A.7) tells us that, for any matrix $A \in \mathbb{R}^{m \times n}$, there are $|\text{sign}(A\mathcal{X})| \leq 2^k \binom{m}{k}$ intersected cells. More generally, if \mathcal{X} is a union of L subspaces, we have $|\text{sign}(A\mathcal{X})| \leq L2^k \binom{m}{k}$. Thus, using the fact that $\binom{m}{k} \leq (\frac{3m}{k})^k$, from the G measurement operators, we can observe up to

$$N \leq GL(\frac{6m}{k})^k \quad (23)$$

different measurement vectors. However, this result only holds for a union of subspaces having each dimension k . The following theorem extends this result to more general low-dimensional sets with small upper box-counting dimension.

Theorem 8. *Let the entries of $A \in \mathbb{R}^{m \times n}$ be sampled from a standard Gaussian distribution, and let $\mathcal{X} \subseteq \mathbb{R}^n$ with $\text{boxdim}(\mathcal{X}) < k$ such that $\mathfrak{N}(\mathcal{X}, \epsilon) \leq \epsilon^{-k}$ for all $\epsilon < \min(\frac{32}{3} \frac{k}{m\sqrt{n}} \log(\frac{3m\sqrt{n}}{32k}), \epsilon_0)$. The cardinality of the measurement set is bounded as*

$$|\text{sign}(A\mathcal{X})| \leq \left(\frac{m\sqrt{n}}{k}\right)^{8k}$$

with probability at least $1 - \frac{1024}{9m^2n}$.

The proof is included in Appendix D. This result depends on the ambient dimension n due to the application of the technical Lemma 9 and can be suboptimal for some signal sets. For example, the bound in (23) avoid this dependency for the case where \mathcal{X} is a union-of-subspaces. Moreover, the squared dependency of the probability in m can be increased by loosening the bound on N (see the proof for more details).

Using this result, if we arbitrarily set m_0 such that $\frac{1024}{9m_0^2n} = 0.01$, we thus obtain the following upper bound on the number of different binary measurement vectors:

$$N \leq G \left(\frac{m\sqrt{n}}{k}\right)^{8k}.$$

which holds with probability exceeding 0.99 for any $m \geq m_0$. Similarly to (23), this bound scales exponentially only in the model dimension k but not in the number of measurements m or operators G . In the setting of a single operator and a k -dimensional invariant signal set, we have the upper bound $N \leq \left(\frac{m\sqrt{n}}{k}\right)^{8k}$.

4 Learning Algorithms

In this section, we present a novel algorithm for learning the reconstruction function $f : (y, A) \mapsto x$ from N binary measurement vectors $\{(y_i, A_{g_i})\}_{i=1}^N$, which is motivated by the analysis in Section 3. We parameterize the reconstruction function using a deep neural network with parameters $\theta \in \mathbb{R}^p$. The learned function can take into account the knowledge about the forward operator by simply applying a linear inverse at the first layer, *i.e.*, $f_\theta(y, A) = \hat{f}_\theta(A^\top y)$, or using more complex unrolled optimization architectures (Monga et al., 2021).

In the case where we observe measurements associated with G different forward operators, we propose the SSBM loss

$$\arg \min_{\theta \in \mathbb{R}^p} \sum_{i=1}^N \left[\mathcal{L}_{\text{MC}}(y_i, A_{g_i} \hat{x}_{\theta,i}) + \alpha \sum_{s \neq g_i} \|\hat{x}_{\theta,i} - f_\theta(\text{sign}(A_s \hat{x}_{\theta,i}), A_s)\|_2^2 \right], \quad (24)$$

where $\hat{x}_{\theta,i} = f_\theta(y_i, A_{g_i})$, the cost $\mathcal{L}_{\text{MC}}(y_i, A_{g_i} \hat{x}_{\theta,i}) \geq 0$ enforces *measurement consistency* (MC), *i.e.*, require that $y_i = \text{sign}(A_{g_i} \hat{x}_{\theta,i})$, and $\alpha \in \mathbb{R}_+$ is a hyperparameter controlling the trade-off between the two terms involved. In the setting where we have a single operator, we aim to learn a reconstruction function $f_\theta : y \mapsto x$ (we remove the dependence of f_θ on A to simplify the notation) via the following self-supervised loss:

$$\arg \min_{\theta \in \mathbb{R}^p} \sum_{i=1}^N \left[\mathcal{L}_{\text{MC}}(y_i, A \hat{x}_{\theta,i}) + \alpha \sum_{g=1}^G \|T_g \hat{x}_{\theta,i} - f_\theta(\text{sign}(AT_g \hat{x}_{\theta,i}))\|_2^2 \right], \quad (25)$$

where $\hat{x}_{\theta,i} = f_\theta(y_i)$ and $\alpha \in \mathbb{R}_+$. In practice, we minimize (24) by mini-batching approaches (*e.g.*, stochastic gradient descent) by using sampling one out of the G operators at random per batch. In both cases, we choose the measurement consistency term to be the logistic loss, *i.e.*,

$$\mathcal{L}_{\text{MC}}(y, \hat{y}) = \log(1 + \exp(-y \circ \hat{y})) \quad (26)$$

which enforces sign-consistent predictions which are far from zero, as the logistic function tends asymptotically towards zero as $|\hat{y}| \rightarrow \infty$. An empirical analysis in Section 5 shows that the logistic loss obtains the best performance across various popular consistency losses.

