

Sensing Theorems for Unsupervised Learning in Linear Inverse Problems

Julián Tachella

*Laboratoire de Physique
CNRS, ENSL
Lyon, F-69364, France*

JULIAN.TACHELLA@CNRS.FR

Dongdong Chen

*School of Engineering
University of Edinburgh
Edinburgh, EH9 3FB, UK*

D.CHEN@ED.AC.UK

Mike Davies

*School of Engineering
University of Edinburgh
Edinburgh, EH9 3FB, UK*

MIKE.DAVIES@ED.AC.UK

Editor: Aapo Hyvärinen

Abstract

Solving an ill-posed linear inverse problem requires knowledge about the underlying signal model. In many applications, this model is a priori unknown and has to be learned from data. However, it is impossible to learn the model using observations obtained via a single incomplete measurement operator, as there is no information about the signal model in the nullspace of the operator, resulting in a chicken-and-egg problem: to learn the model we need reconstructed signals, but to reconstruct the signals we need to know the model. Two ways to overcome this limitation are using multiple measurement operators or assuming that the signal model is invariant to a certain group action. In this paper, we present necessary and sufficient sensing conditions for learning the signal model from measurement data alone which only depend on the dimension of the model and the number of operators or properties of the group action that the model is invariant to. As our results are agnostic of the learning algorithm, they shed light into the fundamental limitations of learning from incomplete data and have implications in a wide range set of practical algorithms, such as dictionary learning, matrix completion and deep neural networks.

Keywords: Inverse problems, unsupervised learning, blind compressed sensing, equivariance, computational imaging.

1. Introduction

Inverse problems are ubiquitous in science and engineering applications such as computed tomography (CT) (Jin et al., 2017), depth ranging (Rapp et al., 2020) and non-line-of-sight imaging (O’Toole et al., 2018). In this paper, we consider linear inverse problems that consist of reconstructing a signal $x \in \mathcal{X} \subset \mathbb{R}^n$ from incomplete and noisy measurements $y \in \mathcal{Y} \subseteq \mathbb{R}^m$, that is

$$y = Ax + \epsilon \tag{1}$$

where ϵ denotes the noise affecting the measurements. This is generally an ill-posed task due to the incomplete forward operator A with $m < n$ and the presence of noise. Classical approaches assume a signal model using some prior knowledge about the underlying signals. For example, the well-known total variation model (Rudin et al., 1992) is built on the prior belief that natural images are approximately piecewise smooth. This approach often yields a loose description of the true model, providing biased and/or suboptimal reconstructions. In recent years, this approach has been replaced by learning the signal model directly from data. For example, a common approach is to learn a model consisting of a dictionary of image patches using a dataset of natural images (Studer and Baraniuk, 2012). In a similar vein, it is possible to learn directly the reconstruction function $f : y \mapsto x$ via deep neural networks using multiple training pairs (x_i, y_i) (Jin et al., 2017).

Despite the appeal and better performance of the learning-based approach, in many sensing applications such as medical and scientific imaging we can only access measurements y_i which hinders supervised learning solutions. Moreover, if the measurement process A is incomplete, it is fundamentally impossible to identify the signal set¹ \mathcal{X} through measurements alone even in the absence of noise, as there is no information about the set of signals in the nullspace of A . In other words, we cannot uniquely identify the signal set \mathcal{X} from the noiseless measurement set $\mathcal{Y} = A\mathcal{X}$. It is also impossible to learn the reconstruction function $f : y \mapsto x$, independently of the number of observed measurement vectors y_i , as explained by the following simple proposition:

Proposition 1 (Chen et al., 2021) *Any reconstruction function $f : \mathbb{R}^m \mapsto \mathbb{R}^n$ of the form*

$$f(y) = A^\dagger y + v(y) \tag{2}$$

where A^\dagger is the linear pseudo-inverse of A and $v(y)$ is any function whose image belongs to the nullspace of A verifies measurement consistency $Af(y_i) = y_i$.

Proof For any f of the form (2), the measurement consistency can be expressed as $Af(y) = AA^\dagger y + Av(y)$ where the first term is simply y as AA^\dagger is the identity matrix, and $Av(y) = 0$ for any $v(y)$ in the nullspace of A . ■

Hence the chicken-and-egg problem: in order to reconstruct x we need the signal model, but to learn this model we require some ground truth training data x_i .

Here we show that this fundamental limitation can be overcome if \mathcal{X} is low-dimensional, either by using information from multiple incomplete measurement operators $A_1, \dots, A_{|\mathcal{G}|}$ or by having a single operator A and exploiting weak prior information associated with group invariance properties of the signal model.

Multiple measurement operators can provide additional information about the model if the operators have different nullspaces. This idea was used in various inverse problems such as image inpainting (Studer and Baraniuk, 2012), magnetic resonance imaging (MRI) (Liu et al., 2020) and hyperspectral imaging (Yang et al., 2015b). Recently, the equivariant imaging framework (Chen et al., 2021, 2022) empirically showed that learning the signal model with a single operator A is also possible if the signal model is invariant to a certain

1. Throughout the paper we use the terms *signal model* and *signal set* interchangeably to denote the support \mathcal{X} of the signal distribution $p(x)$.

Inverse problem	Forward operator	Group action/ Mult. operator	Algorithm
Image inpainting	Binary mask	Permutations 2D shifts	Dictionary learning (Studer and Baraniuk, 2012)
Computed tomography	Sparse-view Radon transform	2D rotations	Equivariant Imaging (Chen et al., 2021, 2022)
MRI	Subsampled Fourier	Multiple op.	Deep networks (Liu et al., 2020)
		2D rotations	Equivariant Imaging (Chen et al., 2022)
High-speed video	Binary mask	2D shifts	Low-rank GMM (Yang et al., 2015b)
Depth completion	Binary mask	Multiple op.	Deep networks (Jaritz et al., 2018)
Hyperspectral imaging	Spectral domain CS matrix	Multiple op.	Low-rank GMM (Yang et al., 2015b)
Electron microscopy	2D projection	3D rotations and shifts	Cryo-GAN (Gupta et al., 2020)
Motion segmentation	Binary mask	Permutations	SSC (Yang et al., 2015a)
Single image view synthesis	Light transport	3D rotations	Coordinate-based neural representation (Tancik et al., 2021)
Compressed Sensing	Random Gaussian	Permutations	Approximate Message Passing (Guo and Davies, 2015) (Metzler et al., 2018)

Table 1: Examples of applications where a low-dimensional signal model is learned via multiple measurement operators or assuming that the model is invariant to a group action.

group of transformations² $T_1, \dots, T_{|\mathcal{G}|}$, as they give access to multiple *virtual* operators $AT_1, \dots, AT_{|\mathcal{G}|}$ with possibly different nullspaces. This strategy offers an appealing way to learn the model, as most real-world signal sets present certain invariances, such as invariance to translations and rotations in images.

Table 1 presents various inverse problems and related reconstruction algorithms where information in the nullspace of A is obtained using multiple measurement operators or exploiting invariance of the signal model to the action of a group. Despite the empirical successes of these methods, theoretical guarantees for model identification are still lacking: what are the requirements on the measurement operators or model invariance properties? When is it possible learn the model and reconstruct the signals? The two fundamental problems of signal recovery and model identification can be formally defined as follows:

2. Formally, these transformations constitute the action of a group (also called the representation of a group).

Signal Recovery Signals x can be uniquely recovered from observations $y = Ax$ if for every measurement vector $y \in \mathcal{Y}$ there is a unique signal $x \in \mathcal{X}$ that verifies the observed measurements $y = Ax$.

Model Identification We define a signal set \mathcal{X} to be identifiable within a given a class \mathcal{C} if it is uniquely defined from the observed measurement sets $\mathcal{Y}_g = A_g\mathcal{X}$ with $g = 1, \dots, |\mathcal{G}|$.

In general, there can be a unique solution for none of the problems, just one or both. There might be a unique solution for signal recovery if the signal model is known, but it might be impossible to identify the model in the first place. For example, this is the case of blind compressed sensing (Gleichman and Eldar, 2011). The converse is also possible, that is, uniquely identifying a model without having enough measurements per sample to uniquely identify the associated signal. For example, multiple single ($m = 1$) measurements $y_i = a_{g_i}^\top x_i$ are sufficient to identify a low-dimensional signal subspace (Chen et al., 2015; Gribonval et al., 2017), however it is not possible to recover the signal x_i linked to each observed measurement y_i .

It is well-known from generalized compressed sensing theory that unique signal recovery is possible if the nullspace of A does not contain any vector in the difference set of \mathcal{X} (Bourrier et al., 2014). This condition can be achieved by the class of *low-dimensional signal sets*, i.e., where the dimension of \mathcal{X} is both smaller than the ambient dimension n and observed measurements m . For example, in standard compressed sensing signals are assumed to belong the k -sparse set, where \mathcal{X} has dimension $k < m$. On the other hand, to the best of our knowledge, the model identification problem has been mostly analysed in the context of matrix completion (Candès and Recht, 2009), where \mathcal{X} is assumed to belong the class of low-dimensional subspaces or union of subspaces (Eriksson et al., 2012), and the observations $y_i = A_{g_i}x$ consist of different permutations of a binary mask, i.e., $A_g = AT_g$ where A is a diagonal matrix indicating the observed entries and $T_1, \dots, T_{|\mathcal{G}|}$ are permutation matrices.

In this paper, we study necessary and sufficient conditions on the number of measurements to guarantee a unique solution to both problems for the general class of low-dimensional signal models, only relying on their intrinsic dimension and weak prior information related to invariance properties of the signal model.

1.1 Summary of Contributions

We consider the class of k -dimensional signal sets (the definition of dimension is made clear in Section 3.2), and study model identifiability from incomplete measurements associated with a finite number of measurement operators where each training sample consists of an observation of a signal from the set \mathcal{X} through one of a set of $|\mathcal{G}|$ measurement operators, i.e., $y_i = A_{g_i}x_i + \epsilon_i$ with $g_i \in \{1, \dots, |\mathcal{G}|\}$, where $A_g \in \mathbb{R}^{m \times n}$ and $m < n$. We also consider the case where we have a single operator A , i.e., $y_i = Ax_i + \epsilon_i$, but the signal set is invariant to the action of a group \mathcal{G} . Our main results are:

- When the dimension of the model is large, i.e., $k = n$, the signal model cannot be uniquely identified if $m < n$.




Figure 1: Model identification and signal recovery regimes as a function of the number of partial observations m per signal, model dimension k , ambient dimension n and number of measurement operators $|\mathcal{G}|$ (when it is possible to access multiple independent operators) or maximum multiplicity of an invariant subspace $\max_j c_j/s_j$ (when the signal is group invariant or the operators are related via a group action).

- When the dimension of the model is small, $k < n$, model uniqueness is possible.
 - We first consider the case where we obtain measurements via a set of $|\mathcal{G}|$ different operators $A_1, \dots, A_{|\mathcal{G}|}$.
 - * A necessary condition for model identification is $m \geq n/|\mathcal{G}|$.
 - * Unique model identification is possible via almost every set of $|\mathcal{G}|$ operators $A_1, \dots, A_{|\mathcal{G}|} \in \mathbb{R}^{m \times n}$ (w.r.t. the Lebesgue measure on $\mathbb{R}^{|\mathcal{G}|mn}$) when $m > k + n/|\mathcal{G}|$.
 - Second, we consider the case where we have a single operator A but the signal set is invariant to a group of transformations. Our results are stated as a function of $\max_j c_j/s_j$, where c_j and s_j are the dimension and multiplicity of the j th irreducible representation of the group action³.
 - * A necessary condition for model identification for a compact group is $m \geq \max_j c_j/s_j$.
 - * We show that, if \mathcal{G} is a cyclic group, unique model identification is possible by almost every operator $A \in \mathbb{R}^{m \times n}$ (w.r.t. the Lebesgue measure on \mathbb{R}^{mn}) with $m > 2k + \max_j c_j + 1$. As cyclic groups have all irreducibles with dimension $s_j = 1$, the bound is equivalent to $m > 2k + \max_j c_j/s_j + 1$. We conjecture that this last bound holds for more general groups with $s_j \geq 1$.
 - * We characterize a subset of operators that fail to provide model uniqueness: if the forward operator A is itself equivariant to the action of the group, it is impossible to uniquely identify the model if $m < n$, even for arbitrarily small k and large number of transformations $|\mathcal{G}|$.
- We experimentally show that our bounds accurately characterize the performance of popular learning algorithms on synthetic and real datasets.

As summarized in Figure 1, our results characterize the sensing regimes for model identification, complementing the existing results for signal recovery. In particular, we provide

³. A detailed explanation of these concepts from linear representation theory can be found in Section 4.

necessary and sufficient conditions for identifying a low-dimensional group invariant signal model from incomplete measurements of a single operator as a function of the properties of the group. For example, consider 1D signals defined on a discrete grid of length n . We can uniquely identify a k dimensional model with invariance to cyclic shifts using $m > 2k + 2$ measurements, whereas for a model with invariance to reflection we can only guarantee identifiability using $m > 2k + n/2 + 1$ measurements. It is important to note that our identifiability results are purely geometric and do not cover the number of samples needed to learn the model structure (the so-called sample complexity). Quantifying the sample complexity would require additional assumptions on the class of signal models, the noise and the choice of learning algorithm, and is out of scope of the present paper. However, in the experiments, we show that our bounds capture the behaviour of various learning algorithms in practical datasets.

1.2 Related Work

Some preliminary results of this work appear in an earlier manuscript (Tachella et al., 2022), which only covers the case of learning from multiple independent operators. This section discusses other related work by topic, from blind compressed sensing to deep neural networks for inverse problems.

Blind Compressed Sensing The fundamental limitation of failing to learn a signal model from incomplete data goes back to blind compressed sensing (Gleichman and Eldar, 2011) for the specific case of models exploiting sparsity on an orthogonal dictionary. In order to learn the dictionary from incomplete observations, (Gleichman and Eldar, 2011) imposes additional constraints on the dictionary, while some subsequent papers (Silva et al., 2011; Aghagolzadeh and Radha, 2015) remove these assumptions by proposing to use multiple operators A_g as studied here. This paper can be seen as a generalization of such results to a much wider class of signal models.

Matrix Completion 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., $y_i = A_{g_i} x_i$ 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 $k \ll n$. This problem can be viewed as the combination of model identification, i.e., identifying the k -dimensional subspace, and signal recovery, i.e., reconstructing the individual columns. If the columns are observed via $|\mathcal{G}|$ sufficiently different patterns A_{g_i} with the same number of entries m , a sufficient condition (Pimentel-Alarcón et al., 2016) for uniquely recovering almost every subspace model is⁴ $m \geq (1 - 1/|\mathcal{G}|)k + n/|\mathcal{G}|$. A similar necessary condition was shown in Pimentel-Alarcon and Nowak (2016) for the case of *high-rank* matrix completion (Eriksson et al., 2012), which arises when the samples x_i belong to a union of k -dimensional subspaces.

4. A larger number of measurements $m = \mathcal{O}(k \log n)$ is required to guarantee a stable recovery when the number of patterns $|\mathcal{G}|$ is large (Candès and Recht, 2009).

We show that unique recovery is possible for almost every set of $|\mathcal{G}|$ operators with $m > k + n/|\mathcal{G}|$ measurements, however the theory presented here goes beyond union of subspaces, being also valid for general low-dimensional models.

Multireference Alignment and Electron Microscopy In the multireference alignment problem (Aguerreberre et al., 2016), the goal is to reconstruct a signal x from shifted and incomplete measurements, i.e.,

$$g_i \sim \mathcal{G} \tag{3}$$

$$y_i = AT_{g_i}x + \epsilon_i \tag{4}$$

for $i = 1, \dots, N$, where the shifts T_{g_i} are a priori unknown. This problem can be generalized beyond shifts by considering transformations related to the action of a group \mathcal{G} . This is the case of electron-microscopy imaging, where \mathcal{G} is the group of rotations and translations of a particle (Gupta et al., 2020). The multireference alignment problem can be seen as a special case of our framework with a single operator A and a \mathcal{G} -invariant signal set of dimension⁵ $k = 0$ which consists of all transformations of a single signal. If \mathcal{G} is cyclic (e.g., shifts), our results guarantee model identification for almost every A with $m > \max_j c_j + 1$ measurements.

