# Linear Regression with an Unknown Permutation: Statistical and Computational Limits 

Ashwin Pananjady ${ }^{\dagger}$ Martin J. Wainwright ${ }^{\dagger}, \star$ Thomas A. Courtade ${ }^{\dagger}$<br>Department of Electrical Engineering and Computer Sciences ${ }^{\dagger}$<br>Department of Statistics ${ }^{\star}$<br>UC Berkeley

August 10, 2016


#### Abstract

Consider a noisy linear observation model with an unknown permutation, based on observing $y=\Pi^{*} A x^{*}+w$, where $x^{*} \in \mathbb{R}^{d}$ is an unknown vector, $\Pi^{*}$ is an unknown $n \times n$ permutation matrix, and $w \in \mathbb{R}^{n}$ is additive Gaussian noise. We analyze the problem of permutation recovery in a random design setting in which the entries of the matrix $A$ are drawn i.i.d. from a standard Gaussian distribution, and establish sharp conditions on the SNR, sample size $n$, and dimension $d$ under which $\Pi^{*}$ is exactly and approximately recoverable. On the computational front, we show that the maximum likelihood estimate of $\Pi^{*}$ is NP-hard to compute, while also providing a polynomial time algorithm when $d=1$.


## 1 Introduction

Recovery of a vector based on noisy linear measurements is the classical problem of linear regression, and is arguably the most basic problem in statistical inference. A variant, the "errors-in-variables" model [LR14], allows for errors in the measurement matrix, but mainly in the form of additive or multiplicative noise [LW12. In this paper, we study a form of errors-in-variables in which the measurement matrix is perturbed by an unknown permutation of its rows.

More concretely, we study an observation model of the form

$$
\begin{equation*}
y=\Pi^{*} A x^{*}+w, \tag{1}
\end{equation*}
$$

where $x^{*} \in \mathbb{R}^{d}$ is an unknown vector, $A \in \mathbb{R}^{n \times d}$ is a measurement (or design) matrix, $\Pi^{*}$ is an unknown $n \times n$ permutation matrix, and $w \in \mathbb{R}^{n}$ is observation noise. We refer to the setting where $w=0$ as the noiseless case. As with linear regression, there are two settings of interest, corresponding to whether the design matrix is (i) deterministic (the fixed design case), or (ii) random (the random design case).

There are also two complementary problems of interest - recovery of the unknown $\Pi^{*}$, and recovery of the unknown $x^{*}$. In this paper, we focus on the former problem; the latter problem is also known as unlabelled sensing [UHV15].

The observation model (1) is frequently encountered in scenarios where there is uncertainty in the order in which measurements are taken. An illustrative example is that of sampling in the presence of jitter Bal62, in which the uncertainty about the instants at which measurements are taken results in an unknown permutation of the measurements. A similar synchronization issue occurs in timing and molecular channels RMS12. Here, identical molecular tokens are received at the receptor at different times, and their signatures are


3D object

$$
\left(\begin{array}{ccc}
p_{1} & q_{1} & r_{1} \\
\vdots & \vdots & \vdots \\
p_{n} & q_{n} & r_{n}
\end{array}\right)
$$



Unknown linear transformation
2D image
$\times\left(\begin{array}{ll}x_{11} & x_{12} \\ x_{21} & x_{22} \\ x_{31} & x_{32}\end{array}\right)=\Pi\left(\begin{array}{cc}u_{1} & u_{2} \\ \vdots & \vdots \\ u_{n} & v_{n}\end{array}\right)$

Figure 1. Example of pose and correspondence estimation. The camera introduces an unknown linear transformation corresponding to the pose. The unknown permututation represents the correspondence between points, which is shown in the picture via coloured shapes, and needs to be estimated.
indistinguishable. The vectors of transmitted and received times correspond to the signal and the observations, repectively, where the latter is some permuted version of the former with additive noise.

Another such scenario arises in multi-target tracking problems [PG06]. For example, SLAM tracking [TL08] is a classical problem in robotics where the environment in which measurements are made is unknown, and part of the problem is to infer relative permutations between measurements. Archaeological measurements Rob51 also suffer from an inherent lack of ordering, which makes inference of chronology hard. Another compelling example of such an observation model is in data anonymization, in which the order, or "labels", of measurements are intentionally deleted to preserve privacy. The inverse problem of data de-anonymization [NS08] is to infer these labels from the observations.