Analysis of the proposed loss We focus on the multi-operator loss in (24), although a similar analysis also holds for the equivariant setting. As the first term of the loss enforces measurement consistency, *i.e.*, requires $y_i = \text{sign}(A_{g_i} f_\theta(y_i, A_{g_i}))$ for every y_i in the dataset. However, in the incomplete setting $m < n$, the simple pseudo-inverse solution

$$f(y, A_g) = A_g^\dagger y \quad (27)$$

with $A_g^\dagger = A_g^\top (A_g A_g^\top)^{-1}$, is measurement consistent for any number of operators G and training data N . Therefore, the first loss does not prevent learning a function $f_\theta(y, A_g)$ which acts independently for each operator (as if there were G independent learning problems). The second loss *bootstraps* the current estimates $\hat{x}_{i,\theta} = f_\theta(y_i, A_{g_i})$ as new ground truth references, mimicking the supervised loss

$$\sum_{i=1}^N \sum_{s=1}^G \|\hat{x}_{i,\theta} - f_\theta(\text{sign}(A_s \hat{x}_{i,\theta}), A_s)\|_2^2, \quad (28)$$

in order to enforce consistency across operators. Importantly, this additional loss avoids the trivial pseudo-inverse solution in (27), as

$$A_g^\dagger y - A_s^\dagger \text{sign}(A_s A_g^\dagger y) \neq 0 \quad (29)$$

for $g \neq s$ if the nullspaces of A_g and A_s are different, *e.g.*, if the necessary condition in Proposition 3 is verified.




Figure 5: Evaluated training losses for enforcing sign measurement consistency $\text{sign}(A\hat{x}) = y$ of reconstructions $f_\theta(y) = \hat{x}$. **Left:** The loss functions are shown for the case $y = 1$. **Right:** Average test PSNR of different measurement consistency losses on the MNIST dataset with $G = 10$ operators.

5 Experiments

For all experiments, we use measurement operators with entries sampled from a standard Gaussian distribution and evaluate the performance of the algorithms using by computing the average peak-to-signal ratio (PSNR) on a test set with N' ground-truth signals, that is:

$$\frac{1}{N'} \sum_{i=1}^{N'} \text{PSNR}(x'_i, f_\theta(\text{sign}(A_{g_i} x'_i), A_{g_i})), \quad (30)$$

where the PSNR is computed after normalizing the reconstructed image such that it has the same norm as the reference image, *i.e.*,

$$\text{PSNR}(x, \hat{x}) = -20 \log \|x - \hat{x} \frac{\|x\|}{\|\hat{x}\|}\|. \quad (31)$$

We choose $f_\theta(y, A) = \tilde{f}_\theta(A^\top y)$ where \tilde{f}_θ is the U-Net network introduced in (Chen et al., 2021) with weights θ , and train for 400 epochs with the Adam optimizer with learning rate 10^{-4} and standard hyperparameters $\beta_1 = 0.9$ and $\beta_2 = 0.99$.

5.1 MNIST experiments

We evaluate the theoretical bounds using the MNIST dataset, which consists of greyscale images with $n = 784$ pixels and whose box-counting dimension is approximately $k \approx 12$ (Hein & Audibert, 2005). We use 6×10^4 images for training and 10^3 for testing.

Multiple operators setting. We start by comparing the logistic consistency loss in (26) with the following alternatives:

- Standard ℓ_p -loss, $\mathcal{L}_{\text{MC}}(y, \hat{y}) = \|y - \hat{y}\|_p^p$. As this loss is zero only if $\hat{y} = y$, it promotes sign consistency, $\text{sign}(\hat{y}) = y$ and unit outputs $|\hat{y}| = 1$.
- One-sided ℓ_p -loss, $\mathcal{L}_{\text{MC}}(y, \hat{y}) = \|\max(-y \circ \hat{y}, 0)\|_p^p$ where \circ denotes element-wise multiplication and the max operation is performed element-wise. This loss is zero as long as $\text{sign}(\hat{y}) = y$ regardless of the value of $|\hat{y}|$.

For each loss we chose the best performance across trade-off parameter $\alpha \in \{0.1, 1, 10\}$. Figure 5 shows the different losses and the test performance for different values of measurements using $G = 10$ operators. The logistic loss obtains the best performance across all sampling regimes, whereas the one-sided ℓ_2 loss obtains the worst results.

Secondly, we compare the logistic loss with the following learning schemes:

- Linear inverse (no learning), defined as $\hat{x}_i = A_{g_i}^\top y_i$. This reconstruction can fail to be measurement consistent (Goyal et al., 1998).




Figure 6: **Left:** Average test PSNR of different supervised and unsupervised algorithms on the MNIST dataset with $G = 10$ operators. **Right:** The performance of the SSBM method follows closely the bounds in Conjecture 7.

- Standard supervised learning loss, defined as $\sum_{i=1}^N \|x_i - f_\theta(y_i, A_{g_i})\|^2$. We also evaluate this loss together with the cross-operator consistency term in (24) which we denote as supervised+.
- Measurement consistency loss, defined as $\sum_{i=1}^N \mathcal{L}_{MC}(y_i, A_{g_i} f_\theta(y_i, A_{g_i}))$ using the logistic loss.
- The binary iterative hard-thresholding (BIHT) reconstruction algorithm (Jacques et al., 2013) with a Daubechies4 orthonormal wavelet basis. The step size and sparsity level of the algorithm were chosen via grid search. It is worth noting that the best-performing sparsity level increases as the number of measurements m is increased.
- Proposed SSBM loss in (24) using the logistic loss for measurement consistency.