Deep Nets for Inverse Problems Despite providing very competitive results, most deep learning based solvers require measurements and signal pairs (x_i, y_i) 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. However, their ideas only apply to denoising settings where there is a trivial nullspace, as the operator A is the identity matrix. In Artifact2Artifact (Liu et al., 2020), it was empirically shown that it is possible to exploit different measurement operators to learn the reconstruction function in the context of MRI. AmbientGAN (Bora et al., 2018) proposed to learn the signal distribution from incomplete measurements of multiple forward operators, however they only provide reconstruction guarantees for the case where an infinite number of operators is available⁶, $|\mathcal{G}| = \infty$, a condition that is not met in practice. Finally, the equivariant imaging framework (Chen et al., 2021) showed empirically that learning the reconstruction function using a single operator A is possible in various imaging inverse problems when the signal model presents some group invariance. Here we provide a theoretical framework to support all of these findings.

2. Signal Recovery Preliminaries

Let $A^\dagger \in \mathbb{R}^{n \times m}$ be the linear pseudo-inverse of A . We denote the range space of A^\dagger as \mathcal{R}_A . Its complement, the nullspace of A , is denoted as \mathcal{N}_A , where $\mathcal{R}_A \oplus \mathcal{N}_A = \mathbb{R}^n$ and \oplus denotes the direct sum. Throughout the paper, we assume that the signals are associated with a distribution $p(x)$ supported on the signal set $\mathcal{X} \subset \mathbb{R}^n$. Signal recovery has a unique solution if and only if the forward operator restricted⁷ to the signal set \mathcal{X} is one-to-one, i.e., if for

5. For infinite but compact groups, $k = \dim(\mathcal{G})$.

6. Their result relies on the Cramér-Wold theorem, which is discussed in Section 3.1

7. Note that if $m < n$, the operator A cannot be one-to-one in the whole \mathbb{R}^n .

every pair of signals $x_1, x_2 \in \mathcal{X}$ where $x_1 \neq x_2$ we have that

$$Ax_1 \neq Ax_2 \tag{5}$$

$$A(x_1 - x_2) \neq 0. \tag{6}$$

In other words, there is no vector $x_1 - x_2 \neq 0$ with $x_1, x_2 \in \mathcal{X}$ in the nullspace of A . It is well-known that this is only possible if the signal set \mathcal{X} is low-dimensional. There are multiple ways to define the notion of dimensionality of a set in \mathbb{R}^n . In this paper, we focus on the upper box-counting dimension (Falconer, 2004, Chapter 2) which is defined for a compact subset $S \subset \mathbb{R}^n$ as

$$\text{boxdim}(S) = \limsup_{\epsilon \rightarrow 0} \frac{\log \mathbf{N}(S, \epsilon)}{-\log \epsilon} \tag{7}$$

where $\mathbf{N}(S, \epsilon)$ is the minimum number of closed balls of radius ϵ with respect to the norm $\|\cdot\|$ that are required to cover S . This definition of dimension is convenient for the theoretical results, and covers both well behaved models such as compact manifolds where the definition coincides with the more intuitive topological dimension, as well as more general sets⁸. The sensing mapping is one-to-one for almost every operator $A \in \mathbb{R}^{m \times n}$ when restricted to \mathcal{X} if (Sauer et al., 1991)

$$m > \text{boxdim}(\Delta\mathcal{X}) \tag{8}$$

where $\Delta\mathcal{X}$ denotes the difference set defined as

$$\Delta\mathcal{X} = \{\Delta x \in \mathbb{R}^n \mid \Delta x = x_2 - x_1, x_1, x_2 \in \mathcal{X}, x_2 \neq x_1\}. \tag{9}$$

The term *almost every* means that the complement has Lebesgue measure 0 in the space of linear measurement operators $\mathbb{R}^{m \times n}$. Almost every statements are used to exclude *small* subsets of operators where one-to-oneness fails. The difference set of models of dimension k generally has dimension $2k$, requiring $m > 2k$ measurements to ensure signal recovery. For example, the union of k -dimensional subspaces requires at least $2k$ measurements to guarantee one-to-oneness (Blumensath and Davies, 2009). This includes well-known models such as k -sparse models (e.g., convolutional sparse coding (Bristow et al., 2013)) and co-sparse models (e.g., total variation (Rudin et al., 1992)). In the regime $k < m \leq 2k$, the subset of signals where one-to-oneness fails is typically at most $(2k - m)$ -dimensional (Sauer et al., 1991). While the bound in (8) guarantees *unique* signal recovery, more measurements are typically necessary in order to have a *stable* inverse $f : y \mapsto x$, i.e., possessing a certain Lipschitz constant. A detailed discussion can be found for example in Puy et al. (2017).

3. Model Identifiability with Multiple Operators

We first focus on the noiseless case to study the intrinsic identifiability problems associated to having only incomplete measurement data. The effect of noise will be discussed in Section 5. In this section, we assume that we observe the measurement sets $\mathcal{Y}_g = A_g\mathcal{X}$ defined as

$$\mathcal{Y}_g = \{y \in \mathbb{R}^m \mid y = A_g x, x \in \mathcal{X}\} \tag{10}$$

8. As discussed in Section 3.2, we can extend the box-counting dimension for non-compact conic sets.

for $g = 1, \dots, |\mathcal{G}|$. In practice, we will only observe a finite number of measurements y_i , where the i th signal is associated to one of $|\mathcal{G}|$ linear operators with $m < n$, i.e.,

$$y_i = A_{g_i} x_i \quad (11)$$

where $g_i \in \{1, \dots, |\mathcal{G}|\}$ and $i = 1, \dots, N$. While we assume that the measurement operator A_{g_i} is known for all observed signals, it is important to note that we do not know a priori if two observations (y_i, A_{g_i}) and $(y_{i'}, A_{g_{i'}})$ are related to the same⁹ underlying signal x . We assume that all operators have the same number of measurements m throughout most of this section, however we also discuss the case of different number of measurements in Section 3.3. We begin with a simple toy example illustrating this observation model:

Example 1 (One-dimensional toy model) *Consider a one-dimensional linear subspace spanned by a triangular signal $\phi \in \mathbb{R}^n$, i.e., $x_i = c_i \phi \in \mathcal{X} \subset \mathbb{R}^n$ with $c_i \in \mathbb{R}$, which is only partially observed through a simple masking operator, as illustrated in Figure 2. The single sensing operator does not provide information outside the observed part, thus there are infinitely many one-dimensional models that fit the observed measurements. However, if we also obtain measurements through multiple masks A_{g_i} , then identifying the model becomes feasible. In this case, solving both model identification and signal recovery problem boils down to low rank matrix completion of the measurement matrix $[y_1, \dots, y_N]$.*

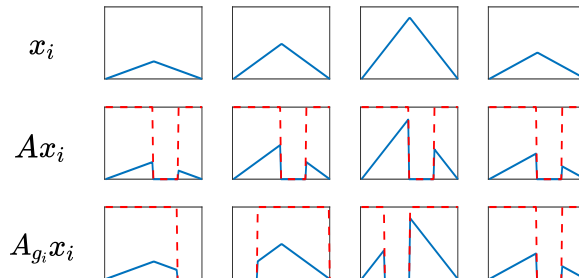


Figure 2: Learning a one-dimensional model from incomplete measurement data. Each column shows a different sample from the signal set. The sensing mask is shown in red.

3.1 Uniqueness of Any Model?

A natural first question when considering uniqueness of the signal model is: can we recover any distribution $p(x)$ observed via a finite number of forward operators A_g , even in the case where the support \mathcal{X} is the full ambient space \mathbb{R}^n ? We show that, in general, the answer is no.

Uniqueness can be analyzed from the point of view of the characteristic function of $p(x)$, defined as $\varphi(w) = \mathbb{E}\{e^{iw^\top x}\}$ where the expectation is taken with respect to $p(x)$ and $i = \sqrt{-1}$ is the imaginary unit. If two distributions have the same characteristic function,

9. The Noise2Noise approach (Lehtinen et al., 2018) does require such pairs of observations, which is impractical in many real-world settings. The theory presented here does not require such an assumption.

then they are necessarily the same almost everywhere. Each forward operator provides information about a subspace of the characteristic function as

$$\mathbb{E}\{e^{iw^\top A_g^\dagger y}\} = \mathbb{E}\{e^{iw^\top A_g^\dagger A_g x}\} \quad (12)$$

$$= \mathbb{E}\{e^{i(A_g^\dagger A_g w)^\top x}\} \quad (13)$$

$$= \varphi(A_g^\dagger A_g w) \quad (14)$$

where $A_g^\dagger A_g$ is a linear projection onto the subspace \mathcal{R}_{A_g} . Given that $m < n$, the characteristic function is only observed in the subspaces \mathcal{R}_{A_g} for all $g \in \{1, \dots, |\mathcal{G}|\}$. For any finite number of operators, the union of these subspaces does not cover the whole \mathbb{R}^n , and hence there is loss of information, as the signal model is not uniquely defined.

In the case of an infinite number of operators, $|\mathcal{G}| = \infty$, there is a well-known case where model uniqueness is possible: the Cramér-Wold theorem guarantees uniqueness of the signal distribution if all possible one dimensional projections ($m = 1$) are available (Cramér and Wold, 1936; Bora et al., 2018). However, in most practical settings we can only access a finite set of operators and many distributions are non-identifiable.

3.2 Uniqueness of Low-Dimensional Models

Most models appearing in signal processing and machine learning are assumed to be approximately low-dimensional, with a dimension k which is much lower than the ambient dimension n . As discussed in Section 2, the low-dimensional property is fundamental to obtain stable reconstructions in ill-posed problems with $m < n$. In the rest of paper, we impose the following assumption on the class of signal models:

A1 The signal set \mathcal{X} is either

- (a) A bounded set with box-counting dimension k .
- (b) An unbounded conic set whose intersection with the unit sphere has box-counting dimension $k - 1$.

This assumption has been widely adopted in the inverse problems literature, as it is necessary to guarantee signal recovery. Our definition of dimension covers most models used in practice, such as single subspace models (principal component analysis), union of subspaces (convolutional sparse coding models, k -sparse models), low-rank matrices and compact manifolds. It also covers signal sets that are well described by a deep generative model (e.g., a variational autoencoder or a generative adversarial network), whose box-counting dimension is given by the dimension of the latent space. Figure 3 shows two examples of common signal sets with small box-counting dimension. In real-world problems, we might not know the box-counting dimension of the dataset a priori. Nonetheless, there exist a range of algorithms that compute the box-counting dimension of real datasets (Hein and Audibert, 2005) which can be applied on measurement data alone¹⁰.

10. If A is one-to-one when restricted to \mathcal{X} (which happens for almost every A with $m > 2k$ measurements, as explained in Section 2), the box-counting dimension of the measurement data equals the box-counting dimension of the underlying signal set.




Figure 3: Examples of low-dimensional signal sets. a) the intersection of a union of k -dimensional subspaces with the unit sphere has box-counting dimension $k - 1$. b) the box-counting dimension of a generative model with a Lipschitz continuous decoder equals the dimension of the latent space.

In the rest of the paper, we focus on conditions for the identification of the support \mathcal{X} instead of the signal distribution $p(x)$, due to the following observation: if there is a one-to-one reconstruction function $f : y \mapsto x$ (which holds for almost every operator with $m > 2k$ measurements, as explained in Section 2), uniqueness of the support necessarily implies uniqueness of the distribution $p(x)$:

Proposition 2 *Assume there is a one-to-one mapping between the measurement set \mathcal{Y} and the support \mathcal{X} of the signal distribution $p(x)$. Further assume that the support \mathcal{X} is known, then the signal distribution $p(x)$ can be uniquely identified from the measurement distribution $p(y)$.*

Proof If \mathcal{X} is identifiable from the measurement set \mathcal{Y} and there is a measurable one-to-one mapping from each observed signal y_i to \mathcal{X} , then it is possible to obtain $p(x)$ as the push-forward of the measurement distribution $p(y)$. ■

We begin with a simple example which provides important geometrical intuition of how a low-dimensional model can be learned via multiple projections A_g :

Example 2 *Consider a toy signal model with support $\mathcal{X} \subset \mathbb{R}^3$ which consists of a one-dimensional linear subspace spanned by $\phi = [1, 1, 1]^\top$, and $|\mathcal{G}| = 3$ measurement operators $A_1, A_2, A_3 \in \mathbb{R}^{2 \times 3}$ which project the signals into the $x(3) = 0$, $x(2) = 0$ and $x(1) = 0$ planes respectively, where $x(i)$ denotes the i th entry of the vector x . The example is illustrated in Figure 4. The first operator A_1 imposes a constraint on \mathcal{X} , that is, every $x \in \mathcal{X}$ should verify $x(1) - x(2) = 0$. Without more operators providing additional information about \mathcal{X} , this constraint yields a plane containing \mathcal{X} , and there are infinitely many one-dimensional models that would fit the measurement data perfectly. However, the additional operator A_2 adds the constraint $x(2) - x(3) = 0$, which is sufficient to uniquely identify \mathcal{X} as*

$$\mathcal{X} = \hat{\mathcal{X}} := \{v \in \mathbb{R}^3 \mid v(1) - v(2) = v(2) - v(3) = 0\}$$

is the desired 1-dimensional subspace. Finally, note that in this case the operator A_3 does not restrict the signal set further, as the constraint $x(1) - x(3) = 0$ is verified by the other two constraints.




Figure 4: Toy example of a 1-dimensional subspace embedded in \mathbb{R}^3 . If we only observe the projection of the signal set into the plane $x(3) = 0$, then there are infinite possible lines that are consistent with the measurements. Adding the projection into the $x(1) = 0$ plane, allows us to uniquely identify the signal model.

3.3 Bounds for Independent Operators

The ideas from Example 2 can be generalized and formalized as follows: for each projection A_g , we can identify the signal set that is consistent with the observations through the measurement operator A_g , that is

$$\hat{\mathcal{X}}_g = \{v \in \mathbb{R}^n \mid v = \hat{x}_g + u, \hat{x}_g \in \mathcal{X}, u \in \mathcal{N}_{A_g}\} \quad (15)$$

which has dimension at most $n - (m - k)$. It can clearly be seen that $\mathcal{X} \subseteq \hat{\mathcal{X}}_g$ by considering the subset with $u = 0$. The inferred signal set belongs to the intersection of these sets

$$\hat{\mathcal{X}} = \bigcap_{g \in \mathcal{G}} \hat{\mathcal{X}}_g \quad (16)$$

which can be expressed concisely as

$$\hat{\mathcal{X}} = \{v \in \mathbb{R}^n \mid A_1(x_1 - v) = \dots = A_{|\mathcal{G}|}(x_{|\mathcal{G}|} - v) = 0, x_1, \dots, x_{|\mathcal{G}|} \in \mathcal{X}\}. \quad (17)$$

Even though we have derived the set $\hat{\mathcal{X}}$ from a purely geometrical argument, the constraints in (17) also offer a simple algebraic intuition: the inferred signal set consists of the points $v \in \mathbb{R}^n$ which verify the following system of equations

$$\begin{bmatrix} A_1 \\ \vdots \\ A_{|\mathcal{G}|} \end{bmatrix} v = \begin{bmatrix} A_1 x_1 \\ \vdots \\ A_{|\mathcal{G}|} x_{|\mathcal{G}|} \end{bmatrix}. \quad (18)$$

for all possible choices of $|\mathcal{G}|$ points $x_1, \dots, x_{|\mathcal{G}|}$ in \mathcal{X} . In other words, given a dataset of N incomplete samples $\{A_{g_i} x_i\}_{i=1}^N$, it is possible to build $\hat{\mathcal{X}}$ by trying all the possible combinations of $|\mathcal{G}|$ samples¹¹ and keeping only the points v which are the solutions of (18).

It is trivial to see that $\mathcal{X} \subseteq \hat{\mathcal{X}}$, but when can we guarantee $\mathcal{X} = \hat{\mathcal{X}}$? As in the previous toy example, if there are not enough constraints, e.g., if we have a single A and no additional transformations, the inferred set will have a dimension larger than k , containing undesired aliases. In particular, we have the following necessary condition on the operators:

11. Despite providing a good intuition, this procedure for estimating \mathcal{X} is far from being practical as it would require an infinite number of observed samples if the dimension of the signal set is not trivial $k > 0$.

Proposition 3 (Theorem 1 in Chen et al. (2021)) *Let the operators $A_1, \dots, A_{|\mathcal{G}|} \in \mathbb{R}^{m \times n}$. The signal set \mathcal{X} can be uniquely identified from the sets $\{A_g \mathcal{X}\}_{g=1}^{|\mathcal{G}|}$ only if*

$$\text{rank}\left(\begin{bmatrix} A_1 \\ \vdots \\ A_{|\mathcal{G}|} \end{bmatrix}\right) = n \quad (19)$$

and thus only if $m \geq n/|\mathcal{G}|$.