Also, in large sensor networks, it is often the case that the number of bits of information that each sensor records and transmits to the server is exceeded by the number of bits it transmits in order to identify itself to the server $\left[\mathrm{KSF}^{+} 09\right]$. In applications where sensor measurements are linear, model (1) corresponds to the case where each sensor only sends its measurement but not its identity. The server is then tasked with recovering sensor identites, or equivalently, with determining the unknown permutation.

The pose and correspondence estimation problem in image processing [DDDS04, MSC09] is also related to the observation model (1). The capture of a 3D object by a 2D image can be modelled by an unknown linear transformation called the "pose", and an unknown permutation representing the "correspondence" between points in the two spaces. One of the central goals in image processing is to identify this correspondence information, which in this case is equivalent to permutation estimation in the linear model. An illustration of the problem is provided in Figure 1.

The discrete analog of the model (1) in which the vectors $x^{*}$ and $y$, and the matrix $A$ are all constrained to belong to some finite alphabet/field corresponds to the permutation channel studied by Schulman and Zuckerman [SZ99], with $A$ representing the (linear) encoding matrix. However, techniques for the discrete problem do not carry over to the continuous problem (11).

Another line of work that is related in spirit to the observation model (1) is the genome assembly problem from shotgun reads [HM99], in which an underlying vector $x^{*} \in\{A, T, G, C\}^{d}$ must be assembled from an unknown permutation of its continuous sub-vector measurements, called "reads". Two aspects, however, render it a particularization of our observation model, besides the obvious fact that $x^{*}$ in the genome assembly problem is constrained to a finite alphabet: (i) in genome assembly, the matrix $A$ is fixed and consists of shifted identity matrices that select sub-vectors of $x^{*}$, and (ii) the permutation matrix of genome assembly is in fact a block permutation matrix that permutes sub-vectors instead of coordinates as in equation (1).

### 1.1 Related work

Previous work related to the observation model (1) can be broadly classified into two categories - those that focus on $x^{*}$ recovery, and those focussed on recovering the underlying permutation. We discuss the most relevant results below.

### 1.1.1 Latent vector estimation

The observation model (1) appears in the context of compressed sensing with an unknown sensor permutation EBDG14]. The authors consider the matrix-based observation model $Y=\Pi^{*} A X^{*}+W$, where $X^{*}$ is a matrix whose columns are composed of multiple unknown vectors. Their contributions include a branch and bound algorithm to recover the underlying $X^{*}$, which they show to perform well empirically for small instances under the setting in which the entries of the matrix $A$ are drawn i.i.d. from a Gaussian distribution.

In the context of pose and correspondence estimation, the paper [MSC09] considers the noiseless observation model (1), and shows that if the permutation matrix maps a sufficiently large number of positions to themselves, then $x^{*}$ can be recovered reliably.

In the context of molecular channels, the model (1) has been analyzed for the case when $x^{*}$ is some random vector, $A=I$, and $w$ represents non-negative noise that models delays introduced between emitter and receptor. Rose et al. [RMS12] provide lower bounds on the capacity of such channels. In particular, their results yield closed-form lower bounds for some special noise distributions, e.g., exponentially random noise.

A more recent paper UHV15 that is most closely related to our model considers the question of when the equation (11) has a unique solution $x^{*}$, i.e., the identifiability of the noiseless model. The authors show that if the entries of $A$ are sampled i.i.d. from any continuous distribution with $n \geq 2 d$, then equation (1) has a unique solution $x^{*}$ with probability 1 . They also provide a converse showing that if $n<2 d$, any matrix $A$ whose entries are sampled i.i.d. from a continuous distribution does not (with probability 1) have a unique solution $x^{*}$ to equation (11). While the paper shows uniqueness, the question of designing an efficient algorithm to recover a solution, unique or not, is left open. The paper also analyzes the stability of the noiseless solution, and establishes that $x^{*}$ can be recovered exactly when the SNR goes to infinity.

We also briefly compare the model (1) with the problem of vector recovery in unions of subspaces, studied widely in the compressive sensing literature [DD08, Blu11]. In the compressive sensing setup, the vector $x^{*}$ lies in the union of finitely many subspaces, and must be recovered from linear measurements with a random matrix, without a permutation. In our model, on the other hand, the vector $x^{*}$ is unrestricted, and the observation $y$ lies in the union of $n!$ subspaces - one for each permutation. While the two models share a superficial connection, results do not carry over from one to the other in any obvious way. In fact,
our model is fundamentally different from traditional compressive sensing, since the unknown permutation acts on the row space of the design matrix $A$. In contrast, restricting $x^{*}$ to a union of subspaces (or more specifically, restricting its sparsity) influences the column space of $A$.