Test PSNR values obtained for the case of $G = 10$ operators are shown in the left subfigure of Figure 6, where the PSNR in dB is plotted against m/n in log-scale representation. The measurement consistency approach obtains performance similar to simply applying a linear inverse for the incomplete $m/n < 1$ setting, whereas it obtains a significant improvement over the linear inverse in the overcomplete case $m/n \geq 1$. This gap can be attributed to the lack of measurement consistency of the linear reconstruction algorithm (Goyal et al., 1998). The proposed loss obtains a performance that is several dBs above the linear inverse and BIHT for all sampling regimes. BIHT relies on the wavelet sparsity prior, which does not capture well enough the MNIST digits. SSBM performs similarly to supervised learning as the sampling ratio tends to 1, and perhaps surprisingly, it obtains slightly better performance than supervised learning for $m/n = 1.28$. However, adding the cross-operator consistency loss to the supervised method (*i.e.*, the method supervised+ in Section 5.1) performs better for all sampling regimes than SSBM.

The right plot in Figure 6 compares the performance of the SSBM with the bounds in Conjecture 7. These bounds behave almost linearly in this log-log plot of both the error—through the PSNR—and the log-scale representation of m/n . We thus observe a good agreement between the predictions in Conjecture 7 and the performance in practice.

Figure 7 shows the average test PSNR and reconstructed images obtained by the proposed self-supervised method for different values of G and m . The method fails to obtain good reconstructions when $G = 1$, as the necessary condition in Proposition 3 is not fulfilled.

Equivariant setting using shifts. We evaluate the setting of learning with a single operator by using the unsupervised equivariant objective in (25) with 2D shifts as the group of transformations (as the MNIST dataset is approximately shift-invariant). Figure 8 shows the average test PSNR and reconstructed images as a function of the measurements m for various algorithms. The proposed unsupervised method significantly outperforms the linear inverse, BIHT, and the measurement consistent network in all sampling regimes, and performs closely to supervised learning for $m/n > 0.4$.




Figure 7: (a) Average test PSNR and (b) reconstructed test images of the proposed unsupervised method for different numbers of operators G and measurements m .

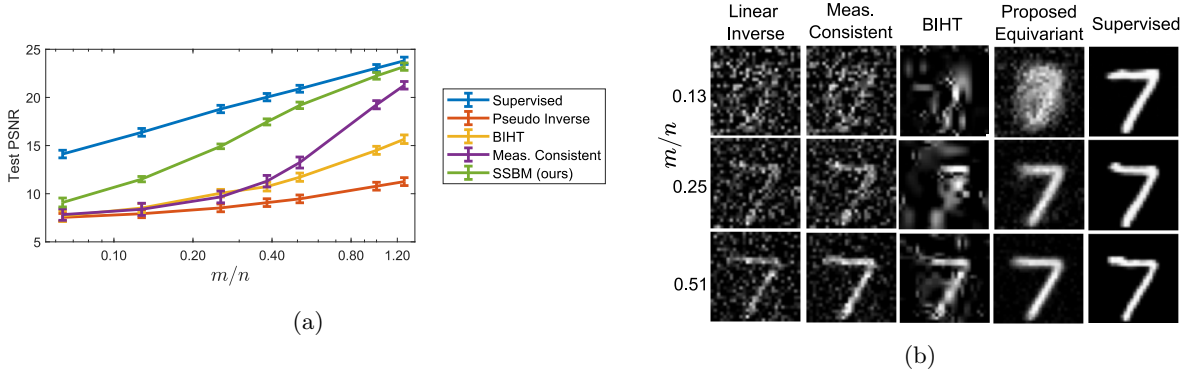


Figure 8: (a) Average test PSNR and (b) reconstructed test images by the compared algorithms with a single operator A as a function of the undersampling ratio m/n .

5.2 Other Datasets

In order to demonstrate the robustness of the proposed method across datasets, we evaluate the proposed unsupervised approach on the FashionMNIST (Xiao et al., 2017), CelebA (Liu et al., 2015) and Flowers (Nilsback & Zisserman, 2008) datasets. The FashionMNIST dataset consists of $6 \cdot 10^5$ grayscale images with 28×28 pixels which are divided across $G = 10$ different forward operators. As with MNIST, we use $N = 6 \cdot 10^4$ per operator for training and 10^4 per operator for testing. For the CelebA dataset, we use $G = 10$ forward operators and choose a subset of 10^4 images for each operator for training and another subset of the same amount for testing. The Flowers dataset consists of 6149 color images for training and 1020 images for testing, all associated with the same forward operator. For both CelebA and Flowers datasets, a center crop of 128×128 pixels of each color image was used for training and testing. Section 5.2 shows the average test PSNR of the proposed unsupervised method, standard supervised learning, BIHT, and the linear inverse. For BIHT, we use the Daubechies4 orthonormal wavelet basis and optimize the step size and sparsity level via grid search.

The self-supervised method obtains an average test PSNR which is only 1 to 2 dB below the supervised approach. Figures 9 and 10 show reconstructed test images by the evaluated approaches for each forward operator. The proposed unsupervised method is able to provide good estimates of the images, while only having access to highly incomplete binary information. The supervised method obtains sharper images, however at the cost of hallucinating details, whereas the proposed method obtains blurrier estimates with less hallucinated details.

Dataset	n	m	G	Linear Inverse	BIHT	Supervised	SSBM(ours)
FashionMNIST	784	300	10	6.38 ± 0.23	10.68 ± 0.31	17.63 ± 0.33	16.47 ± 0.22
CelebA	49152	9830	10	4.81 ± 0.32	16.26 ± 0.40	21.59 ± 0.31	19.53 ± 0.3
Flowers	49152	9830	shifts	5.31 ± 0.72	14.62 ± 0.92	18.26 ± 0.75	16.45 ± 0.71

Table 2: Average test PSNR in dB obtained by the compared methods for the FashionMNIST, CelebA and Flowers datasets.




Figure 9: Reconstructed test images using the FashionMNIST dataset. Each column corresponds to a different forward operator.