Proof In order to have model uniqueness, the system in (18) should only admit a solution if $v = x_1 = \dots = x_{|\mathcal{G}|}$. If the rank condition in (19) is not satisfied, there is more than one solution for any choice of $x_1, \dots, x_{|\mathcal{G}|} \in \mathcal{X}$. \blacksquare

Note that this necessary condition does not take into account the dimension of the model. As discussed in Section 3.1, a sufficient condition for model uniqueness must depend on the dimension of the signal set k . The next theorem shows that k additional measurements per operator are sufficient for model identification:

Theorem 4 *Assume that the signal set \mathcal{X} verifies assumption **A1**. For almost every set of $|\mathcal{G}|$ operators $A_1, \dots, A_{|\mathcal{G}|} \in \mathbb{R}^{m \times n}$ w.r.t. the Lebesgue measure on $\mathbb{R}^{|\mathcal{G}|nm}$, the signal model \mathcal{X} can be uniquely identified from the measurement sets $\{A_g \mathcal{X}\}_{g=1}^{|\mathcal{G}|}$ if the number of measurements per operator verifies $m > k + n/|\mathcal{G}|$.*

The proof is included in Appendix B. It is worth noting that almost every set of $|\mathcal{G}|$ operators verifies the necessary condition in Proposition 3. If we have a large number of independent operators $|\mathcal{G}| \geq n$, Theorem 4 states that only $m > k + 1$ measurements are sufficient for model identification, which is slightly smaller (if the model is not trivial, i.e., $k > 1$) than the number of measurements typically needed for signal recovery $m > 2k$. In this case, it is possible to uniquely identify the model, without necessarily having a perfect reconstruction of each observed signal. However, as discussed in Section 2, for $k < m \leq 2k$, the subset of signals which cannot be uniquely recovered is for almost all A_g at most $(2k - m)$ -dimensional.

Operators with Different Number of Measurements The results of the previous subsections can be easily extended to the setting where each measurement operator has a different number of measurements, i.e., $A_1 \in \mathbb{R}^{m_1 \times n}, \dots, A_{|\mathcal{G}|} \in \mathbb{R}^{m_{|\mathcal{G}|} \times n}$. In this case, the necessary condition in Proposition 3 is $\sum_{g=1}^{|\mathcal{G}|} m_g \geq n$, and the sufficient condition in Theorem 4 is $\frac{1}{|\mathcal{G}|} \sum_{g=1}^{|\mathcal{G}|} m_g > k + n/|\mathcal{G}|$. As the proofs mirror the ones of Proposition 3 and Theorem 4, we leave the details to the reader.

4. Learning with a Single Operator and Group Symmetry

We begin with some notation on groups and their associated linear representations. A group \mathcal{G} with multiplication denoted by \cdot is a set of $|\mathcal{G}|$ elements g , such that the product of two elements is equal to another element in \mathcal{G} ($g_1 \cdot g_2 \in \mathcal{G}$), \mathcal{G} contains the identity element $1_{\mathcal{G}}$

and the inverse of every element is contained in \mathcal{G} , i.e., $g^{-1} \in \mathcal{G}$ and $g \cdot g^{-1} = 1_{\mathcal{G}}$. The action of a group in a vector space V (e.g., \mathbb{R}^n or \mathbb{C}^n) can be represented by invertible transformations T_g acting on V , where the \cdot operation is just matrix multiplication, and the transformations are compatible with the group multiplication, i.e., $T_g T_{g'} = T_{g \cdot g'}$.

In this section, we assume that we observe the measurement set $\mathcal{Y} = A\mathcal{X}$ defined as

$$\mathcal{Y} = \{y \in \mathbb{R}^m \mid y = Ax, x \in \mathcal{X}\} \quad (20)$$

and that the signal set \mathcal{X} is invariant to the action of a group \mathcal{G} , that is $T_g \mathcal{X} = \mathcal{X}$ for all transformations $T_1, \dots, T_{|\mathcal{G}|} \in \mathbb{R}^{n \times n}$. As with Section 3, in practice we will only observe a set of N training measurement data y_i via a single operator, i.e.,

$$y_i = Ax_i \quad (21)$$

for $i = 1, \dots, N$. We do not assume that the distribution $p(x)$ is also \mathcal{G} -invariant, which is a more stringent condition, although this setting is still covered here as a special case, since any \mathcal{G} -invariant $p(x)$ has a \mathcal{G} -invariant support. As with the multiple operator case, we first focus on the noiseless case, leaving the discussion about the effect of noise for Section 5.

The invariance property provides us with a simple but powerful way of learning the model with incomplete measurements (Chen et al., 2021): for all group elements the following holds

$$y_i = Ax_i \quad (22)$$

$$= AT_g T_g^{-1} x \quad (23)$$

$$= A_g x_{i'} \quad (24)$$

where both x_i and $x_{i'}$ belong to the signal set due to the invariance property. Equation (24) tells us that we can implicitly access $|\mathcal{G}|$ different measurement operators $A_g = AT_g$, each one associated with a nullspace \mathcal{N}_{A_g} which might differ from \mathcal{N}_A . Thus, we have a similar setting as in Section 3, with the additional constraint that the operators A_g are related by some transform $A_g = AT_g$, where the transforms $\{T_g\}_{g \in \mathcal{G}}$ can be seen as the action of some group \mathcal{G} . As the virtual forward operators A_g are related by the group action, we cannot directly apply the results in Section 3, which assume that the operators are independent, and we require further analysis. The following toy example provides intuition of how the invariance property can help in learning the model:

Example 3 (Shift invariant toy model) *Consider again the one-dimensional subspace model presented in Example 1. We now assume that the model is shift invariant, including all shifts of every signal, but we only partially observe signals through a fixed masking operator. While the original non-invariant set \mathcal{X}_0 is a one-dimensional linear subspace, the full set \mathcal{X} is now a union of linear one-dimensional subspaces $\bigcup_{g \in \mathcal{G}} T_g \mathcal{X}_0$. While \mathcal{X} is much larger than \mathcal{X}_0 , the intrinsic degrees of freedom are the same, i.e., the triangular wave ϕ . Despite having a single mask, we can access to multiple virtual operators AT_g via (24), as illustrated in Figure 5. Solving both model identification and signal recovery problem reduces to high rank (union of subspaces) matrix completion (Eriksson et al., 2012) of the measurement matrix $[y_1, \dots, y_N]$.*

Before delving into the main results of this section, we briefly introduce some fundamental concepts of linear representation theory and discuss some common group actions.




Figure 5: Learning a one-dimensional shift invariant model from masked observations. Each column shows a different observation from the signal set. The sensing mask is shown in red. Following (24), the samples can be reinterpreted as having measured another valid signal $x_{i'}$ through a different operator A_{g_i} .

4.1 Linear Representation Theory Preliminaries

A classical result in linear representation theory states that \mathbb{C}^n admits a unique decomposition into invariant¹² subspaces (Serre, 1977):

Theorem 5 (Decomposition of a representation, adapted from Serre (1977)) *Let \mathcal{G} be a compact group acting on \mathbb{C}^n , whose action is represented by the invertible linear mappings T_g . There is a unique canonical decomposition of \mathbb{C}^n into J invariant subspaces as*

$$\mathbb{C}^n = V_1 \oplus \cdots \oplus V_J \quad (25)$$

where each subspace has dimension $s_j c_j$ with s_j the dimension of the j th irreducible representation and c_j its multiplicity. Moreover, any linear representation admits a block-diagonal form

$$T_g = F \Sigma_g F^{-1} \quad (26)$$

where F is a basis for (25) and

$$\Sigma_g = \begin{bmatrix} \Lambda_1(g) & & \\ & \ddots & \\ & & \Lambda_J(g) \end{bmatrix} \quad (27)$$

with

$$\Lambda_j(g) = \begin{bmatrix} \rho_j(g) & & \\ & \ddots & \\ & & \rho_j(g) \end{bmatrix} \in \mathbb{C}^{c_j s_j \times c_j s_j} \text{ and } \rho_j(g) \in \mathbb{C}^{s_j \times s_j}. \quad (28)$$

The matrices $\rho_j(g)$ are unique (independent of the basis F) and correspond to the j th irreducible representation of \mathcal{G} .

In words, a group is linked to a unique set of J irreducible representations of dimension s_1, \dots, s_J , which serve as fundamental building blocks of any linear group action. A linear

¹². An invariant subspace is a subspace of \mathbb{C}^n which is invariant to all transformations $\{T_g\}_{g \in \mathcal{G}}$.

group action on \mathbb{C}^n consists of a specific basis F and a set of multiplicities c_j such that $\sum_{j=1}^J s_j c_j = n$. The vector space on which the group acts admits a decomposition into J invariant subspaces of dimension $s_1 c_1, \dots, s_J c_J$, each associated with a different irreducible representation. To build some intuition about Theorem 5, we discuss its implications for common group actions such as shifts, reflection, rotations and permutations on 1D (e.g., audio) or 2D discretely sampled signals (e.g., images):

4.1.1 FINITE GROUPS

Shifts The cyclic group of n elements associated with circulant matrices T_{g_i} which shift the signal by i taps/pixels, with $i \in \{0, \dots, n-1\}$. In this case, we have $|\mathcal{G}| = J = n$ and $c_1 = \dots = c_J = 1$ and $s_j = 1$, where the subspaces are the Fourier modes. Here, F is the discrete Fourier transform and Λ is a diagonal matrix containing the Fourier transform of the discrete shift operation.

Reflection The cyclic group of 2 elements, the identity and a reflection of a 1D signal about its central entry ($|\mathcal{G}| = 2$). For even¹³ n , we have $J = 2$, $s_1 = 1$ with $c_1 = n/2$ composed of vectors $[v_1, \dots, v_{n/2}, v_{n/2}, \dots, v_1]^\top$ with $v \in \mathbb{C}^{n/2}$, and $s_2 = 1$ with $c_2 = n/2$, composed of vectors $[v_1, \dots, v_{n/2}, -v_{n/2}, \dots, -v_1]^\top$. The matrix F can be built as

$$\frac{1}{\sqrt{2}} \begin{bmatrix} \tilde{F} & \tilde{F} \\ T_1 \tilde{F} & -T_1 \tilde{F} \end{bmatrix} \quad (29)$$

where $T_1 \in \mathbb{R}^{n/2 \times n/2}$ is a reflection matrix with entries $[T_1]_{i,j} = 1$ if $i + j = n/2 + 1$ and 0 otherwise, and $\tilde{F} \in \mathbb{C}^{n/2 \times n/2}$ is any unitary matrix.

Rotations To simplify the exposition, here we consider a 2D circular sensing pattern with n_1 pixels of diameter and angle sampled every $360/n_2$ degrees (hence $n = n_1 n_2$). The cyclic group \mathcal{G} is then represented as the set of all $|\mathcal{G}| = n_2$ possible rotations. In this case, we have $J = n_2$, $s_1 = \dots = s_J = 1$ and $c_1 = \dots = c_J = n_1$. The matrix F is given by the Kronecker product between a 1D discrete Fourier transform along the angular dimension $F_1 \in \mathbb{C}^{n_2 \times n_2}$ and any unitary matrix along the radial dimension $F_2 \in \mathbb{C}^{n_1 \times n_1}$, i.e., $F_1 \otimes F_2$.

Permutations Full set of permutations of the n entries of the signal with $|\mathcal{G}| = n!$ elements. This group can model permutations of a masking operator, exchangeable distributions (Eaton, 1989) or patches in an image. Commonly-used separable distributions (Guo and Davies, 2015) are also permutation invariant.. This is a significantly larger group (it contains all the rest as subgroups), where $J = 2$, $s_1 = c_1 = 1$ for the subspace spanned by $[1, \dots, 1]^\top$ and its complement $s_2 = n - 1$ with $c_2 = 1$. Here, F is given by any unitary matrix whose first column is the constant vector $[\frac{1}{\sqrt{n}}, \dots, \frac{1}{\sqrt{n}}]^\top$.

4.1.2 INFINITE GROUPS

Natural signals are often described as continuous functions $x : S \mapsto \mathbb{R}^p$, where S is a compact domain and \mathbb{R}^p is the range. For example, an RGB image can be defined using

13. The extension to odd n is trivial.

$S = (0, 1]^2$ and $p = 3$. In this description, the ambient space is generally an infinite-dimensional Hilbert space \mathcal{H} and group actions are generally continuous, such as translations or arbitrary rotations.

Formally, the linear representation of a group acting on \mathcal{H} is given by a mapping from \mathcal{G} to the space of unitary linear mappings $T_g : \mathcal{H} \mapsto \mathcal{H}$. The decomposition into invariant subspaces of Theorem 5 also holds in this setting, where the number of irreducible representations with non-zero multiplicity J can be infinite (Serre, 1977, Chapter 4).

In practice, we generally assume that the signals of interest are bandlimited and, according to Nyquist theorem, can be represented on a discrete grid associated with \mathbb{R}^n . Restricted to the bandlimited subspace, the group action¹⁴ has a finite number of non-zero multiplicities c_j and Theorem 5 holds despite the infinite number of group elements ($|\mathcal{G}| = \infty$). As Nyquist theorem guarantees a one-to-one mapping between the signal model in \mathbb{R}^n and the signal model defined in \mathcal{H} , we only need to prove model identifiability using signals and (possibly continuous) group actions on \mathbb{R}^n in order to prove model identifiability in the continuous case. We illustrate these concepts with two simple examples:

Translations Consider the space $L_2(0, 1]$ of square integrable signals taking values in the $(0, 1]$ interval where the group of translations is represented via the left action

$$T_g x = x(t - g) \quad (30)$$

with $g \in (0, 1]$, where $t - g$ is modulo 1. This representation has $s_j = c_j = 1$ for all $j \in \mathbb{Z}$, where each invariant subspace corresponds to the complex exponential $e^{-i2\pi jt}$. Assuming that the signals of interest have a maximum frequency $J_{\max} \in \mathbb{N}$, we can represent them on a grid as a vector $\tilde{x} \in \mathbb{R}^n$, $n = 2J_{\max} + 1$, with entries given by

$$\tilde{x}(r) = \int_0^1 x(t) k(t - \frac{r}{n}) dt \quad (31)$$

for $r = 1, \dots, n$ where $k(t) = \frac{\sin(\pi t n)}{\sin(\pi t)}$ is the periodic sinc kernel (the ideal¹⁵ anti-aliasing filter). On \mathbb{R}^n , the group action is given by diagonal matrices

$$\tilde{T}_g = F \begin{bmatrix} e^{-i2\pi j_1 g} & & & \\ & \ddots & & \\ & & \ddots & \\ & & & e^{-i2\pi j_n g} \end{bmatrix} F^{-1} \quad (32)$$

where $j_1 < \dots < j_n$ are the integers in $[-J_{\max}, J_{\max}]$, F is the discrete Fourier transform and $g \in (0, 1]$. This representation has $s_j = c_j = 1$ for all $j = 1, \dots, n$. It is worth noting that, restricted to a subgroup of n equispaced elements, this group action is equivalent to the discrete shifts example in Section 4.1.1.

14. If an antialiasing filter is used, the discretization process is an equivariant mapping, i.e., there is a valid representation of \mathcal{G} acting on \mathbb{R}^n .

15. In practice, images are usually sampled on a grid and non ideal anti-aliasing filters are used hence there is only approximate equivariance.

Continuous Rotations Consider the set of continuous signals whose domain S is the unit circle and the range is \mathbb{R} . The signals can be written as $x(r, \theta)$ with radius $r \in (0, 1]$ and angle $\theta \in (0, 2\pi]$. The left action of the group of rotations is given by

$$T_g x = x(r, \theta - g) \quad (33)$$

with $g \in (0, 2\pi]$, where $\theta - g$ is modulo 2π . Assuming that the signals of interest have a maximum frequency J_{\max_1} , $n_1 = 2J_{\max_1} + 1$, on the radial axis and a maximum frequency J_{\max_2} , $n_2 = 2J_{\max_2} + 1$, on the angular axis, we can represent them as vectors $\tilde{x} \in \mathbb{R}^{n_1 n_2}$ on a discrete grid by the convolution with the periodic sinc kernel every $\Delta r = 1/n_1$ and $\Delta\theta = 2\pi/n_2$. The action of rotations in $\mathbb{R}^{n_1 n_2}$ has $s_1 = \dots = s_J = 1$ and $c_1 = \dots = c_J = n_1$. Restricted to the subgroup of rotations by $360/n_2$ degrees, the representation is equivalent to the discrete rotations example in Section 4.1.1.