### 1.1.2 Latent permutation estimation

While our paper seems to be the first to consider permutation recovery in the linear regression model (1), there are many related problems for which permutation recovery has been studied. We mention only those that are most closely related to our work.

The problem of feature matching in machine learning [D16 bears a superficial resemblance to our observation model. There, observations take the form $Y=X^{*}+W$ and $Y^{\prime}=\Pi^{*} X^{*}+W^{\prime}$, with all of $\left(X^{*}, Y, Y^{\prime}, W, W^{\prime}\right)$ representing matrices of appropriate dimensions, and the goal is to recover $\Pi^{*}$ from the tuple $\left(Y, Y^{\prime}\right)$. The paper [D16] establishes minimax rates on the separation between the rows of $X^{*}$ (as a function of problem parameters $n, d, \sigma)$ that allow for exact permutation recovery.

The problem of statistical seriation [FMR16] involves an observation model of the form $Y=\Pi^{*} X^{*}+W$, with the matrix $X^{*}$ obeying some shape constraint. In particular, if the columns of $X^{*}$ are unimodal (or, as a special case, monotone), then Flammarion et al. FMR16 establish minimax rates for the problem in the prediction error metric $\left\|\widehat{\Pi} \widehat{X}-\Pi^{*} X^{*}\right\|_{F}^{2}$ by analyzing the least squares estimator. The seriation problem was also considered by Fogel et al. [FJBd13] in the context of designing convex relaxations to permutation problems.

Permutation estimation has also been considered in other observation models involving matrices with structure, particularly in the context of ranking [SBGW15, Cha15, or even more generally, in the context of identity management [HGG09]. While we mention both of these problems because are related in spirit to permutation recovery, the problem setups do not bear too much resemblance to our linear model (1).

Algorithmic approaches to solving for $\Pi^{*}$ in equation (1) are related to the multi-dimensional assignment problem. In particular, while finding the correct permutation mapping between two vectors minimizing some loss function between them corresponds to the 1-dimensional assignment problem, here we are faced with an assignment problem between subspaces. While we do not elaborate on the vast literature that exists on solving variants on assignment problems, we note that broadly speaking, assignment problems in higher dimensions are much harder than the 1-D assignment problem. A survey on the quadratic assignment problem $\left[\mathrm{LdABN}^{+} 07\right]$ and references therein provide examples and methods that are currently used to solve these problems.

### 1.2 Contributions

Our primary contribution addresses permutation recovery in the noisy version of observation model (1), with a random design matrix $A$. In particular, when the entries of $A$ are drawn i.i.d. from a standard Gaussian matrix, we show sharp conditions on the SNR under which exact permutation recovery is possible. We also derive necessary conditions for approximate permutation recovery to within a prescribed Hamming distortion.

We also briefly address the computational aspect of the permutation recovery problem. We show that the information theoretically optimal estimator we propose for exact permutation recovery is NP-hard to compute in the worst case. For the special case of $d=1$, however, we
show that it can be computed in polynomial time. Our results are corroborated by numerical simulations.

### 1.3 Organization

The paper is organized as follows. In the next section, we set up notation and formally state the problem. In Section 3, we state our main results and discuss some of their implications. We provide proofs of the main results in Section 4, deferring the more technical lemmas to the appendices.

## 2 Background and problem setting

In this section, we set up notation, state the formal problem, and provide concrete examples of the noiseless version of our observation model by considering some fixed design matrices.

### 2.1 Notation

Since most of our analysis involves metrics involving permutations, we introduce all the relevant notation in this section. Permutations are denoted by $\pi$ and permutation matrices by $\Pi$. We use $\pi(i)$ to denote the image of an element $i$ under the permutation $\pi$. With a minor abuse of notation, we let $\mathcal{P}_{n}$ denote both the set of permutations on $n$ objects as well as the corresponding set of permutation matrices. We sometimes use the compact notation $y_{\pi}$ (or $y_{\Pi}$ ) to denote the vector $y$ with entries permuted according to the permutation $\pi$ (or $\Pi$ ).