6 Conclusions and Future Work

The theoretical analysis in this work characterizes the best approximation of a low-dimensional set that can be obtained from binary measurements. The model identification bounds presented here apply to a large class of signal models, as they only rely on the box-counting dimension, and complement those existing for signal recovery from binary measurements (Goyal et al., 1998; Jacques et al., 2013). Moreover, the proposed self-supervised loss provides a practical algorithm for learning to reconstruct signals from binary measurements alone, which performs closely to fully supervised learning. This work paves the way for deploying machine learning algorithms in scientific and medical imaging applications with quantized observations, where no ground-truth references are available for training.

We leave the proof of Conjecture 7, and a study of the effect of noise in the observations and related dithering techniques for future work. Another avenue of future research is the extension of Theorem 6 for the case of operators related through the action of a group.

Acknowledgments

Part of this research was supported by the Fonds de la Recherche Scientifique – FNRS under Grant T.0136.20 (Project Learn2Sense).

References

- G Alberti, G Schirinzi, G Franceschetti, and V Pascazio. Time-domain convolution of one-bit coded radar signals. In *IEEE Proceedings F (Radar and Signal Processing)*, volume 138, pp. 438–444. IET, 1991.
- Keith Ball et al. An elementary introduction to modern convex geometry. *Flavors of geometry*, 31(1-58):26, 1997.




Figure 10: CelebA results. Reconstructed test images using the CelebA dataset. Each column corresponds to a different forward operator.

R Baraniuk, S Foucart, D Needell, Y Plan, and M Wootters. One-bit compressive sensing of dictionary-sparse signals. *Information and Inference: A Journal of the IMA*, 7(1):83–104, 08 2017. ISSN 2049-8764. doi: 10.1093/imaiai/iax009. URL <https://doi.org/10.1093/imaiai/iax009>.

Richard G Baraniuk and Michael B Wakin. Random projections of smooth manifolds. *Foundations of computational mathematics*, 9(1):51–77, 2009.

T. Blumensath and M. E. Davies. Sampling theorems for signals from the union of finite-dimensional linear subspaces. *IEEE Transactions on Information Theory*, 55(4):1872–1882, 2009. doi: 10.1109/TIT.2009.2013003.

Petros T Boufounos, Laurent Jacques, Felix Krahmer, and Rayan Saab. Quantization and compressive sensing. In *Compressed Sensing and its Applications: MATHEON Workshop 2013*, pp. 193–237. Springer, 2015.

Anthony Bourrier, Mike Davies, Tomer Peleg, Patrick Pérez, and Rémi Gribonval. Fundamental performance limits for ideal decoders in high-dimensional linear inverse problems. *IEEE Transactions on Information Theory*, 60(12):7928–7946, 2014.

Ching-Hsien Chen and Jwo-Yuh Wu. Amplitude-aided 1-bit compressive sensing over noisy wireless sensor networks. *IEEE Wireless Communications Letters*, 4(5):473–476, 2015.

Dongdong Chen, Julián Tachella, and Mike Davies. Equivariant imaging: Learning beyond the range space. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 4379–4388, October 2021.

Dongdong Chen, Julián Tachella, and Mike Davies. Robust equivariant imaging: a fully unsupervised framework for learning to image from noisy and partial measurements. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.

Mark A Davenport, Yaniv Plan, Ewout Van Den Berg, and Mary Wootters. 1-bit matrix completion. *Information and Inference: A Journal of the IMA*, 3(3):189–223, 2014.

Kenneth Falconer. *Fractal geometry: mathematical foundations and applications*. John Wiley & Sons, 2004.

V.K. Goyal, M. Vetterli, and N.T. Thao. Quantized overcomplete expansions in \mathbb{R}^N : analysis, synthesis, and algorithms. *IEEE Transactions on Information Theory*, 44(1):16–31, 1998. doi: 10.1109/18.650985.

- Matthias Hein and Jean-Yves Audibert. Intrinsic dimensionality estimation of submanifolds in \mathbb{R}^d . In *Proceedings of the 22nd international conference on Machine learning (ICML)*, pp. 289–296, 2005.
- Laurent Jacques, Jason N. Laska, Petros T. Boufounos, and Richard G. Baraniuk. Robust 1-bit compressive sensing via binary stable embeddings of sparse vectors. *IEEE transactions on information theory*, 59(4): 2082–2102, 2013.
- Ahmed Kirmani, Dheera Venkatraman, Donggeek Shin, Andrea Colaço, Franco NC Wong, Jeffrey H Shapiro, and Vivek K Goyal. First-photon imaging. *Science*, 343(6166):58–61, 2014.
- Jaakko Lehtinen, Jacob Munkberg, Jon Hasselgren, Samuli Laine, Tero Karras, Miika Aittala, Timo Aila, et al. Noise2Noise. In *International Conference on Machine Learning (ICML)*. PMLR, 2018.
- Jiaming Liu, Yu Sun, Cihat Eldeniz, Weijie Gan, Hongyu An, and Ulugbek S Kamilov. RARE: Image reconstruction using deep priors learned without groundtruth. *IEEE Journal of Selected Topics in Signal Processing*, 14(6):1088–1099, 2020.
- Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.
- Vishal Monga, Yuelong Li, and Yonina C Eldar. Algorithm unrolling: Interpretable, efficient deep learning for signal and image processing. *IEEE Signal Processing Magazine*, 38(2):18–44, 2021.
- Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *2008 Sixth Indian Conference on Computer Vision, Graphics & Image Processing*, pp. 722–729. IEEE, 2008.
- Samet Oymak and Ben Recht. Near-optimal bounds for binary embeddings of arbitrary sets. *arXiv preprint arXiv:1512.04433*, 2015.
- Gilles Puy, Mike E Davies, and Rémi Gribonval. Recipes for stable linear embeddings from hilbert spaces to \mathbb{R}^m . *IEEE Transactions on Information Theory*, 63(4):2171–2187, 2017.
- Lucas Rencker, Francis Bach, Wenwu Wang, and Mark D Plumbley. Sparse recovery and dictionary learning from nonlinear compressive measurements. *IEEE Transactions on Signal Processing*, 67(21):5659–5670, 2019.
- Leonid I Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: nonlinear phenomena*, 60(1-4):259–268, 1992.
- Julián Tachella, Dongdong Chen, and Mike Davies. Unsupervised learning from incomplete measurements for inverse problems. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), *Advances in Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=aV9WSvM6N3>.
- Julián Tachella, Dongdong Chen, and Mike Davies. Sensing theorems for learning from incomplete measurements. *Journal of Machine Learning Research*, 24(39):1–45, 2023.
- Nguyen T Thao and Martin Vetterli. Lower bound on the mean-squared error in oversampled quantization of periodic signals using vector quantization analysis. *IEEE Transactions on Information Theory*, 42(2): 469–479, 1996.
- Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.
- Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms, 2017.
- Burhaneddin Yaman, Seyed Amir Hossein Hosseini, Steen Moeller, Jutta Ellermann, Kâmil Uğurbil, and Mehmet Akçakaya. Self-supervised learning of physics-guided reconstruction neural networks without fully sampled reference data. *Magnetic resonance in medicine*, 84(6):3172–3191, 2020.