4.2 Uniqueness of Any \mathcal{G} -Invariant Model?

As with the case of multiple independent operators, we first ask the question of whether it is possible to uniquely identify any \mathcal{G} -invariant model from measurements of a single operator with $m < n$, regardless of the dimension of its support \mathcal{X} . We show that even under the assumption that $p(x)$ is \mathcal{G} -invariant, uniqueness is not possible if the support has dimension close to the ambient dimension. If $p(x)$ is \mathcal{G} -invariant, all points in the same orbit, \mathcal{O}_x , defined as

$$\mathcal{O}_x = \{T_g x \mid x \in \mathbb{C}^n, \forall g \in \mathcal{G}\} \quad (34)$$

will necessarily have the same probability. Hence, the distribution is uniquely identified by the distribution over the orbits \mathcal{O}_x , instead of the whole space \mathbb{C}^n . This notion is made precise by a fundamental theorem of invariant statistics:

Theorem 6 (Decomposable measures, adapted from Eaton (1989)) *A \mathcal{G} -invariant distribution on \mathbb{C}^n can be decomposed into a uniform distribution over the elements in \mathcal{G} and a distribution over the quotient space \mathbb{C}^n/\mathcal{G} .*

The theorem tells us that we can potentially design a measurement operator with $m < n$ which is one-to-one with respect to the quotient space (having model uniqueness) but not necessarily one-to-one with respect to the signal space (failing to have signal recovery). For infinite compact groups, the quotient space might be significantly smaller than the full signal space. For example, consider distributions which are invariant to the action of the orthogonal group O_n composed of all orthogonal matrices. These are necessarily isotropic distributions only dependent on the radius $\|x\|$. In this case, the full space is n -dimensional, but the quotient space is only one-dimensional (the sufficient statistics are functions of $\|x\|$), and (non-linear) one-dimensional measurements $A(x) = \|x\|$ are sufficient to identify any \mathcal{G} -invariant distribution. On the other hand, the orbits of finite groups¹⁶, which are the main focus of this paper, are sets consisting of a finite number of signals, and the quotient space is intrinsically n -dimensional:

16. A similar statement holds for smaller infinite and compact groups, such as continuous translations of 1D or 2D signals.

Theorem 7 (Dimension of quotient space (Sturmfels, 2008)) *The quotient space \mathbb{C}^n/\mathcal{G} of any finite group acting on \mathbb{C}^n is n -dimensional.*

As \mathbb{C}^n/\mathcal{G} is n -dimensional for any finite group, any forward operator $A \in \mathbb{C}^{m \times n}$ with $m < n$ cannot be one-to-one on \mathbb{C}^n/\mathcal{G} , hence some invariant distributions will not be unique. Despite not being able to learn any \mathcal{G} -invariant model, we show that, as in Section 3, uniqueness among low-dimensional \mathcal{G} -invariant models is indeed possible.

4.3 Low-Dimensional \mathcal{G} -Invariant Models

The observation in (24) tells us that \mathcal{G} -invariant models provide access to $|\mathcal{G}|$ (virtual) operators AT_g with $g \in \mathcal{G}$. Thus, we can apply similar ideas to those used in Theorem 4 to analyze the problem of learning a \mathcal{G} -invariant model, with the additional constraint that the operators are \mathcal{G} -related. Instead of depending on the ratio between ambient dimension n and the number of transformations $|\mathcal{G}|$, the bounds depend on the largest ratio of multiplicity and dimension of the irreducible components of the group action (as defined in Theorem 5), that is

$$\max_j \frac{c_j}{s_j}. \tag{35}$$

With this definition in mind, we begin by reformulating the necessary condition in Proposition 3:

Proposition 8 *Let \mathcal{G} be a compact group. A necessary condition for model uniqueness with operators $\{AT_g\}_{g \in \mathcal{G}}$ is $m \geq \max_j \frac{c_j}{s_j}$.*

The proof is detailed in Appendix C, and similarly to the proof of Proposition 3, requires analyzing the rank of the matrix (or its infinite dimensional equivalent for infinite groups)

$$\begin{bmatrix} A_1 \\ \vdots \\ A_{|\mathcal{G}|} \end{bmatrix} \tag{36}$$

with the additional constraint that $A_g = AT_g$. Note that for any group action $n/|\mathcal{G}| \leq \max_j c_j/s_j \leq n$, hence this is always a more restrictive bound than the one in Proposition 3. The following two examples present group actions where $\max_j c_j/s_j \gg n/|\mathcal{G}|$:

Example 4 *Consider the group of permutations of the last 4 elements of a vector of 10 elements where $|\mathcal{G}| = 4! = 24$. The group action has $\max_j c_j/s_j = 6$ associated with the subspace containing the first 6 elements, whereas $n/|\mathcal{G}| < 1$.*

Example 5 *Consider the infinite group of rotations of a signal introduced in Section 4.1.2. The group action has $\max_j c_j/s_j = n_2$, however $|\mathcal{G}| = \infty$.*

We now study a sufficient condition for unique model identification. As we have that $\max_j c_j/s_j \geq n/|\mathcal{G}|$, one might be tempted to extrapolate the bound in Theorem 4, and expect that $m > k + \max_j c_j/s_j$ measurements are sufficient for model identification. However, the following counter-example shows that this is not always sufficient:

Example 6 Let \mathcal{G} be the group of reflection with 2 elements $\{e, r\}$ of a signal with even length n (see Section 4.1.1). Consider the representation given by $T_e = I_n$ where I_n is the $n \times n$ identity matrix and

$$T_r = \begin{bmatrix} I_{n/2} & 0 \\ 0 & -I_{n/2} \end{bmatrix}. \quad (37)$$

Note that the representation has $J = 2$ invariant subspaces with $s_1 = s_2 = 1$ and $c_1 = c_2 = n/2$. Let \mathcal{X} be the set of d -sparse signals where $n > 2d$. This conic set has dimension $k = d - 1$ according to the definition in Section 3.2. In order to have model uniqueness, we require that the inferred signal set $\hat{\mathcal{X}}$ defined in (17) equals the true set \mathcal{X} , or equivalently that their difference

$$\hat{\mathcal{X}} \setminus \mathcal{X} = \{v \in \mathbb{R}^n \setminus \mathcal{X} \mid AT_e(x_1 - v) = AT_r(x_2 - v) = 0, x_1, x_2 \in \mathcal{X}\} \quad (38)$$

is empty. This condition can be written in matrix form as

$$A[\tilde{x}_1 - T_e v, \tilde{x}_2 - T_r v] \neq 0 \quad (39)$$

where $\tilde{x}_1 = T_e x_1$ and $\tilde{x}_2 = T_r x_2$ are d -sparse and thus also belong to the signal set \mathcal{X} . Let $\Phi_z = [\tilde{x}_1 - T_e v, \tilde{x}_2 - T_r v]$, such that condition (39) can be written as $A\Phi_z \neq 0$. We consider the (extreme) case where \tilde{x}_1 and \tilde{x}_2 have disjoint supports, i.e. $\text{supp}(\tilde{x}_1) = [n/2+1, \dots, n/2+d]$ and $\text{supp}(\tilde{x}_2) = [n/2+d+1, \dots, n/2+2d]$. Using coordinates $[a_1, \dots, a_d]^\top \in \mathbb{R}^d$ for \tilde{x}_1 and $[b_1, \dots, b_d]^\top \in \mathbb{R}^d$ for \tilde{x}_2 , we obtain

$$[\tilde{x}_1 - T_e v, \tilde{x}_2 - T_r v] = \begin{bmatrix} -v_1 & -v_1 \\ \vdots & \vdots \\ -v_{n/2} & -v_{n/2} \\ a_1 - v_{n/2+1} & v_{n/2+1} \\ \vdots & \vdots \\ a_d - v_{n/2+d} & v_{n/2+d} \\ -v_{n/2+d+1} & b_1 + v_{n/2+d+1} \\ \vdots & \vdots \\ -v_{n/2+2d} & b_d + v_{n/2+2d} \\ \vdots & \vdots \\ -v_n & v_n \end{bmatrix}. \quad (40)$$

Setting $a_i = 2v_{n/2+i}$, $b_i = -2v_{n/2+d+i}$ for $i = 1, \dots, d$ and $v_j = 0$ for $j = n/2+2d+1, \dots, n$, we have that $[\tilde{x}_1 - T_e v, \tilde{x}_2 - T_r v] = [\mu, \mu]$ with $\mu = -[v_1, \dots, v_{n/2+2d}, 0, \dots, 0]^\top$. Hence, condition (39) can be simplified to

$$\tilde{A} \begin{bmatrix} v_1 \\ \vdots \\ v_{n/2+2d} \end{bmatrix} \neq 0 \quad (41)$$

where \tilde{A} is a $m \times (n/2 + 2d)$ submatrix of A . Thus, the condition in (41) can only hold for all $[v_1, \dots, v_{n/2+2d}]^\top \in \mathbb{R}^{n/2+2d} - \{0\}$ if $m \geq n/2 + 2d$. As this linear representation has $\max_j c_j/s_j = n/2$, and $k = d - 1$, we have

$$m \geq \max_j c_j/s_j + 2k - 2. \quad (42)$$

Our main theorem shows that, for a cyclic group, $2k + 1$ additional measurements are actually sufficient to guarantee model uniqueness for almost every A :

Theorem 9 *Let \mathcal{G} be a compact cyclic group acting on \mathbb{R}^n and assume the signal set \mathcal{X} verifies assumption **A1**. For almost every $A \in \mathbb{R}^{m \times n}$ w.r.t. to the Lebesgue measure on \mathbb{R}^{mn} , the signal set can be uniquely identified from the sets $\{AT_g \mathcal{X}\}_{g \in \mathcal{G}}$ if the number of measurements verifies $m > 2k + \max_j c_j + 1$. In particular, for almost any operator $A \in \mathbb{R}^{m \times n}$, a \mathcal{G} -invariant signal set can be identified from the measurement set $\mathcal{Y} = A\mathcal{X}$ if $m > 2k + \max_j c_j + 1$.*

The proof is detailed in Appendix C. As all the irreducibles of a cyclic group have dimension $s_j = 1$ (Serre, 1977), the bound can also be expressed as $m > 2k + \max_j c_j/s_j + 1$. Unlike the case of multiple independent operators, having enough measurements to guarantee model identification also guarantees unique signal recovery due to the additional k measurements. However, the experiments in Section 7 suggest that $m > 2k + \max_j c_j/s_j$ measurements are only required in worst-case scenarios such as Example 6, whereas $m > k + \max_j c_j/s_j$ are required for typical models.

Although Theorem 9 is specific to actions of cyclic groups, it can be applied to any group action, as all groups contain at least one cyclic subgroup. For example, if the group action consists of the composition of rotations and reflection, which is not cyclic nor abelian, we can apply Theorem 9 by only considering rotations or reflection, which are cyclic when considered separately. However, if the cyclic subgroup is significantly smaller than the full group, Theorem 9 might not be tight. We conjecture that the bound in Theorem 9 holds for any compact group:

Conjecture 10 *Let \mathcal{G} be any compact group acting on \mathbb{R}^n and assume the signal set \mathcal{X} verifies assumption **A1**. For almost every $A \in \mathbb{R}^{m \times n}$ w.r.t. to the Lebesgue measure on \mathbb{R}^{mn} , the signal set can be uniquely identified from the sets $\{AT_g \mathcal{X}\}_{g \in \mathcal{G}}$ if the number of measurements verifies $m > 2k + \max_j c_j/s_j + 1$. In particular, for almost any operator $A \in \mathbb{R}^{m \times n}$, a \mathcal{G} -invariant signal set can be identified from the measurement set $\mathcal{Y} = A\mathcal{X}$ if $m > 2k + \max_j c_j/s_j + 1$.*

The *almost every* property can be hard to interpret in practical scenarios. A simple example of matrices that will hold this property with probability 1 are compressive sensing matrices (e.g., iid random Gaussian entries). However, there is an important subset of measurement operators of Lebesgue measure zero in $\mathbb{R}^{m \times n}$ that do not verify this property and are treated in the next subsection.

4.4 Equivariant Measurement Operators

In order to obtain additional information in the nullspace of A , Proposition 3 requires that the measurements operators $AT_1, \dots, AT_{|\mathcal{G}|}$ do not share the same nullspace. Operators which fail to bring information in the nullspace of A possess the property of being \mathcal{G} -equivariant:




Figure 6: Identification of models with continuous signals and continuous group actions. Continuous signals are well described as points in an infinite dimensional Hilbert space \mathcal{H} , and symmetries in the model are described as a continuous group acting on \mathcal{H} . As discussed in Section 4.1.2, if the signals are bandlimited, we can represent them on a discrete grid \mathbb{R}^n . Nyquist theorem ensures this mapping is equivariant and one-to-one, thus we have a valid group action on \mathbb{R}^n . Theorem 9 guarantees that we can uniquely identify the model from measurement data alone via almost every (thus non-equivariant) mapping A with $m > 2k + \max_j c_j + 1$ measurements if $|\mathcal{G}|$ is a cyclic group.

Theorem 11 *The set of forward operators $\{AT_g\}_{g \in \mathcal{G}}$ share the same nullspace for all $g \in \mathcal{G}$ if and only if A is an equivariant map, i.e., it verifies*

$$AT_g = \tilde{T}_g A \quad (43)$$

for all $g \in \mathcal{G}$, where $\tilde{T}_g : \mathcal{G} \mapsto \mathbb{C}^{m \times m}$ is a linear representation of \mathcal{G} acting on \mathbb{C}^m . Moreover, any equivariant linear operator can be decomposed as $A = \tilde{F} \Lambda F^{-1}$, where F is the basis associated to the group representation in Theorem 5, \tilde{F} is any basis of \mathbb{C}^m and $\Lambda \in \mathbb{C}^{m \times n}$ has the following block-diagonal structure

$$\Lambda = \begin{bmatrix} \Lambda_1 & & \\ & \ddots & \\ & & \Lambda_J \end{bmatrix} \text{ with } \Lambda_j = \begin{bmatrix} B_j & & \\ & \ddots & \\ & & B_j \end{bmatrix} \in \mathbb{C}^{\tilde{c}_j s_j \times c_j s_j} \text{ and } B_j \in \mathbb{C}^{\tilde{c}_j \times c_j} \quad (44)$$

where \tilde{c}_j are the multiplicities of the representation \tilde{T}_g , such that $m = \sum_{j=1}^J s_j \tilde{c}_j$.

The proof is detailed in Appendix C. This necessary condition on A has important practical implications, as equivariant operators appear often in real-world settings which are discussed in next subsection.

We end this subsection with a remark concerning continuous signals and groups actions (see Figure 6) defined in an infinite-dimensional Hilbert space. As discussed in Section 4.1.2, bandlimited signals x can be represented in \mathbb{R}^n on a discrete grid after an anti-aliasing filter. This procedure is an equivariant mapping, as we have a valid linear representation of the group acting on \mathbb{R}^n . Contrary to the grid sensing step, the measurement process A should be non-equivariant in order to be able to uniquely identify the model from measurement data alone.

4.5 Consequences for Common Group Actions

We discuss the implications of Theorem 9 and Theorem 11 for the group actions introduced in Section 4.1.




Figure 7: Forward operators consisting of partial Fourier measurements. **(a)** The Radon transform appearing in tomography problems consists of lines in Fourier space. **(b)** Spiral frequency strategies are commonly used in MRI. **(c)** Circular Fourier pattern of measurements. Shift invariance does not provide any information in the nullspace. Rotation invariance provides additional information in both **(a)** and **(b)** but does not provide additional information in **(c)**.

Translations/Shifts Continuous translations and shifts have the irreducible representations with non-zero multiplicities. A generic sufficient condition for model identification in both cases is $m > 2k + 2$. Unique model identification is impossible with measurement operators consisting of a subset of Fourier basis vectors which are shift equivariant. This is the case of multiple inverse problems such as deblurring, super-resolution, accelerated MRI and limited angle CT. However, it is possible to uniquely identify the signal model using translation invariance in other inverse problems such as image inpainting or compressed sensing.

Reflection As this is a small group, in order to guarantee model identification our theory requires $m > 2k + n/2 + 1$ measurements, significantly more than those needed for signal recovery.

Rotations If we consider signals defined on the unit circle and a 2D radial sensing pattern with n_1 pixels of diameter and angle sampled every $360/n_2$ degrees (hence $n = n_1 n_2$), we have $\max_j c_j/s_j = n_1$ and thus need $m > 2k + n_1 + 1$ for model identification, which is also more measurements than required for signal recovery.

General rotations of a Cartesian grid require some degree of interpolation for off-of-grid values. Thus, this set of transformations do not form an exact group action in general. However, if the signals are sufficiently band-limited and supported within a radius from the center, the Cartesian grid can be approximated by the radial grid, and thus we need $m > 2k + n_1 + 1$ measurements.