We let $\mathrm{d}_{\mathrm{H}}\left(\pi, \pi^{\prime}\right)$ denote the Hamming distance between two permutations. More formally, we have $\mathrm{d}_{\mathrm{H}}\left(\pi, \pi^{\prime}\right):=\#\left\{i \mid \pi(i) \neq \pi^{\prime}(i)\right\}$. For convenience, we let $\mathrm{d}_{\mathrm{H}}\left(\Pi, \Pi^{\prime}\right)$ denote the Hamming distance between two permutation matrices, which is to be interpreted as the Hamming distance between the corresponding permutations.

The notation $v_{i}$ denotes the $i$ th entry of a vector $v$. We denote the $i$ th standard basis vector in $\mathbb{R}^{d}$ by $e_{i}$. We use the notation $a_{i}^{\top}$ to refer to the $i$ th row of $A$. We also use the standard shorthand notation $[n]:=\{1,2, \ldots, n\}$.

We also make use of standard asymptotic $\mathcal{O}$ notation. Specifically, for two real sequences $f_{n}$ and $g_{n}, f_{n}=\mathcal{O}\left(g_{n}\right)$ means that $f_{n} \leq C g_{n}$ for a universal constant $C>0$. Lastly, all logarithms denoted by $\log$ are to the base $e$, and we use $c_{1}, c_{2}$, etc. to denote absolute constants that are independent of other problem parameters.

### 2.2 Formal problem setting and permutation recovery

As mentioned in the introduction, we focus exclusively on the noisy observation model in the random design setting. In other words, we obtain an $n$-vector of observations $y$ from the model (1) with $n \geq d$ to ensure identifiability, and with the following assumptions:

Signal model The vector $x^{*} \in \mathbb{R}^{d}$ is fixed, but unknown. We note that this is different from the adversarial signal model of Unnikrishnan et al. [UHV15, and we provide clarifying examples in Section 2.3.

Measurement matrix The measurement matrix $A \in \mathbb{R}^{n \times d}$ is a random matrix of i.i.d. standard Gaussian variables chosen without knowledge of $x^{*}$. Our assumption on i.i.d. standard Gaussian designs easily extends to accommodate the more general case when rows of
$A$ are drawn i.i.d. from the distribution $\mathcal{N}(0, \Sigma)$. In particular, writing $A=W \sqrt{\Sigma}$, where $W$ in an $n \times d$ standard Gaussian matrix and $\sqrt{\Sigma}$ denotes the symmetric square root of the (non-singular) covariance matrix $\Sigma$, our observation model takes the form

$$
y=\Pi^{*} W \sqrt{\Sigma} x^{*}+w,
$$

and the unknown vector is now $\sqrt{\Sigma} x^{*}$ in the model (1).
Noise variables The vector $w \sim \mathcal{N}\left(0, \sigma^{2} I_{n}\right)$ represents uncorrelated noise variables, each of (possibly unknown) variance $\sigma^{2}$. As will be made clear in the analysis, our assumption that the noise is Gaussian also readily extends to accommodate i.i.d. $\sigma$-sub-Gaussian noise. Additionally, the permutation noise represented by the unknown permutation matrix $\Pi^{*}$ is arbitrary.

The main recovery criterion addressed in this paper is that of exact permutation recovery, which is formally described below. Following that, we also discuss two other relevant recovery criteria.

Exact permutation recovery The problem of exact permutation recovery is to recover $\Pi^{*}$, and the risk of an estimator is evaluated on the $0-1$. More formally, given an estimator of $\Pi^{*}$ denoted by $\widehat{\Pi}:(y, A) \rightarrow \mathcal{P}_{n}$, we evaluate its risk by

$$
\begin{equation*}
\operatorname{Pr}\left\{\widehat{\Pi} \neq \Pi^{*}\right\}=\mathbb{E}\left[\mathbf{1}\left\{\widehat{\Pi} \neq \Pi^{*}\right\}\right], \tag{2}
\end{equation*}
$$

where the probability in the LHS is taken over the randomness in $y$ induced by both $A$ and $w$.

Approximate permutation recovery It is reasonable to think that recovering $\Pi^{*}$ up to some distortion is sufficient for many applications. Such a relaxation of exact permutation recovery allows the estimator to output a $\widehat{\Pi}$ such that $\mathrm{d}_{\mathrm{H}}\left(\widehat{\Pi}, \Pi^{*}\right) \leq D$, for some distortion $D$ to be specified. The risk of such an estimator is again evaluated on the 0-1 loss of this error metric, given by $\operatorname{Pr}\left\{\mathrm{d}_{\mathrm{H}}\left(\widehat{\Pi}, \Pi^{*}\right) \geq D\right\}$, with the probability again taken over both $A$ and $w$. While our results are derived mainly in the context of exact permutation recovery, they can be suitably modified to also yields results for approximate permutation recovery.