Hadi Zayyani, Mehdi Korki, and Farrokh Marvasti. Dictionary learning for blind one bit compressed sensing. *IEEE Signal Processing Letters*, 23(2):187–191, 2015.

A Technical Lemmas

We begin by introducing some technical results that play an important role in the main theorems of the paper. We start with a result from (Jacques et al., 2013).

Lemma 9 (Lemma 9 in (Jacques et al., 2013)). *Given $0 \leq \epsilon < 1$ and two unit vectors $\tilde{x}, \tilde{v} \in \mathbb{S}^{n-1} \subset \mathbb{R}^n$ and $a \in \mathbb{R}^n$ with $a_i \sim_{\text{i.i.d.}} \mathcal{N}(0, 1)$, we have*

$$p_0 = \mathbb{P} [\forall x \in B_\epsilon(\tilde{x}), \forall v \in B_\epsilon(\tilde{v}) \mid \text{sign}(a^\top v) = \text{sign}(a^\top x)] \geq 1 - d(\tilde{x}, \tilde{v}) - \sqrt{n \frac{\pi}{2}} \epsilon \quad (32)$$

$$p_1 = \mathbb{P} [\forall x \in B_\epsilon(\tilde{x}), \forall v \in B_\epsilon(\tilde{v}) \mid \text{sign}(a^\top v) \neq \text{sign}(a^\top x)] \geq d(\tilde{x}, \tilde{v}) - \sqrt{n \frac{\pi}{2}} \epsilon. \quad (33)$$

where $d(\cdot, \cdot)$ denotes the angular distance.

Let $C^0(S)$ denote the set of continuous functions on the set S . This lemma has the following corollary:

Corollary 10. *Given $\tilde{x} \in \mathbb{S}^{n-1}$, $0 < \epsilon < 1/2$, $a \in \mathbb{R}^n$ with $a \sim_{\text{i.i.d.}} \mathcal{N}(0, 1)$, we have*

$$\mathbb{P} [\text{sign}(a^\top \cdot) \notin C^0(B_\epsilon(\tilde{x}) \cap \mathbb{S}^{n-1})] \leq \eta(\epsilon) := \sqrt{n} \epsilon.$$

Proof. The proof can be derived from the complement of the event associated with p_0 in (32) when $\tilde{x} = \tilde{v}$. Here is, however, a simplified proof for completeness. We first observe that $\text{sign}(a^\top \cdot)$ is discontinuous over $B_\epsilon(\tilde{x}) \cap \mathbb{S}^{n-1}$ iff $|\frac{a^\top \tilde{x}}{\|a\|}| \leq \epsilon$. Therefore, by the rotational invariance of the Gaussian distribution we can choose $\tilde{x} = [1, 0, \dots, 0]^\top$ and the probability above amounts to computing

$$p := \mathbb{P} [|\frac{a_1}{\|a\|}| \leq \epsilon] = \mathbb{P} [a_1^2 \leq \epsilon^2 \|a\|^2] = \mathbb{P} [a_1^2 \leq \frac{\epsilon^2}{(1-\epsilon^2)} (a_2^2 + \dots + a_n^2)] = \mathbb{E}_\xi \mathbb{P} [a_1^2 \leq \frac{\epsilon^2}{(1-\epsilon^2)} \xi],$$

where $\xi \sim \chi^2(n-1)$. Since $\mathbb{P} [a_1^2 \leq \frac{\epsilon^2}{(1-\epsilon^2)} \xi] \leq \frac{\sqrt{2}}{\sqrt{\pi}} \frac{\epsilon}{\sqrt{1-\epsilon^2}} \sqrt{\xi}$, and $\mathbb{E}_\xi \sqrt{\xi} \leq \sqrt{\mathbb{E}_\xi \xi} \leq \sqrt{n-1} \leq \sqrt{n}$ by Jensen, we finally get $p \leq \frac{\sqrt{2}}{\sqrt{\pi}} \frac{\epsilon}{\sqrt{1-\epsilon^2}} \sqrt{n} \leq \frac{2\sqrt{2}}{\sqrt{\pi}\sqrt{3}} \epsilon \sqrt{n} < \epsilon \sqrt{n}$. \square