Rotations are useful for learning the signal model in inverse problems which depend on Fourier measurements, as long as these measurements are not of circular stripes of Fourier space, as illustrated in Figure 7. Two important examples where learning is possible are limited angle CT (Chen et al., 2021) and accelerated MRI (Chen et al., 2022), as the Fourier patterns are generally not rotationally invariant.

Rotations by 90 degrees and Reflection Consider the non-commutative group of 8 elements whose action is composed of 90 degree rotations and reflection of a 2D signal defined on a regular grid. This group action has $\max_j c_j/s_j \approx n/8$ for large n . We conjecture

that $m > 2k + n/8 + 1$ measurements are sufficient for model identification, significantly more than the case of rotations of a radial grid.

Permutations Group of all permutations of n entries which has $\max_j c_j/s_j = 1$. Note that the previous examples are subgroups of this group. In this case, signal recovery also guarantees model identification. The only equivariant measurement operator which does not provide additional information is the constant vector $A = \alpha[1, \dots, 1]^\top$ with $\alpha \in \mathbb{R}$.

5. Noisy Measurement Data

Surprisingly, the results of this paper are also theoretically valid if the measurements are corrupted by independent additive noise ϵ , i.e., $y = A_g x + \epsilon$, as long as the noise distribution is *known* and has a nowhere zero characteristic function (e.g., Gaussian noise):

Proposition 12 *For a fixed noise distribution, if its characteristic function is nowhere zero, then there is a one-to-one mapping between the space of clean measurement distributions and noise measurement distributions.*

Proof Consider the noisy measurements associated to the g th operator A_g , as $z = y + \epsilon$, where z are the observed noisy measurements, y are the clean measurements and ϵ is additive noise (independent of y). The characteristic function of the sum of two independent random variables is given by the multiplication of their characteristic functions, i.e.,

$$\varphi_z(w) = \varphi_y(w)\varphi_\epsilon(w) \quad (45)$$

where φ_z , φ_y and φ_ϵ are the characteristic functions the noisy measurement, clean measurements and noise distributions, respectively. If the characteristic function of the noise distribution is nowhere zero, we can uniquely identify the characteristic function of the clean measurement distribution as

$$\varphi_y(w) = \varphi_z(w)/\varphi_\epsilon(w). \quad (46)$$

The clean measurement distribution is fully characterized by its characteristic function $\varphi_y(w)$. We end the proof by noting that the same reasoning applies to the measurements of every operator A_g with $g \in \{1, \dots, |\mathcal{G}|\}$. ■

If the clean measurement distribution can be uniquely identified, we can then apply Theorem 4 and Proposition 3 for the case of independent operators, and Theorems 9 and 11 and Proposition 8 for the case of \mathcal{G} -related operators. Note that this only guarantees model identifiability and makes no claims on the sample complexity of any learning process.

6. Algorithms

In this section, we summarize some different approaches for solving model identification and/or signal recovery problems. We recall the learning settings:

Multiple operators We observe N measurement vectors y_i , each associated with one out of $|\mathcal{G}|$ operators, $y_i = A_{g_i} x_i + \epsilon_i$, where A_{g_i} is known.

Single operator, \mathcal{G} -invariant model We observe N measurement vectors y_i , associated with a single (known) operator, $y_i = Ax_i + \epsilon_i$, and assume that the set of signals \mathcal{X} is \mathcal{G} -invariant.

6.1 Learning a Generative Model

One approach is to first learn the low-dimensional support \mathcal{X} and then reconstruct the signals. This can be done in practice by estimating $\hat{\mathcal{X}}$ within a class of low-dimensional models \mathcal{C} while satisfying measurement consistency:

$$\arg \min_{\hat{\mathcal{X}}, x_1, \dots, x_N} \sum_{i=1}^N \|y_i - A_{g_i} x_i\|^2 \tag{47}$$

$$\text{s.t. } \hat{\mathcal{X}} \in \mathcal{C}, x_i \in \hat{\mathcal{X}}. \tag{48}$$

For example, low-rank matrix completion algorithms (Cai et al., 2010; Candès and Recht, 2009) assume that the signal set (columns of the matrix) consists of a k -dimensional linear subspace. Sparse subspace clustering with missing entries methods (Yang et al., 2015a; Ongie et al., 2017) assume that the signal set is a union of low-dimensional subspaces. These approaches are closely related to k -sparse dictionary learning approaches (Studer and Baraniuk, 2012), which can also be trained using only incomplete measurements. Union of subspace models are also closely related to a mixture of Gaussians with low-rank covariance matrices, which can also be used to learn from incomplete measurements (Yang et al., 2015b).

In some inverse problems, a direct inversion $y \mapsto x$ might be hard to obtain due to a high level of noise affecting the measurements. In this setting, it might be possible to identify the signal distribution from incomplete data, without recovering each individual signal. For example, the signal distribution $p(x)$ can be modelled with a generative deep network. This strategy can be applied in learning problems with independent operators, \mathcal{G} -structured operators or \mathcal{G} -invariant signal sets with a single operator.

The case of multiple operators is summarized as follows. Let $q_g^{\text{data}}(y)$ be the empirical measurement distribution of observations linked to the g th operator. The learning strategy can be written as

$$\arg \min_{\hat{p}(x)} \sum_{g=1}^{|\mathcal{G}|} d(\hat{q}_g(y), q_g^{\text{data}}(y)) \tag{49}$$

$$\text{s.t. } \hat{q}_g(y) = \int \ell(y|A_g x) \hat{p}(x) dx \quad g = 1, \dots, |\mathcal{G}| \tag{50}$$

where $\ell(y|A_g x)$ denotes the distribution of measurements given $A_g x$ (i.e., the likelihood), and $d(\cdot, \cdot)$ is a (pseudo) divergence, which is also generally learned using generative adversarial networks (GANs). This approach has been adopted by AmbientGAN (Bora et al., 2018) for various imaging problems.

The case of a single operator and a \mathcal{G} -invariant model follows a similar learning strategy:

$$\arg \min_{\hat{p}(x)} d(\hat{q}(y), q^{\text{data}}(y)) \quad (51)$$

$$\text{s.t. } \hat{q}(y) = \int \ell(y|Ax)\hat{p}(x)dx \quad (52)$$

$$\hat{p}(x) \text{ is } \mathcal{G}\text{-invariant} \quad (53)$$

where q^{data} denotes the empirical distribution of the observations. This approach has been adopted by Cryo-GAN (Gupta et al., 2020) for the cryo electron-microscopy problem.

6.2 Learning to Invert

If we observe signals through multiple forward operators, an alternative approach consists of directly learning the reconstruction function $f : (g, y) \mapsto x$ without explicitly learning the signal model. To this end, we can use the following training loss:

$$\arg \min_{f \in \mathcal{F}} \sum_{i=1}^N \|y_i - A_{g_i} f(g_i, y_i)\|^2. \quad (54)$$

The family of functions \mathcal{F} is some parameterized function space which captures the low-dimensionality of the signal set \mathcal{X} (e.g., a deep neural network with an autoencoder architecture). Without the low-dimensional constraint on \mathcal{F} , we can have $f(g, y) = A_g^\dagger y$ which is a minimizer of (54) but fails to learn the signal set. For example, we can choose $f(g, y) = \phi(A_g^\dagger y)$ where $\phi(\cdot)$ plays the role of a denoiser function (independent of the forward operator) which projects $A_g^\dagger y$ to a low-dimensional set $\hat{\mathcal{X}}$. More generally, the network architecture can be an unrolled optimization algorithm (Monga et al., 2021) which incorporates the forward operator A_g and a denoiser subnetwork which is operator independent, e.g., acting as a (low-dimensional) proximal operator in the unrolled algorithm.

In cases where the model is assumed to be \mathcal{G} -invariant and A is not itself equivariant (c.f., Theorem 11), the reconstruction function $f : y \mapsto x$ can be learned by enforcing that the composition of the forward operator and reconstruction $f \circ A$ yields a \mathcal{G} -equivariant mapping on \mathcal{X} . To this end, we can enforce approximate equivariance via the following training loss:

$$\arg \min_{f \in \mathcal{F}} \sum_{i=1}^N \|y_i - Af(y_i)\|^2 \quad (55)$$

$$\text{s.t. } f(AT_{g_i}\tilde{x}_i) = T_{g_i}\tilde{x}_i \quad \forall \tilde{x}_i = f(y_i). \quad (56)$$

The additional equivariance constraint in (56) stops the network from learning the identity mapping, since $I \circ A = A$ is not equivariant. This strategy was proposed in the equivariant imaging framework (Chen et al., 2021) and a similar approach is presented in Tachella et al. (2022) for the case of multiple forward operators. In Chen et al. (2021), the authors also consider adding an adversarial network to try to impose similarity between the estimated sample signal model $\hat{\mathcal{X}}$ and its transforms $\{T_g \hat{\mathcal{X}}\}_{g \in \mathcal{G}}$. The equivariant imaging loss in (55) can be extended to handle noisy measurements (Chen et al., 2022) using Stein’s unbiased risk estimator (Stein, 1981).




Figure 8: Reconstruction probability of a k -dimensional subspace using incomplete measurements arising from $|\mathcal{G}|$ independent operators for different k . The curve in red shows the bound of Theorem 4, $m > k + n/|\mathcal{G}|$.

7. Experiments

We perform a series of numerical experiments to illustrate the theoretical bounds presented in Sections 3 and 4.

7.1 Low-Dimensional Subspace Model

We consider the problem of learning a k -dimensional subspace model from incomplete observations, where the signals x_i are generated using a standard Gaussian distribution on the subspace. We first evaluate the case where each observation y_i is obtained by randomly choosing one out of $|\mathcal{G}|$ operators $A_1, \dots, A_{|\mathcal{G}|} \in \mathbb{R}^{m \times n}$, each composed of iid Gaussian entries of mean 0 and standard deviation $n^{-1/2}$. In order to recover the signal matrix $X = [x_1, \dots, x_N]$, we solve the following matrix completion problem

$$\begin{aligned} & \arg \min_X \|X\|_* \\ & \text{s.t. } A_{g_i} x_i = y_i \quad \forall i = 1, \dots, N \end{aligned} \tag{57}$$

where $\|\cdot\|_*$ denotes the nuclear norm. A recovery is considered successful if $\frac{\sum_i \|\hat{x}_i - x_i\|^2}{\sum_i \|x_i\|^2} < 10^{-1}$, where \hat{x}_i is the estimated signal for the i th sample. We use a standard matrix completion algorithm (Cai et al., 2010) to solve (57). The ambient dimension is fixed at $n = 50$, and the experiment is repeated for $k = 1, 10, 40$. For each experiment we set $N = 150k$ in order to have enough samples to estimate the subspaces. Figure 8 shows the probability of recovery over 25 Monte Carlo trials for different numbers of measurements m and operators $|\mathcal{G}|$. The reconstruction probability exhibits a sharp transition which follows the generic sufficient condition $m > k + n/|\mathcal{G}|$ of Theorem 4.

Second, we perform the same experiment using \mathcal{G} -related operators $A_g = AT_g$. Figure 9a shows the probability of correct recovery for the group of shifts, where $\max_j c_j/s_j = 1$, whereas Figure 9b shows the recovery probability for the reflection group where $\max_j c_j/s_j = n/2$. In both cases, the transition in probability of recovery follows $m \approx k + \max c_j/s_j$, i.e., requiring approximately k less measurements than the sufficient condition of Theorem 9. As discussed in Section 4.3, it is likely that these additional k measurements are only required in worst-case scenarios such as Example 6.

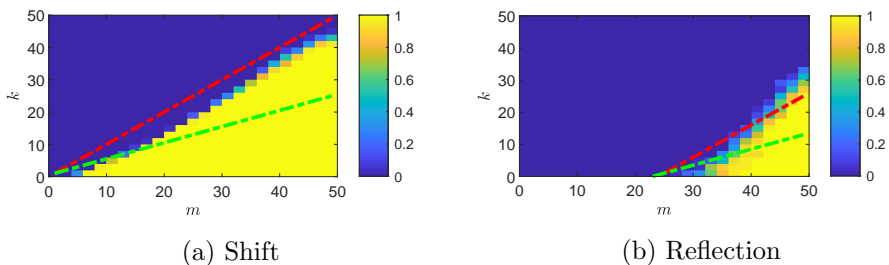


Figure 9: Reconstruction probability of a k -dimensional subspace using incomplete measurements arising from (a) shifted and (b) reflected operators AT_g . The bounds $m > k + \max_j c_j/s_j + 1$ and $m > 2k + \max_j c_j/s_j + 1$ are plotted in the red and green dashed lines respectively.




Figure 10: Schematic of the fully connected autoencoder used in the MNIST experiments. Number of nodes are indicated at the bottom of each layer.

7.2 Deep Networks

7.2.1 MNIST VIA MULTIPLE OPERATORS

We next consider the problem of directly learning the reconstruction function $f : (g, y) \mapsto x$ using measurements from multiple compressed sensing operators, as described in Section 6.2. We train a network that aims to achieve data consistency for all the training data via the unsupervised loss in (54). We use the standard MNIST dataset which has an approximate box-counting dimension $k = 12$ (Hein and Audibert, 2005). The dataset contains $N = 60000$ training samples, and these are partitioned such that $N/|\mathcal{G}|$ different samples are observed via each operator. The entries of the forward operators are sampled from a Gaussian distribution with zero mean and variance n^{-1} . The test set consists of 10000 samples, which are also randomly divided into $|\mathcal{G}|$ parts, one per operator. We define $f(g, y) = \phi \circ A_g^\dagger$ where $\phi : \mathbb{R}^n \mapsto \mathbb{R}^n$ is a trainable network whose aim is to map $A_g^\dagger y$ to the signal set \mathcal{X} . The networks are trained using the Adam optimizer.

	$\phi = \text{Id}$	ϕ w. residual	ϕ w/o residual
PSNR	10.02	10.32	21.36

Table 2: Average PSNR in dB achieved by a residual network and a non-residual network for the MNIST reconstruction task with $|\mathcal{G}| = 40$ operators, each with $m = 100$ measurements.

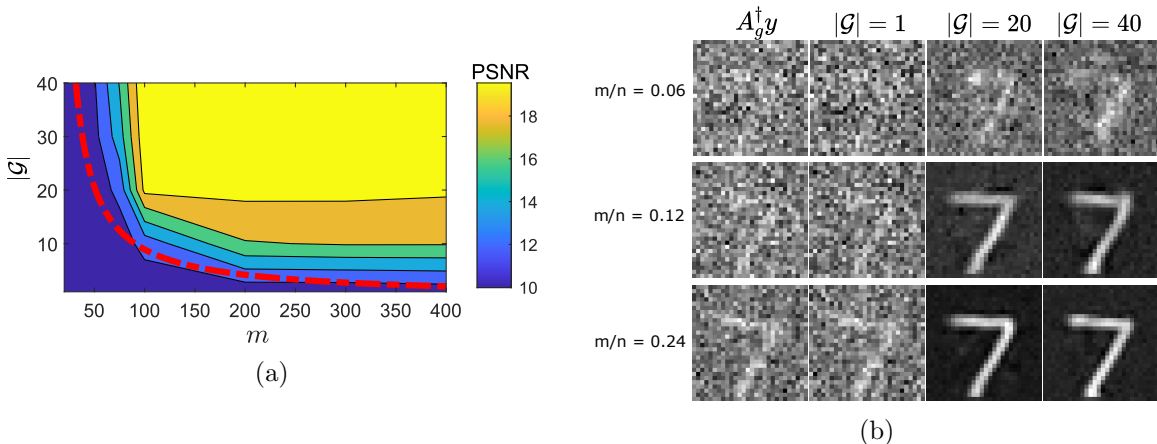


Figure 11: (a) Average test PSNR and (b) reconstructed images for the MNIST dataset, for different number of random Gaussian operators and measurements per operator. The curve in red shows the sufficient condition of Theorem 4, $m > k + n/|\mathcal{G}|$.

When evaluating the ability of neural networks to perform unsupervised learning from incomplete data it is necessary to ensure that we are not just observing the inductive bias of the network, which has been shown can provide a powerful image model without any training (Ulyanov et al., 2018; Tachella et al., 2021). In order to minimize the impact of the inductive bias of the networks’ architecture, we use fully connected layers which do not exploit any spatial image prior. It is easy to see that a valid solution to the optimization problem in (54) is just the identity $\phi(x) = x$. This reflects the fact from the theory that in general the problem is not solvable and that we need to impose a low-dimensional prior on the signal. To avoid learning the identity, we require that $\phi(x)$ exploits the low-dimensionality of the signal set, so we use an autoencoder architecture with 3 hidden layers with 1000, 32 and 1000 neurons, as shown in Figure 10. We use relu non-linearities between layers, except at the output of the last layer.