B Signal Recovery Proof

Proof of Theorem 1. As $\text{boxdim}(\mathcal{X}) < k$, there exist a constant $\epsilon_0 \in (0, \frac{1}{2})$ such that $\mathfrak{N}(\mathcal{X}, \epsilon) \leq \epsilon^{-k}$ for all $\epsilon \leq \epsilon_0$. Thus, there is a set Q_ϵ of ϵ^{-k} points, such that for every $x \in \mathcal{X}$, there exist a point $q \in Q_\epsilon$ which verifies $\|x - q\| < \epsilon$. Applying Lemma 9, we have for any two distinct points $\tilde{q}_1, \tilde{q}_2 \in Q_\epsilon$ and $a_i \in \mathbb{R}^n$ drawn from a standard Gaussian distribution such that

$$\mathbb{P} [\forall q_1 \in B_\epsilon(\tilde{q}_1), \forall q_2 \in B_\epsilon(\tilde{q}_2) : \text{sign}(a_i^\top q_1) \neq \text{sign}(a_i^\top q_2)] \geq \frac{1}{\pi} \|\tilde{q}_1 - \tilde{q}_2\| - \sqrt{n \frac{\pi}{2}} \epsilon \quad (34)$$

Since $\|\tilde{q}_1 - \tilde{q}_2\| \geq \|q_1 - q_2\| - 2\epsilon$, for any $\delta > 0$, we can write

$$\mathbb{P} [\forall q_1 \in B_\epsilon(\tilde{q}_1), \forall q_2 \in B_\epsilon(\tilde{q}_2) : \text{sign}(a_i^\top q_1) \neq \text{sign}(a_i^\top q_2) \mid \|q_1 - q_2\| > \delta] \geq \frac{\delta}{\pi} - (\frac{2}{\pi} + \sqrt{n \frac{\pi}{2}} \epsilon)$$

By setting $\epsilon = \frac{\pi \delta}{4 + \pi \sqrt{2\pi n}}$ and reversing the inequality we obtain

$$\mathbb{P} [\exists q_1 \in B_\epsilon(\tilde{q}_1), \exists q_2 \in B_\epsilon(\tilde{q}_2) : \text{sign}(a_i^\top q_1) = \text{sign}(a_i^\top q_2) \mid \|q_1 - q_2\| > \delta] \leq 1 - \delta/2 \quad (35)$$

Extending these bound to all the rows of $A \in \mathbb{R}^{m \times n}$ drawn from a standard Gaussian distribution, we have

$$\mathbb{P}[\exists q_1 \in B_\epsilon(\tilde{q}_1), \exists q_2 \in B_\epsilon(\tilde{q}_2) : \text{sign}(Aq_1) = \text{sign}(Aq_2) \mid \|q_1 - q_2\| > \delta] \leq (1 - \delta/2)^m \quad (36)$$

Applying a union bound to all points $q \in Q_\epsilon$. Since there are $\binom{|Q_\epsilon|}{2} \leq |Q_\epsilon|^2 \leq \epsilon^{-2k}$ pairs of points, we obtain

$$\mathbb{P}[\exists x_1, x_2 \in \mathcal{X} : \text{sign}(Ax_1) = \text{sign}(Ax_2) \mid \|x_1 - x_2\| > \delta] \leq \left(\frac{4 + \pi\sqrt{2\pi n}}{\pi\delta}\right)^{2k} (1 - \delta/2)^m \quad (37)$$

$$\leq \exp\left(2k \log\left(\frac{4 + \pi\sqrt{2\pi n}}{\pi\delta}\right) + m \log(1 - \delta/2)\right) \quad (38)$$

Upper bounding this probability by ξ and using the fact that $1 - \delta/2 \leq \exp(\delta/2)$ and $\frac{4 + \pi\sqrt{2\pi n}}{\pi} \leq \sqrt{n}(\frac{4}{\pi} + \sqrt{2\pi}) \leq 4\sqrt{n}$, we obtain

$$2k \log \frac{4\sqrt{n}}{\delta} + m \frac{\delta}{2} \geq \log \xi \quad (39)$$

$$m \geq \frac{2}{\delta} \left(2k \log \frac{4\sqrt{n}}{\delta} + \log \frac{1}{\xi}\right) \quad (40)$$

for all $\delta \leq 4\epsilon_0\sqrt{n}$. \square

C Model Identification Proof

Proof of Theorem 6. We want to prove that $\hat{\mathcal{X}} \subseteq \mathcal{X}_\delta$ holds with high probability w.r.t. a random draw of the operators A_1, \dots, A_G . Equivalently, we need to show that

$$\text{sign}(A_g x_g) = \text{sign}(A_g v) \quad \forall g = 1, \dots, G \quad (41)$$

holds for some $v \in \mathbb{S}^{n-1} \setminus \mathcal{X}_\delta$ and some $x_1, \dots, x_G \in \mathcal{X}$ with probability at most ξ with respect to a random draw of the Gaussian matrices A_1, \dots, A_G . This proof adapts some of the procedures given in (Jacques et al., 2013) to our specific setting. We start by bounding this probability for ϵ -balls around vectors $\tilde{v} \in \mathbb{S}^{n-1} \setminus \mathcal{X}$, $\tilde{x}_1, \dots, \tilde{x}_G \in \mathcal{X}$, that is

$$p_0 = \mathbb{P}[\exists x_1 \in B_\epsilon(\tilde{x}_1), \dots, \exists x_G \in B_\epsilon(\tilde{x}_G), \exists v \in B_\epsilon(\tilde{v}) \mid \forall g = 1, \dots, G, \text{sign}(A_g v) = \text{sign}(A_g x_g)] \quad (42)$$