Table 2 shows the performance for the case of $|\mathcal{G}| = 40$ operators with $m = 100$ measurements each for $\phi(x)$ being (i) the identity, (ii) the autoencoder with a residual connection which allows it to learn the identity mapping and (iii) the autoencoder with no residual connection which enforces a low-dimensional representation. The network with residual connection fails to learn, obtaining a similar performance to the simple pseudo-inverse, whereas the autoencoder without residual obtains a significant improvement of more than 10 dB. This is in line with our theory where we have seen that a requirement for unsupervised learning from incomplete measurements is that the model is low-dimensional. We next explore the ability of the autoencoder without residual connection to perform unsupervised learning as a function of the number of operators $|\mathcal{G}|$ and measurements m .

Figure 11a shows the average test peak-signal-to-noise ratio (PSNR) achieved by the trained model for $|\mathcal{G}| = 1, 10, 20, 30, 40$ and $m = 1, 100, 200, 300, 400$. The results follow closely the bound presented in Section 3.2 which is indicated by the red dashed line, as the network is only able to learn the reconstruction mapping when the sufficient condition $m > k + n/|\mathcal{G}|$ is verified. In sensing regimes below this condition, the performance is similar

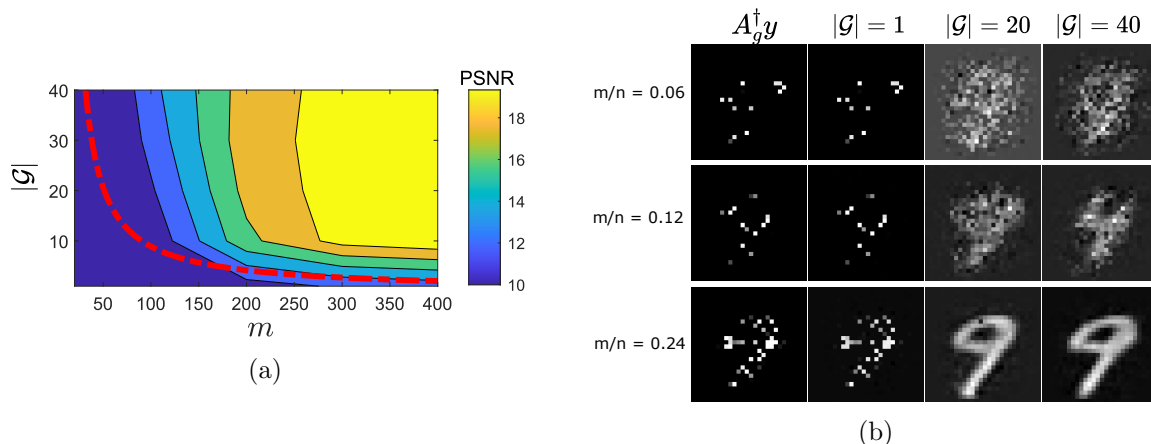


Figure 12: (a) Average test PSNR and (b) reconstructed images for the MNIST dataset, for different number of random inpainting operators and measurements per operator. The curve in red shows the necessary condition of Theorem 4, $m > k + n/|\mathcal{G}|$.

to simply applying the pseudo-inverse A_g^\dagger . Figure 11b shows examples of reconstructed images for networks trained with different number of operators and measurements.

Second, we replace Gaussian operators for $|\mathcal{G}|$ different random inpainting masks. The inpainting operators have a diagonal structure which has zero measure on $\mathbb{R}^{m \times n}$, however our sufficient condition still provides a reasonable lower bound on predicting the performance, as shown in Figure 12a. It is likely that due to the coherence between measurement operators and images (both operators and MNIST images are sparse), more measurements are required to obtain good reconstructions than in the case of Gaussian operators. Figure 12b shows examples of reconstructed images for different number of operators and measurements.

7.2.2 \mathcal{G} -INVARIANT MNIST

Lastly, we consider the problem of learning a \mathcal{G} -invariant version of the MNIST dataset from incomplete measurements using deep convolutional neural networks. In this setting, the goal is to learn the reconstruction function $x_i = f(y_i)$ where f is a trainable network. In order to learn from measurement data alone, we leverage the equivariant imaging approach of Section 6.2. This technique has been very successfully applied to large scale CT and MRI imaging problems (Chen et al., 2021, 2022). Here we demonstrate on a simple example that it also respects the measurement bounds presented in this paper. As in Chen et al. (2021, 2022), we compare 4 learning algorithms:

- **Pseudo-inverse:** Reconstruction by applying the pseudo-inverse to the observed measurements $y_i = A^\dagger x_i$ (baseline of no learning).
- **Supervised:** Standard training using ground truth pairs $\{(x_i, y_i)\}_{i=1}^N$.
- **Unsupervised:** Training from incomplete measurements $\{y_i\}_{i=1}^N$ alone without enforcing \mathcal{G} -invariance of the reconstructed signal set.




Figure 13: Shift-invariant MNIST experiments. (a) Average test PSNR and (b) reconstructions obtained by the competing methods for different compression ratios. The unsupervised method with no equivariance and pseudo-inverse perform similarly. The green and red dash-dotted lines show the necessary and sufficient model identifiability conditions, respectively.

- **Unsupervised equivariant:** Training from incomplete measurements $\{y_i\}_{i=1}^N$ alone. In order to enforce \mathcal{G} -invariance of the reconstructed signal set, we use the equivariant training loss of (55).

In order to generate a \mathcal{G} -invariant dataset, we augment the standard MNIST dataset by applying the transformations associated with the group action. In all cases, we use a single forward operator A with iid Gaussian entries, and the U-Net network of (Chen et al., 2021).

First, we evaluate the group of horizontal and vertical shifts. This is an abelian group with irreducibles of dimension $s_j = 1$, however it is not cyclic. As $\max_j c_j/s_j = 1$ for this group action, assuming Conjecture 10 holds, we expect that both supervised and unsupervised methods are able to reconstruct as long as $m > 2k + 1$. Figure 13 shows the reconstructed images and test peak-to-signal ratio (PSNR) for different compression ratios m/n . The unsupervised method with no equivariance fails to learn for all m , converging to the linear pseudo-inverse. When the number of measurements is not enough for reconstruction, i.e., $m/n < 2k/n \approx 0.03$, both supervised and unsupervised methods fail. For $m/n > 0.07$, the unsupervised equivariant method is able to perform as well as the fully supervised one as expected. There is an intermediate regime $0.03 < m/n < 0.07$ where the supervised method performs better than the unsupervised equivariant. While there are enough measurements for model uniqueness, this discrepancy might be attributed to not having enough measurements for stable model identification.

Secondly, we evaluate the group of 90 degree rotations and reflection, which has $|\mathcal{G}| = 8$ elements. This is a not a cyclic group (nor abelian) which is not covered by Theorem 9. However, we can obtain a bound from Theorem 9 by considering the cyclic subgroup of 90 degree rotations. The full group action has $\max_j c_j/s_j = 0.134n$, whereas the largest cyclic subgroup has $\max_j c_j/s_j = 0.25n$. In both cases, the unsupervised equivariant method cannot achieve the same performance as the supervised one for small m . Figure 14 shows the reconstructed images and test PSNR for different compression ratios m/n . As expected, the unsupervised method with no equivariance fails to learn for all m , converging to the

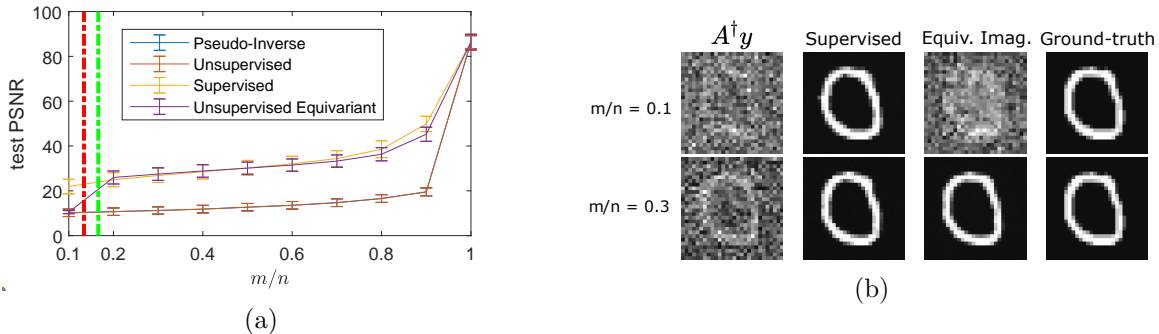


Figure 14: Reflection-invariant MNIST experiments. (a) Average test PSNR and (b) reconstructions obtained by the competing methods for different compression ratios. The unsupervised method with no equivariance and pseudo-inverse perform similarly. The green and red dash-dotted lines show the necessary and sufficient model identifiability conditions, respectively.

linear pseudo-inverse. For $m \geq 0.2n$, both supervised and unsupervised equivariant perform equally well. This result is aligned with Conjecture 10, as the full group action offers a tighter sufficient condition of $m > 0.16n$, in comparison to the condition associated with the largest cyclic subgroup, $m > 0.28n$.

8. Conclusions and Future Work

We have presented fundamental bounds for learning models from incomplete measurements, either by using multiple measurement operators, or by exploiting a group symmetries of the signal set. Our bounds characterize the interplay between the fundamental properties of the inverse problem: the ambient dimension, the signal model dimension, and the number of measurement operators or symmetry of the model. Moreover, the bounds are agnostic of the learning algorithms and provide useful necessary and sufficient conditions for designing principled sensing strategies.

An interesting avenue of research is to extend the bounds in this paper to more general classes of signal models, including compressible distributions (Gribonval et al., 2012), e.g., Bernoulli-Gaussian or heavy tailed models, and other approximately low-dimensional models, e.g., models whose support is close to a low-dimensional set. We leave the study of robustness to noise as well as the extension of the theory for future work. Another interesting avenue of research is to extend our results to more general semigroups, whose elements do not necessarily have an inverse, but can be useful to capture certain symmetries in the signal model such as invariance to scale (Florack et al., 1994; Worrall and Welling, 2019).

Acknowledgments

This work is supported by the ERC C-SENSE project (ERCADG-2015-694888).

Appendix A.

Appendix A. Proofs Preliminaries

Let M be a matrix in $\mathbb{C}^{m \times n}$ and x be a vector in \mathbb{C}^n . In the proofs, we will often use the ‘vec trick’

$$Mx = (x^\top \otimes I_m) \text{vec}(M) \quad (58)$$

where \otimes denotes the Kronecker product, I_m is an $m \times m$ identity matrix and $\text{vec}(M) \in \mathbb{C}^{mn}$ is the column-wise vectorization of the matrix M . We will also use the following result by Sauer et al. (1991):

Lemma 13 (Lemmas 4.5 and 4.6 in Sauer et al. (1991)) *Let S be a bounded subset of \mathbb{R}^n , and let G_0, G_1, \dots, G_t be Lipschitz maps from S to \mathbb{R}^m . For each integer $r \geq 0$, let S_r be the subset of $z \in S$ such that the rank of the $m \times t$ matrix*

$$\Phi_z = [G_1(z), \dots, G_t(z)] \quad (59)$$

is r , and let $\text{boxdim}(S_r) = k_r$. For each $\alpha \in \mathbb{R}^t$ define $G_\alpha(z) = G_0 + \sum_{i=1}^t \alpha_i G_i(z) = G_0 + \Phi_z \alpha$. If for all integers $r \geq 0$ we have that $r > k_r$, then $G_\alpha^{-1}(0)$ is empty for almost every $\alpha \in \mathbb{R}^t$.

The proof of this result follows standard covering arguments and may be sketched as follows. From the dimensionality assumption, the set S_r can be essentially covered by $\mathcal{O}(\epsilon^{-k_r})$ ϵ -balls. Furthermore, for any $z \in S_r$, the probability (measured w.r.t. $\alpha \in \mathbb{R}^t$) that $G_\alpha(z)$ maps to the neighborhood of 0 scales as ϵ^r . Hence the probability of this happening for any of the points in the cover scales as ϵ^{r-k_r} . If we take $r > k_r$ then the probability of such an event tends to zero as we shrink ϵ . Full details can be found in the proofs in Sauer et al. (1991).

Appendix B. Proof of Theorem 4

Proof In order to have model uniqueness, we require that the inferred signal set $\hat{\mathcal{X}}$ defined in (17) equals the true set \mathcal{X} , or equivalently that their difference

$$\hat{\mathcal{X}} \setminus \mathcal{X} = \{v \in \mathbb{R}^n \setminus \mathcal{X} \mid A_1(x_1 - v) = \dots = A_{|\mathcal{G}|}(x_{|\mathcal{G}|} - v) = 0, x_1, \dots, x_{|\mathcal{G}|} \in \mathcal{X}\} \quad (60)$$

is empty, where \setminus denotes set difference. Let $S \subset \mathbb{R}^{n(|\mathcal{G}|+1)}$ be the set of all vectors $z = [v, x_1, \dots, x_{|\mathcal{G}|}]^\top$ with $v \in \mathbb{R}^n \setminus \mathcal{X}$ and $x_1, \dots, x_{|\mathcal{G}|} \in \mathcal{X}$. The difference set defined in (60) is empty if and only if for any $z \in S$ we have

$$\underbrace{\begin{bmatrix} -A_1 & A_1 & & \\ \vdots & & \ddots & \\ -A_{|\mathcal{G}|} & & & A_{|\mathcal{G}|} \end{bmatrix}}_{G_\alpha \in \mathbb{R}^{m|\mathcal{G}| \times n(|\mathcal{G}|+1)}} \underbrace{\begin{bmatrix} v \\ x_1 \\ \vdots \\ x_{|\mathcal{G}|} \end{bmatrix}}_{z \in S \subset \mathbb{R}^{n(|\mathcal{G}|+1)}} \neq 0 \quad (61)$$

$$G_\alpha(z) \neq 0 \quad (62)$$

where G_α maps $z \in S \subset \mathbb{R}^{n(|\mathcal{G}|+1)}$ to $\mathbb{R}^{m|\mathcal{G}|}$. Let $\alpha = [\text{vec}(A_1)^\top, \dots, \text{vec}(A_{|\mathcal{G}|})^\top]^\top \in \mathbb{R}^{mn|\mathcal{G}|}$, then as a function of α we can also write (61) as

$$\begin{bmatrix} (x_1 - v)^\top \otimes I_m & & \\ & \ddots & \\ & & (x_{|\mathcal{G}|} - v)^\top \otimes I_m \end{bmatrix} \alpha \neq 0 \quad (63)$$

where we used the ‘vec trick’ in (58). As v does not belong to the signal set, the matrix on the left hand side of (63) has rank $m|\mathcal{G}|$ for all $z \in S$. We treat the cases of bounded and conic signal sets separately, showing in both cases that, for almost every $\alpha \in \mathbb{R}^{mn|\mathcal{G}|}$, the condition in (63) holds for all $z \in S$ if $m > k + n/|\mathcal{G}|$:

Bounded signal set Decompose S into a countable union of bounded subsets $S = \bigcup_{\rho \geq 1} S_\rho$ defined as

$$S_\rho = \{z \in \mathbb{R}^{n(|\mathcal{G}|+1)} \mid z = [v^\top, x_1^\top, \dots, x_{|\mathcal{G}|}^\top]^\top, x_1, \dots, x_{|\mathcal{G}|} \in \mathcal{X}, \|v\|_2 \leq \rho\}. \quad (64)$$

where $\text{boxdim}(S_\rho) \leq k|\mathcal{G}| + n$. Thus, Lemma 13 states that for almost every α all $z \in S$ verifies (63) if $m > k + n/|\mathcal{G}|$.

Conic signal set If the signal set is conic, then S is also conic. Let B be a bounded set B containing an open neighbourhood of 0. As $\text{boxdim}(S \cap B) \leq |\mathcal{G}|k + n$, Lemma 13 states that for almost every α , all $z \in S$ verifies (63) as long as $m > k + n/|\mathcal{G}|$. ■

Appendix C. Proofs for \mathcal{G} -invariant models

We next present a lemma which will be useful for the proofs of Proposition 8 and Theorem 9.

Lemma 14 *Let $v \in \mathbb{C}^n$ and let the decomposition of v into the J invariant subspaces of a \mathcal{G} -action can be written as:*

$$v = \begin{bmatrix} v_1 \\ \vdots \\ v_J \end{bmatrix} \in \mathbb{C}^n \text{ where } v_j = \begin{bmatrix} v_j^1 \\ \vdots \\ v_j^{c_j} \end{bmatrix} \in \mathbb{C}^{s_j c_j}. \quad (65)$$

where $v_j^\ell \in \mathbb{C}^{s_j}$ corresponds to the ℓ th copy out of c_j multiplicities of the j th invariant subspace. We have

$$\frac{1}{|\mathcal{G}|} \sum_{g \in \mathcal{G}} T_g v v^\top T_g^\top = M_v M_v^\top \quad (66)$$

with

$$M_v = \begin{bmatrix} [v_1^1, v_1^2, \dots, v_1^{c_1}]^\top \otimes I_{s_1} & & \\ & \ddots & \\ & & [v_J^1, v_J^2, \dots, v_J^{c_J}]^\top \otimes I_{s_J} \end{bmatrix}. \quad (67)$$

The result also holds for infinite compact groups by replacing the sum over the group with an integral.