As every row of each operator A_g is independent, we have

$$p_0 = \prod_{g=1}^G \mathbb{P}[\exists x_g \in B_\epsilon(\tilde{x}_g), v \in B_\epsilon(\tilde{v}) \mid \text{sign}(a_{g,i}^\top v) = \text{sign}(a_{g,i}^\top x_g)]^m \quad (43)$$

$$= \prod_{g=1}^G (1 - \mathbb{P}[\forall x_g \in B_\epsilon(\tilde{x}_g), v \in B_\epsilon(\tilde{v}) \mid \text{sign}(a_{g,i}^\top v) \neq \text{sign}(a_{g,i}^\top x_g)])^m \quad (44)$$

Applying Lemma 9, we have that

$$\mathbb{P}[\forall x_g \in B_\epsilon(\tilde{x}_g), \forall v \in B_\epsilon(\tilde{v}) \mid \text{sign}(a_{g,i}^\top v) \neq \text{sign}(a_{g,i}^\top x_g)] \geq d(\tilde{x}_g, \tilde{v}) - \sqrt{n\frac{\pi}{2}\epsilon} \quad (45)$$

where the angular distance can be bounded by the Euclidean distance

$$\pi d(\tilde{x}_g, \tilde{v}) \geq 2 \sin\left(\frac{\pi}{2} d(\tilde{x}_g, \tilde{v})\right) = \|\tilde{x}_g - \tilde{v}\| \geq \delta \quad (46)$$

for all $x_g \in B_\epsilon(\tilde{x}_g)$ and all $v \in B_\epsilon(\tilde{v})$. Plugging this into (44), we have

$$p_0 \leq \left(1 - \frac{\delta}{\pi} + \sqrt{n\frac{\pi}{2}\epsilon}\right)^{mG}. \quad (47)$$

We can extend this result to all vectors $v \in \mathbb{S}^{n-1} \setminus \mathcal{X}_\delta$ and $x_1, \dots, x_G \in \mathcal{X}$ by applying a union bound over a covering of the product set $\mathcal{X}^G \times (\mathbb{S}^{n-1} \setminus \mathcal{X}_\delta)$. Since we can cover \mathcal{X} with ϵ^{-k} balls with $\epsilon \leq \epsilon_0$ due to the assumption that $\text{boxdim}(\mathcal{X}) < k$, and also cover $\mathbb{S}^{n-1} \setminus \mathcal{X}_\delta$ with $(3/\epsilon)^n$ balls, we have

$$\mathbb{P}[\exists x_1, x_2, \dots, x_G \in \mathcal{X}, \exists v \in (\mathbb{S}^{n-1} \setminus \mathcal{X}_\delta) \mid \forall g = 1, \dots, G, \text{sign}(A_g v) = \text{sign}(A_g x_g)] \leq \epsilon^{-kG} (\epsilon/3)^{-n} p_0 \quad (48)$$

Bounding this probability by ξ , we obtain the following inequality

$$\epsilon^{-kG} (\epsilon/3)^{-n} (1 - \frac{\delta}{\pi} + \sqrt{n \frac{\pi}{2}} \epsilon)^{mG} \leq \xi \quad (49)$$

$$\epsilon^{-kG+n} 3^n (1 - \frac{\delta}{\pi} + \sqrt{n \frac{\pi}{2}} \epsilon)^{mG} \leq \xi \quad (50)$$

Solving for m and choosing $\epsilon = \sqrt{\frac{2(4-\pi)^2}{\pi^3 n}} \delta \approx 0.23 \sqrt{\frac{1}{n}} \delta$ we get

$$m \geq \frac{1}{\log(1 - \frac{\delta}{4})} \left\{ \left(k + \frac{n}{G}\right) \log \frac{5\sqrt{n}}{\delta} + \frac{1}{G} \log \frac{1}{\xi} + \frac{n}{G} \log 3 \right\} \quad (51)$$

for $\delta < 4\sqrt{n}\epsilon_0$. Finally, using the fact that $\log(1 - \frac{\delta}{4}) \geq \delta/4$, we obtain the desired bound,

$$m \geq \frac{4}{\delta} \left\{ \left(k + \frac{n}{G}\right) \log \frac{5\sqrt{n}}{\delta} + \frac{1}{G} \log \frac{1}{\xi} + \frac{n}{G} \log 3 \right\}. \quad (52)$$

□

Derivation of δ . The bound in Theorem 6, that is

$$m \geq \frac{4}{\delta} \left\{ \left(k + \frac{n}{G}\right) \log \frac{5\sqrt{n}}{\delta} + \frac{1}{G} \log \frac{1}{\xi} + \frac{n}{G} \log 3 \right\} \quad (53)$$

can be rewritten as a function of δ as

$$\log \delta + \delta a \geq b \quad (54)$$

where

$$a = \frac{m}{4(k + \frac{n}{G})} \quad (55)$$

$$b = \log 5\sqrt{n} + \frac{1}{(Gk + n)} \log \frac{1}{\xi} + \frac{n}{(Gk + n)} \log 3 \quad (56)$$