Proof Without loss of generality, we assume that the linear representation of \mathcal{G} is block-diagonalized in the canonical basis (F in Theorem 5 equals the identity). Using the ‘vec trick’ in (58) and the decomposition of v into invariant subspaces, we have

$$T_g v = \begin{bmatrix} \rho_1(g) v_1^1 \\ \vdots \\ \rho_J(g) v_J^{c_J} \end{bmatrix} \quad (68)$$

$$= \begin{bmatrix} ([v_1^1, v_1^2, \dots, v_1^{c_1}]^\top \otimes I_{s_1}) \text{vec} \rho_1(g) \\ \vdots \\ ([v_J^1, v_J^2, \dots, v_J^{c_J}]^\top \otimes I_{s_J}) \text{vec} \rho_J(g) \end{bmatrix} \quad (69)$$

$$= \begin{bmatrix} R_1 \text{vec} \rho_1(g) \\ \vdots \\ R_J \text{vec} \rho_J(g) \end{bmatrix} \quad (70)$$

where the last line defines $R_j := ([v_j^1, v_j^2, \dots, v_j^{c_j}]^\top \otimes I_{s_j}) \in \mathbb{C}^{s_j c_j \times s_j^2}$ for $j = 1, \dots, J$. Using this decomposition, we can compute the (j, j') block of $\sum_{g \in \mathcal{G}} T_g v v^\top T_g^\top$ as

$$\frac{1}{|\mathcal{G}|} \sum_{g \in \mathcal{G}} R_j \text{vec} \rho_j(g) \text{vec} \rho_{j'}(g)^\top R_{j'}^\top = R_j \left(\frac{1}{|\mathcal{G}|} \sum_{g \in \mathcal{G}} \text{vec} \rho_j(g) \text{vec} \rho_{j'}(g)^\top \right) R_{j'}^\top \quad (71)$$

where the middle term can be evaluated using the orthogonality relations of irreducible representations (Serre, 1977, Chapter 2):

$$\frac{1}{|\mathcal{G}|} \sum_{g \in \mathcal{G}} \text{vec} \rho_j(g) \text{vec} \rho_{j'}(g)^\top = \begin{cases} I_{s_j^2} & \text{if } j = j' \\ 0 & \text{otherwise} \end{cases} \quad (72)$$

Considering all blocks, we have¹⁷ $\frac{1}{|\mathcal{G}|} \sum_{g \in \mathcal{G}} T_g v v^\top T_g^\top = M_v M_v^\top$ with

$$M_v = \begin{bmatrix} R_1 & & \\ & \ddots & \\ & & R_J \end{bmatrix}. \quad (73)$$

For infinite compact groups acting on \mathbb{C}^n , we can obtain the same result by replacing the sum over the group by an integral (Serre, 1977, Chapter 4), where the orthogonality relations are given by

$$\int_{\mathcal{G}} \text{vec} \rho_j(g) \text{vec} \rho_{j'}(g)^\top dg = \begin{cases} I_{s_j^2} & \text{if } j = j' \\ 0 & \text{otherwise} \end{cases} \quad (74)$$

■

17. For complex matrices M , M^\top denotes the conjugate transpose.

Proof of Proposition 8.

Proof We begin with the case of finite groups, and then extend it to infinite (but compact) groups. In order to have model uniqueness it is necessary that the matrix

$$M^\top = \begin{bmatrix} AT_1 \\ \vdots \\ AT_{|\mathcal{G}|} \end{bmatrix} \quad (75)$$

has rank n . This matrix contains the orbits of the measurement vectors $\{a_i\}_{i=1}^m$ (the rows of A). Letting $R = \frac{1}{|\mathcal{G}|}MM^\top$ and using Lemma 14, we have

$$R = \sum_{i=1}^m \frac{1}{|\mathcal{G}|} \sum_{g \in \mathcal{G}} T_g a_i a_i^\top T_g^\top \quad (76)$$

$$= \sum_{i=1}^m M_{a_i} M_{a_i}^\top \quad (77)$$

$$= [M_{a_1}, \dots, M_{a_m}] [M_{a_1}, \dots, M_{a_m}]^\top \quad (78)$$

where $M_{a_i} \in \mathbb{C}^{n \times \sum_j s_j^2}$ is the block-diagonal matrix in (67) associated to a_i . Note that $\text{rank}(M^\top) = \text{rank}(R) = \text{rank}([M_{a_1}, \dots, M_{a_m}])$. Furthermore, due to the block-diagonal structure of the submatrices M_{a_i} , $[M_{a_1}, \dots, M_{a_m}]$ can also be rearranged in block diagonal form with blocks of size $s_j c_j \times m s_j^2$. Thus, M^\top has rank n only if all the blocks verify $m s_j^2 \geq s_j c_j$, which yields the bound

$$m \geq \max_j c_j / s_j. \quad (79)$$

If the group has infinite elements, the matrix in (75) is not well-defined as it would have infinite entries. However, as the decomposition into finite invariant subspaces and Lemma 14 still hold, we can compute the dimension of the subspace spanned by $\{T_g a_i\}_{g \in \mathcal{G}, i=1, \dots, m}$ with (76) by replacing the finite sum over group elements by an integral and obtain the bound $m \geq \max_j c_j / s_j$. ■

We now present two useful technical lemmas for proving Theorem 9.

Lemma 15 *Let $T \in \mathbb{R}^{n \times n}$ be the linear representation of the generator of a finite cyclic group of order $|\mathcal{G}|$. Let $v \in \mathbb{R}^n$ and $B \in \mathbb{R}^{n \times r+1}$ with $r+1 \leq n$, such that $M = B - [Tv, \dots, T^{r+1}v]$ has rank r and the first r columns are linearly independent. The set $\Omega \subset \mathbb{R}^n$ of v verifying the rank assumption is contained in an affine variety of dimension at most $r + \sum_{j \in \mathcal{J}} c_j$ where c_j denotes the multiplicity of the j th irreducible representation of the group action and \mathcal{J} is a subset of r out of $|\mathcal{G}|$ irreducibles.*

Proof By assumption, $M = B - [Tv, \dots, T^{r+1}v]$ has rank r . All $(r+1) \times (r+1)$ minors of matrices with rank at most r are necessarily zero. Minors are given by polynomial equations

on the entries of M . Thus, the (at most) rank- r condition on M can be translated into a set of polynomial equations on v which must equal zero. As the zero set of polynomial equations, the subset of $v \in \mathbb{R}^n$ which verify this constraint is a variety (Cox et al., 2013). We will use these polynomial equations to show that this set has dimension at most $r + \sum_{j \in \mathcal{J}} c_j$ where \mathcal{J} is a subset of r out of $|\mathcal{G}|$ irreducibles. The polynomials depend on both B and v , thus the v -variety will vary smoothly as a function of B . In order to simplify the analysis, we study the set of complex $v \in \mathbb{C}^n$ that verify the rank condition, noting that the set of real $v \in \mathbb{R}^n$ is just a subset of the complex setting.

Using Theorem 5, the matrix M can be written as

$$M = \begin{bmatrix} B_{1,1} - \rho_1(1)v_1 & B_{1,2} - \rho_1(2)v_1 & \dots & B_{1,r+1} - \rho_1(r+1)v_1 \\ B_{2,1} - \rho_{j_2}(1)v_2 & B_{2,2} - \rho_{j_2}(2)v_2 & \dots & B_{2,r+1} - \rho_{j_2}(r+1)v_2 \\ \vdots & \vdots & \ddots & \vdots \\ B_{n,1} - \rho_J(1)v_n & B_{n,2} - \rho_J(2)v_n & \dots & B_{n,r+1} - \rho_J(r+1)v_n \end{bmatrix} \quad (80)$$

where we use the index j_i to indicate the irreducible representation associated with the i th row. For example, if $J = 2$ and $c_1 = c_2 = 2$, we have that $j_1 = j_2 = 1$ and $j_3 = j_4 = 2$. Moreover, as we are dealing with cyclic groups, we have that $\rho_j(g) = e^{-i2\pi jg/|\mathcal{G}|}$ (Serre, 1977, Chapter 5).

By assumption, there is at least one $r \times r$ invertible submatrix within the first r columns. Without loss of generality, we assume that this submatrix is given by the first r rows. For any other choice of rows, the results are identical with a re-indexing. We then consider the $n - r$ minors of size $(r+1) \times (r+1)$ corresponding to the first $r+1$ columns together with the first r rows and the i th row:

$$\det \begin{bmatrix} B_{1,1} - \rho_1(1)v_1 & \dots & B_{1,r+1} - \rho_1(r+1)v_1 \\ \vdots & \ddots & \vdots \\ B_{r,1} - \rho_{j_r}(1)v_r & \dots & B_{r,r+1} - \rho_{j_r}(r+1)v_r \\ B_{i,1} - \rho_{j_i}(1)v_i & \dots & B_{i,r+1} - \rho_{j_i}(r+1)v_i \end{bmatrix} \quad (81)$$

for $i = r+1, \dots, n$. Applying Laplace's expansion on the last row of each minor, we have that

$$\sum_{g=1}^{r+1} (B_{i,g} - \rho_{j_i}(g)v_i) \mu_g = 0 \quad (82)$$

where μ_1, \dots, μ_{r+1} are the determinants of $r \times r$ matrices which only depend on the first $r+1$ columns of B and v_1, \dots, v_r . Let $\mu = [\mu_1, \dots, \mu_{r+1}, 0, \dots, 0]^\top \in \mathbb{C}^{|\mathcal{G}|}$, where $\mu_g = 0$ for all $g > r+1$. Note that μ cannot be identically zero, as $\mu_{r+1} \neq 0$ is the determinant of the submatrix with the first r columns and rows, which is invertible. The equations can be rewritten as

$$\begin{bmatrix} \hat{\mu}_{j_{r+1}} & & \\ & \ddots & \\ & & \hat{\mu}_{j_n} \end{bmatrix} \begin{bmatrix} v_{r+1} \\ \vdots \\ v_n \end{bmatrix} = \begin{bmatrix} d_{r+1} \\ \vdots \\ d_n \end{bmatrix} \quad (83)$$

where $\hat{\mu}_j = \sum_{g=1}^{|\mathcal{G}|} \rho_j(g) \mu_g$ are the j th coefficient of the discrete Fourier transform of the vector μ (which has at most $r+1$ consecutive nonzero elements) and $d_i = \sum_{g=1}^{r+1} B_{i,g} \mu_g$.

For all $i \in \{r+1, \dots, n\}$ where $\hat{\mu}_{j_i} \neq 0$, we have $v_i = d_i/\hat{\mu}_{j_i}$ in order to satisfy the rank- r constraint. As the Fourier transform of a vector with at most $r+1$ non-zero elements which are consecutive has at most r zero coefficients (Donoho and Stark, 1989, Lemma 5), we have that at most r coefficients $\hat{\mu}_j = 0$. As each $\hat{\mu}_j$ is repeated at most c_j times in (83), there are at most $\sum_{j \in \mathcal{J}} c_j$ zeros along the diagonal in (83), where \mathcal{J} is a subset of r out of $|\mathcal{G}|$ irreducibles. Locally, we are free to vary the first r components v_1, \dots, v_r without changing the rank of the submatrix. From (83), we are also free to locally vary at most $\sum_{j \in \mathcal{J}} c_j$ components of the remaining v_i associated with the zero set of $\hat{\mu}_{j_i}$. Hence, the set of v that verifies the rank- r constraint has a dimension equal or smaller than $r + \sum_{j \in \mathcal{J}} c_j$. ■

Lemma 16 *For any infinite compact cyclic group \mathcal{G}_1 acting on \mathbb{R}^n , there is a finite cyclic subgroup $\mathcal{G}_2 \subset \mathcal{G}_1$ such that the restriction of the linear representation of \mathcal{G}_1 to \mathcal{G}_2 has the same multiplicities of the irreducible representations.*

Proof Following Theorem 5, the representation of a compact infinite cyclic group \mathcal{G}_1 in \mathbb{C}^n is given by

$$T_g = F \begin{bmatrix} e^{-i2\pi j_1 g} & & \\ & \ddots & \\ & & e^{-i2\pi j_J g} \end{bmatrix} F^{-1} \quad (84)$$

for the elements $g \in (0, 1]$. This representation contains $J = n/\sum_j c_j$ distinct irreducibles given by $\rho_j(g) = e^{-i2\pi j g}$ with integers $j_1 < \dots < j_J$. The linear representation has the same form on \mathbb{R}^n as in (84), with the additional constraint that for each j_i the conjugate pair $-j_i$ is also present in the linear representation (and thus $j_1 = -j_J$). The restriction to a finite subgroup \mathcal{G}_2 with p elements is given by

$$\tilde{T}_r = F \begin{bmatrix} e^{-i2\pi j_1 \frac{r}{p}} & & \\ & \ddots & \\ & & e^{-i2\pi j_J \frac{r}{p}} \end{bmatrix} F^{-1} \quad (85)$$

for the elements $r = 1, \dots, p$. In order to ensure that the representation of \mathcal{G}_2 has the same multiplicities than the one of \mathcal{G}_1 , we need that no two distinct irreducibles of \mathcal{G}_1 are mapped to a common irreducible of the finite subgroup. For any subgroup of size $p > 2j_J$, the diagonal entries in (85), $\rho_j(r) = e^{-i2\pi j \frac{r}{p}}$, are orthogonal as vectors in \mathbb{C}^p for any $j \neq j'$ as

$$\frac{1}{p} \sum_{r=1}^p \rho_j(r) \rho_{j'}^*(r) = \begin{cases} 1 & \text{if } j = j' \\ 0 & \text{otherwise} \end{cases} \quad (86)$$

where $*$ denotes complex conjugation. Thus, they are distinct irreducible representations of \mathcal{G}_2 . By Theorem 5, the representation in (85) is in diagonal form, and (84) and (85) have the same multiplicities. ■

Proof of Theorem 9.

Proof Our proof focuses on finite cyclic groups, whose irreducible representations have dimension $s_j = 1$. If the cyclic group is infinite and compact, we can restrict its action

to a finite subgroup with the same multiplicities c_j of the irreducible representations using Lemma 16. Similarly to Theorem 4, we have to prove that $\hat{\mathcal{X}} \setminus \mathcal{X}$ is empty for almost every $A \in \mathbb{R}^{m \times n}$ if $m > 2k + \max_j c_j + 1$. In this case, this is equivalent to

$$\underbrace{\begin{bmatrix} -AT_1 & AT_1 & & \\ \vdots & & \ddots & \\ -AT_{|\mathcal{G}|} & & & AT_{|\mathcal{G}|} \end{bmatrix}}_{G_\alpha \in \mathbb{R}^{n(|\mathcal{G}|+1) \times m|\mathcal{G}|}} \underbrace{\begin{bmatrix} v \\ x_1 \\ \vdots \\ x_{|\mathcal{G}|} \end{bmatrix}}_{z \in \mathbb{R}^{n(|\mathcal{G}|+1)}} \neq 0 \quad (87)$$

$$G_\alpha(z) \neq 0 \quad (88)$$

for any $x_1, \dots, x_{|\mathcal{G}|} \in \mathcal{X}$ and $v \in \mathbb{R}^n \setminus \mathcal{X}$. (87) can be rewritten as

$$A\Phi_z \neq 0 \quad (89)$$

where

$$\Phi_z = [T_1(x_1 - v), \dots, T_{|\mathcal{G}|}(x_{|\mathcal{G}|} - v)] \in \mathbb{R}^{n \times |\mathcal{G}|}. \quad (90)$$

Moreover, letting $T_1 = T$ be the linear representation of the generator of the group, we can write $T_r = T^r$ where T^r denotes the r th power of T . Thus, we can rewrite (90) as

$$\Phi_z = [T(x_1 - v), \dots, T^{|\mathcal{G}|}(x_{|\mathcal{G}|} - v)] \in \mathbb{R}^{n \times |\mathcal{G}|}. \quad (91)$$

Equation (89) requires that the nullspace of A does not contain the range of Φ_z for any choice of z . We perform a separate analysis for the cases of full rank or rank-deficient Φ_z . We decompose $S = S_1 \cup S_2$, where S_1 is the set of z such that Φ_z has full rank (i.e., $\text{rank}(\Phi_z) = \min\{n, |\mathcal{G}|\}$) and S_2 is the set of z such that Φ_z is rank-deficient.

We begin with analyzing the (simpler) full-rank case associated with S_1 . If $|\mathcal{G}| \geq n$, then $\text{rank}(\Phi_z) = n$ and (89) implies the trivial inequality $A \neq 0$ which has measure zero in $\mathbb{R}^{m \times n}$. If $|\mathcal{G}| < n$, then, in order to apply Lemma 13, we need to compute the dimension of

$$S_1 = \{z \in (\mathbb{R}^n \setminus \mathcal{X}) \times \mathcal{X}^{|\mathcal{G}|} \mid \text{rank}(\Phi_z) = |\mathcal{G}|\} \quad (92)$$

As being full-rank is an open condition, S_1 has the same dimension as $(\mathbb{R}^n \setminus \mathcal{X}) \times \mathcal{X}^{|\mathcal{G}|}$. The analysis of this set can be done in the same way as in the proof of Theorem 4 (which we do not repeat here), where we have in both cases that almost every $A \in \mathbb{R}^{m \times n}$ with $m > k + n/|\mathcal{G}|$ verifies $A\Phi_z \neq 0$. Moreover, as $n/|\mathcal{G}| \geq \max_j c_j$ for any linear representation of a cyclic group, we have that

$$m > k + n/|\mathcal{G}| > 2k + \max_j c_j + 1. \quad (93)$$

We now treat the rank-deficient case associated with S_2 . Let $\Phi_{z,r} \in \mathbb{R}^{n \times r}$ denote the first r columns of Φ_z . As no column of Φ_z can be exactly 0, we can decompose S_2 as

$$S_2 = \bigcup_{r=0}^{\min\{n-1, |\mathcal{G}|-1\}} S_{2,r} \quad (94)$$

where $S_{2,r}$ is the subset of $z \in S$ such that the first r columns are linearly independent (i.e., $\text{rank}(\Phi_{z,r}) = r$) but the first $r+1$ columns of Φ_z are rank-deficient (i.e., $\text{rank}(\Phi_{z,r+1}) = r$). We will consider $A\Phi_{z,r} \neq 0$, as this necessarily implies $A\Phi_z \neq 0$. In order to apply Lemma 13, we need to bound the dimension of the set¹⁸

$$S_{2,r} = \{\tilde{z} = [v^\top, x_1^\top, \dots, x_{r+1}^\top]^\top \in (\mathbb{R}^n \setminus \mathcal{X}) \times \mathcal{X}^{r+1} \mid \text{rank}(\Phi_{z,r}) = \text{rank}(\Phi_{z,r+1}) = r\}. \quad (95)$$

By Lemma 15, the set of v in z verifying the constraint $\text{rank}(\Phi_{z,r}) = \text{rank}(\Phi_{z,r+1}) = r$ is a variety of \mathbb{R}^n of dimension $d \leq r + \sum_{\mathcal{J}} c_j$ which varies smoothly as a function of x_1, \dots, x_{r+1} , where c_j denotes the multiplicity of the j th irreducible representation of the group action and \mathcal{J} is a subset of r out of $|\mathcal{G}|$ irreducibles. A variety consists of a finite union of irreducible¹⁹ varieties, where each irreducible has the structure of a smooth manifold (Cox et al., 2013). Thus, we decompose $S_{2,r}$ as a finite union of irreducible varieties, $S_{2,r} \subseteq \bigcup_{\ell} S_{2,r,\ell}$, and divide each variety further as a countable union of local (bounded) neighbourhoods, yielding

$$S_{2,r} \subseteq \bigcup_{\ell} \bigcup_{i=1}^{\infty} S_{2,r,\ell,i} \quad (96)$$

As the v -variety has a smooth dependence on x_1, \dots, x_{r+1} by Lemma 15, for each subset $S_{2,r,\ell,i}$, we can build a smooth locally Lipschitz mapping $f : S_{2,r,\ell,i} \subset \mathbb{R}^{n(r+2)} \mapsto \mathbb{R}^{n(r+2)}$ such that $f(S_{2,r,\ell,i})$ has a simple product structure $\{0\}^{n-d} \times (0,1)^d \times (W \cap \mathcal{X}^{r+1})$ where W is local neighborhood of $\mathbb{R}^{n(r+1)}$ (see (Falconer, 2004, Chapter 7)). As a Lipschitz mapping does not increase the box-counting dimension of a set (Robinson, 2010), we have

$$\text{boxdim}(S_{2,r,\ell,i}) \leq \text{boxdim}(\{0\}^{n-d}) + \text{boxdim}((0,1)^d) + \text{boxdim}(W \cap \mathcal{X}^{r+1}) \quad (97)$$

$$\leq d + k(r+1) \quad (98)$$

$$\leq r + \sum_{j \in \mathcal{J}} c_j + k(r+1) \quad (99)$$

for both bounded and conic signal sets \mathcal{X} . As we have $\text{rank}(\Phi_{z,r}) = r$, according to Lemma 13, for almost every $A \in \mathbb{R}^{m \times n}$ there is no $z \in S_{2,r,\ell,i}$ which verifies $A\Phi_{z,r} \neq 0$ as long as m verifies

$$mr > k(r+1) + \sum_{j \in \mathcal{J}} c_j + r \quad (100)$$

$$m > k(1 + 1/r) + \sum_{j \in \mathcal{J}} c_j/r + 1 \quad (101)$$

$$m > 2k + \max_j c_j + 1 \quad (102)$$

where the last inequality is obtained by using $r \geq 1$ and uses the fact that $|\mathcal{J}| = r$. As the countable union of events of measure zero have measure zero, then for almost every $A \in \mathbb{R}^{m \times n}$ with $m > 2k + \max_j c_j + 1$, all $z \in S_{2,r}$ verifies $A\Phi_z \neq 0$ for all possible ranks r .

18. Note that we do not consider $x_{r+2}, \dots, x_{|\mathcal{G}|}$ in $S_{2,r}$ as, for any fixed v, x_1, \dots, x_{r+1} , they yield the same submatrix $\Phi_{z,r}$.

19. Despite having the same name, the term *irreducible* variety is not related to the irreducible group representations.

■

Proof of Theorem 11.

Proof We need to show that the following statements are equivalent:

- (i) The range space of $(AT_g)^\dagger$ is the same for all $g \in \mathcal{G}$.
- (ii) A is a \mathcal{G} -equivariant map.
- (iii) A can be decomposed as $A = \tilde{F}\Sigma F^{-1}$ where Σ has the block-structure in Theorem 5.

We first prove (i) \implies (ii). Let $V \in \mathbb{R}^{m \times n}$ be an orthogonal basis for the range space of A^\dagger , such that $A = MV^\top$ where $M \in \mathbb{R}^{m \times m}$ is an invertible matrix. Due to (i), we have

$$AT_g = MQ_gV^\top \tag{103}$$

$$= \underbrace{MQ_gM^{-1}}_{U_g} MV^\top \tag{104}$$

$$= U_gA \tag{105}$$

where $Q_g \in \mathbb{R}^{m \times m}$ is an orthogonal matrix and $U_g = MQ_gM^{-1} \in \mathbb{R}^{m \times m}$ is invertible. The mapping $g \mapsto U_g$ is a valid linear representation (e.g., see (Serre, 1977, Chapter 1)) as we have that for any two group elements g and g' , $U_{g \cdot g'}A = AT_{g \cdot g'} = AT_gT_{g'} = U_{g'}AT_g = U_gU_{g'}A$ and thus $U_{g \cdot g'} = U_gU_{g'}$. As we have $U_gA = AT_g$ where U_g and T_g are two linear representations of \mathcal{G} , A is an equivariant map.

We now prove (ii) \implies (i). Using the decomposition $A = MV^\top$ and (ii) we have

$$AT_g = \tilde{T}_gA \tag{106}$$

$$= \tilde{T}_gMV^\top \tag{107}$$

$$= \tilde{U}_gV^\top \tag{108}$$

where $\tilde{U}_g = \tilde{T}_gM$ is an invertible matrix. As A and AT_g share the same right singular vectors given by the matrix V , they have the same range space for all $g \in \mathcal{G}$.

The proof of (ii) \iff (iii) follows from Theorem 5 and is a standard result of linear representation theory which can be found in, for example, (Reeder, 2014) and (Stiefel and Fässler, 2012). ■

References

Mohammad Aghagolzadeh and Hayder Radha. New guarantees for blind compressed sensing. In *2015 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 1227–1234. IEEE, 2015.

Cecilia Aguerrebere, Mauricio Delbracio, Alberto Bartesaghi, and Guillermo Sapiro. Fundamental limits in multi-image alignment. *IEEE Transactions on Signal Processing*, 64(21):5707–5722, 2016.

- 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.
- Ashish Bora, Eric Price, and Alexandros G. Dimakis. AmbientGAN: Generative models from lossy measurements. In *International Conference on Learning Representations*, 2018.
- Anthony Bourrier, Mike E 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.
- Hilton Bristow, Anders Eriksson, and Simon Lucey. Fast convolutional sparse coding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 391–398, 2013.
- Jian-Feng Cai, Emmanuel J Candès, and Zuowei Shen. A singular value thresholding algorithm for matrix completion. *SIAM Journal on optimization*, 20(4):1956–1982, 2010.
- Emmanuel J Candès and Benjamin Recht. Exact matrix completion via convex optimization. *Foundations of Computational mathematics*, 9(6):717, 2009.
- Dongdong Chen, Julián Tachella, and Mike E. Davies. Equivariant imaging: Learning beyond the range space. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4379–4388, October 2021.
- Dongdong Chen, Julián Tachella, and Mike E Davies. Robust equivariant imaging: a fully unsupervised framework for learning to image from noisy and partial measurements. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5647–5656, June 2022.
- Y. Chen, Y. Chi, and A. J. Goldsmith. Exact and stable covariance estimation from quadratic sampling via convex programming. *IEEE Transactions on Information Theory*, 61(7):4034–4059, 2015. doi: 10.1109/TIT.2015.2429594.
- David Cox, John Little, and Donal OShea. *Ideals, varieties, and algorithms: an introduction to computational algebraic geometry and commutative algebra*. Springer Science & Business Media, 2013.
- Harald Cramér and Herman Wold. Some theorems on distribution functions. *Journal of the London Mathematical Society*, 1(4):290–294, 1936.
- David L Donoho and Philip B Stark. Uncertainty principles and signal recovery. *SIAM Journal on Applied Mathematics*, 49(3):906–931, 1989.
- Morris L Eaton. Group invariance applications in statistics. In *Regional conference series in Probability and Statistics*, pages i–133. JSTOR, 1989.
- Brian Eriksson, Laura Balzano, and Robert Nowak. High-rank matrix completion. In *Artificial Intelligence and Statistics*, pages 373–381. PMLR, 2012.

- Kenneth Falconer. *Fractal geometry: mathematical foundations and applications*. John Wiley & Sons, 2004.
- Luc MJ Florack, Bart M ter Haar Romeny, Jan J Koenderink, and Max A Viergever. Linear scale-space. *Journal of Mathematical Imaging and Vision*, 4(4):325–351, 1994.
- Sivan Gleichman and Yonina C Eldar. Blind compressed sensing. *IEEE Transactions on Information Theory*, 57(10):6958–6975, 2011.
- Rémi Gribonval, Volkan Cevher, and Mike E Davies. Compressible distributions for high-dimensional statistics. *IEEE Transactions on Information Theory*, 58(8):5016–5034, 2012.
- Rémi Gribonval, Gilles Blanchard, Nicolas Keriven, and Yann Traonmilin. Compressive statistical learning with random feature moments. *arXiv preprint arXiv:1706.07180*, 2017.
- Chunli Guo and Mike E Davies. Near optimal compressed sensing without priors: Parametric sure approximate message passing. *IEEE Transactions on Signal Processing*, 63(8):2130–2141, 2015.
- Harshit Gupta, Thong H Phan, Jaejun Yoo, and Michael Unser. Multi-cryoGAN: Reconstruction of continuous conformations in cryo-EM using generative adversarial networks. In *European Conference on Computer Vision*, pages 429–444. Springer, 2020.
- 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)*, pages 289–296, 2005.
- Maximilian Jaritz, Raoul De Charette, Emilie Wirbel, Xavier Perrotton, and Fawzi Nashashibi. Sparse and dense data with CNNs: Depth completion and semantic segmentation. In *2018 International Conference on 3D Vision (3DV)*, pages 52–60, 2018. doi: 10.1109/3DV.2018.00017.
- Kyong Hwan Jin, Michael T McCann, Emmanuel Froustey, and Michael Unser. Deep convolutional neural network for inverse problems in imaging. *IEEE Transactions on Image Processing*, 26(9):4509–4522, 2017.
- 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.
- Christopher A Metzler, Ali Mousavi, Reinhard Heckel, and Richard G Baraniuk. Unsupervised learning with stein’s unbiased risk estimator. *arXiv preprint arXiv:1805.10531*, 2018.
- 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.

- Greg Ongie, Rebecca Willett, Robert D Nowak, and Laura Balzano. Algebraic variety models for high-rank matrix completion. In *International Conference on Machine Learning*, pages 2691–2700. PMLR, 2017.
- Matthew O’Toole, David B Lindell, and Gordon Wetzstein. Confocal non-line-of-sight imaging based on the light-cone transform. *Nature*, 555(7696):338–341, 2018.
- Daniel Pimentel-Alarcon and Robert Nowak. The information-theoretic requirements of subspace clustering with missing data. In *International Conference on Machine Learning*, pages 802–810. PMLR, 2016.
- Daniel L Pimentel-Alarcón, Nigel Boston, and Robert D Nowak. A characterization of deterministic sampling patterns for low-rank matrix completion. *IEEE Journal of Selected Topics in Signal Processing*, 10(4):623–636, 2016.
- 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.
- Joshua Rapp, Julian Tachella, Yoann Altmann, Stephen McLaughlin, and Vivek K Goyal. Advances in single-photon lidar for autonomous vehicles: Working principles, challenges, and recent advances. *IEEE Signal Processing Magazine*, 37(4):62–71, 2020.
- Mark Reeder. *Notes on representations of finite groups*. 2014.
- James C Robinson. *Dimensions, embeddings, and attractors*, volume 186. Cambridge University Press, 2010.
- 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.
- Tim Sauer, James A Yorke, and Martin Casdagli. Embedology. *Journal of statistical Physics*, 65(3):579–616, 1991.
- Jean-Pierre Serre. *Linear representations of finite groups*, volume 42. Springer, 1977.
- Jorge Silva, Minhua Chen, Yonina C Eldar, Guillermo Sapiro, and Lawrence Carin. Blind compressed sensing over a structured union of subspaces. *arXiv preprint arXiv:1103.2469*, 2011.
- Charles M Stein. Estimation of the mean of a multivariate normal distribution. *The annals of Statistics*, pages 1135–1151, 1981.
- E Stiefel and A Fässler. *Group theoretical methods and their applications*. Springer Science & Business Media, 2012.
- Christoph Studer and Richard G Baraniuk. Dictionary learning from sparsely corrupted or compressed signals. In *2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3341–3344. IEEE, 2012.
- Bernd Sturmfels. *Algorithms in invariant theory*. Springer Science & Business Media, 2008.

- Julian Tachella, Junqi Tang, and Mike Davies. The neural tangent link between CNN denoisers and non-local filters. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8618–8627, June 2021.
- Julián Tachella, Dongdong Chen, and Mike Davies. Unsupervised learning from incomplete measurements for inverse problems. In *Advances in Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=aV9WSvM6N3>.
- Matthew Tancik, Ben Mildenhall, Terrance Wang, Divi Schmidt, Pratul P Srinivasan, Jonathan T Barron, and Ren Ng. Learned initializations for optimizing coordinate-based neural representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2846–2855, 2021.
- Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Deep image prior. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9446–9454, 2018.
- Daniel E Worrall and Max Welling. Deep scale-spaces: Equivariance over scale. *arXiv preprint arXiv:1905.11697*, 2019.
- Congyuan Yang, Daniel Robinson, and Rene Vidal. Sparse subspace clustering with missing entries. In *International Conference on Machine Learning*, pages 2463–2472. PMLR, 2015a.
- Jianbo Yang, Xuejun Liao, Xin Yuan, Patrick Llull, David J. Brady, Guillermo Sapiro, and Lawrence Carin. Compressive sensing by learning a Gaussian mixture model from measurements. *IEEE Transactions on Image Processing*, 24(1):106–119, 2015b. doi: 10.1109/TIP.2014.2365720.