Using this notation, the inequality in (54) can be further simplified as

$$\log \delta + \delta a \geq b \quad (57)$$

$$\delta e^{a\delta} \geq e^b \quad (58)$$

$$a\delta \exp a\delta \geq a e^b \quad (59)$$

$$a\delta \geq W(a e^b) \quad (60)$$

$$\delta \geq \frac{1}{a} W(a e^b) \quad (61)$$

where the line before the last uses the fact that the inverse of xe^x is the Lambert W function denoted as $W(\cdot)$. Since $W(x) \geq \log x - \log \log x$ for all $x \geq e$, we can write

$$\delta \geq \frac{1}{a} \log(a e^b) - \frac{1}{a} (\log \log a + \log b) \quad (62)$$

$$\delta \geq \frac{1}{a} (\log a + b) \quad (63)$$

Finally, observing that $a \approx \frac{m}{k + \frac{n}{G}}$ and $b \approx \log n$ for large m, n and G , we get

$$\delta = \mathcal{O} \left(\frac{k + \frac{n}{G}}{m} \log \frac{mn}{k + \frac{n}{G}} \right) \quad (64)$$

which, for $n/G \ll k$ is

$$\delta = \mathcal{O} \left(\frac{k}{m} \log \frac{mn}{k} \right). \quad (65)$$

D Sample Complexity Proof

Proof of Theorem 8. We aim to bound the number of different cells associated with the binary mapping $\text{sign}(A \cdot)$ which contain at least one element from the signal set \mathcal{X} , i.e., $|\text{sign}(A\mathcal{X})|$. Our strategy consists in obtaining a global bound on the number of discontinuities of the binary mapping over the image of a covering of \mathcal{X} , which can then be related to the number of different cells that contain at least one element of \mathcal{X} .

For $\epsilon < \epsilon_0$, let $\mathcal{Q}_\epsilon \subset \mathcal{X}$ be an optimal ϵ covering of \mathcal{X} , i.e., $\mathcal{X} \subset \mathcal{X}_\epsilon + \epsilon \mathbb{B}^n$. If $\text{boxdim}(\mathcal{X}) < k$, then there exists an $\epsilon_0 \in (0, \frac{1}{2})$ such that $|Q_\epsilon| \leq \epsilon^{-k}$ for all $\epsilon < \epsilon_0$. Let us define the binary mapping $\Phi(\cdot) := \text{sign}(A \cdot)$ and the number $Z(S)$ of its discontinuous components over a set $S \subset \mathbb{S}^{n-1}$, i.e.,

$$Z(S) := |\{i : \Phi_i \notin C^0(S)\}|.$$

Combining Corollary 10 with a union bound argument and setting $t = \eta(\epsilon)$, we get that with probability exceeding

$$1 - |Q_\epsilon| \exp(-\frac{3}{2} \frac{mt^2}{3\eta(\epsilon)+t}) \geq 1 - \exp(k \log(\frac{1}{\epsilon}) - \frac{3}{8} m\eta(\epsilon)),$$

we have, for $V_\epsilon(q) := B_\epsilon(q) \cap \mathbb{S}^{n-1}$

$$Z(V_\epsilon(q)) \leq 2m\eta(\epsilon),$$

for all $q \in \mathcal{X}_\epsilon$. Therefore, provided $\eta(\epsilon)m = \epsilon m \sqrt{\frac{\pi}{2}n} \geq \frac{16}{3} k \log(\frac{1}{\epsilon})$, the above probability exceeds $1 - \exp(-\frac{3}{16} m\eta(\epsilon))$.

Given $s > 1$, with $s = \mathcal{O}(1)$, let us set ϵ as

$$\epsilon = \frac{16}{3} s \frac{k}{m\sqrt{n}} \log(\frac{3m\sqrt{n}}{16sk}),$$

or $\eta(\epsilon)m = \frac{16}{3} sk \log(\frac{3m\sqrt{n}}{16sk})$, under the assumption that $\epsilon_0 > \frac{16}{3} s \frac{k}{m\sqrt{n}} \log(\frac{3m\sqrt{n}}{16sk})$. Assuming $\frac{m\sqrt{n}}{sk} > e$, this means that $\log(1/\epsilon) \leq \log(\frac{3m\sqrt{n}}{16sk})$, so that the previous requirement on $\eta(\epsilon)m$ holds if

$$sk \log(\frac{3m\sqrt{n}}{16sk}) \geq k \log(\frac{3m\sqrt{n}}{16sk}),$$

which is reached as soon as $s > 1$.

Therefore, with probability exceeding $1 - \exp(-\frac{3}{16} m\eta(\epsilon)) = 1 - \exp(-sk \log(\frac{3m\sqrt{n}}{16sk})) = 1 - (\frac{3m\sqrt{n}}{16sk})^{-sk} \geq 1 - (\frac{3m\sqrt{n}}{16s})^{-s}$, we thus know that for any $q \in Q_\epsilon$,

$$|\Phi(V_\epsilon(q))| \leq 2^{2sk \log(\frac{3m\sqrt{n}}{16sk})}.$$

Since $\epsilon^{-k} = 2^{k \log_2(1/\epsilon)} \leq 2^{(\log 2)^{-1} k \log(\frac{3m\sqrt{n}}{16sk})}$ for some $c > 0$,

$$|\text{sign}(A\mathcal{X})| \leq \sum_{q \in \mathcal{X}_\epsilon} |\Phi(V_\epsilon(q))| \leq \epsilon^{-k} 2^{2sk \log(\frac{3m\sqrt{n}}{16sk})} \leq 2^{((\log 2)^{-1} + 1) sk \log(\frac{3m\sqrt{n}}{16sk})} \leq \left(\frac{2m\sqrt{n}}{sk} \right)^{4sk}.$$

We conclude by choosing $s = 2$ to obtain the bound $|\text{sign}(A\mathcal{X})| \leq \left(\frac{m\sqrt{n}}{k} \right)^{8k}$ with probability exceeding $1 - \frac{1024}{9m^2n}$ which holds as long as $\epsilon_0 > \frac{32}{3} \frac{k}{m\sqrt{n}} \log(\frac{3m\sqrt{n}}{32k})$.

□