Recovery with side information In this variation, the unknown permutation matrix is not arbitrary, but known to be in some Hamming ball around the identity matrix. In other words, the estimator is provided with side information that $\mathrm{d}_{\mathrm{H}}\left(\Pi^{*}, I\right) \leq \bar{h}$, for some $\bar{h}<n$. In many applications, this may constitute a prior that leads us to believe that the permutation matrix is not arbitrary. In multi-target tracking, for example, we may be sure that at any given time, a certain number of measurements correspond to the true sensors that made them (that are close to the target, perhaps). Our results also address the exact permutation recovery problem with side information.

We now briefly give some examples in which the noiseless version is identifiable.

### 2.3 Illustrative examples of the noiseless model

In this section, we present two examples to illustrate the problem of permutation recovery and highlight the difference between our signal model and that of Unnikrishnan et al. UHV15.

Example 1 Consider the noiseless case of the observation model (1). Let $\nu_{i}, \nu_{i}^{\prime} \quad(i=1,2, \ldots, d)$ represent i.i.d. continuous random variables, and form the $\operatorname{design}$ matrix $A$ by choosing

$$
a_{2 i-1}^{\top}:=\nu_{i} e_{i}^{\top} \text { and } a_{2 i}^{\top}=\nu_{i}^{\prime} e_{i}^{\top}, \quad i=1,2, \ldots, d .
$$

Note that $n=2 d$. Now consider our fixed but unknown signal model for $x^{*}$. Since the permutation is arbitrary, our observations can be thought of as the unordered set $\left\{\nu_{i} x_{i}^{*}, \nu_{i}^{\prime} x_{i}^{*} \mid i \in[d]\right\}$. With probability 1 , the ratios $r_{i}:=\nu_{i} / \nu_{i}^{\prime}$ are distinct for each $i$, and also that $\nu_{i} x_{i}^{*} \neq \nu_{j} x_{j}^{*}$ with probability 1 , by assumption of a fixed $x^{*}$. Therefore, there is a one to one correspondence between the ratios $r_{i}$ and $x_{i}^{*}$. All ratios are computable in time $\mathcal{O}\left(n^{2}\right)$, and $x^{*}$ can be exactly recovered. Using this information, we can also exactly recover $\Pi^{*}$.

Example 2 A particular case of this example was already observed by Unnikrishnan et al. [UHV15], but we include it to illustrate the difference between our signal model and the adversarial signal model. Form the fixed design matrix $A$ by including $2^{i-1}$ copies of the vector $e_{i}$ among its rows. We therefore ${ }^{1}$ have $n=\sum_{i=1}^{d} 2^{i-1}=2^{d}-1$.

Our observations therefore consist of $2^{i-1}$ repetitions of $x_{i}^{*}$ for each $i \in[d]$. The value of $x_{i}^{*}$ can therefore be recovered by simply counting the number of times it is repeated, with our choice of the number of repetitions also accounting for cases when $x_{i}^{*}=x_{j}^{*}$ for some $i \neq j$. Notice that we can now recover any vector $x^{*}$, even those chosen adversarially with knowledge of the $A$ matrix. Therefore, such a design matrix allows for an adversarial signal model, in the flavor of compressive sensing CT06.

Having provided examples of the noiseless observation model, we now return to the noisy setting of Section 2.2, and state our main results.

## 3 Main results

In this section, we state our main theorems and discuss their consequences. Proofs of the theorems can be found in Section 4 .

### 3.1 Statistical limits of exact permutation recovery

Our main theorems in this section provide necessary and sufficient conditions under which the probability of error in exactly recovering the true permutation goes to zero.

In brief, provided that $d$ is sufficiently small, we establish a threshold phenomenon that characterizes how the signal-to-noise ratio snr $:=\frac{\left\|x^{*}\right\|_{2}^{2}}{\sigma^{2}}$ must scale relative to $n$ in order to ensure identifiability. More specifically, defining the ratio

$$
\Gamma(n, \mathrm{snr}):=\frac{\log (1+\mathrm{snr})}{\log n},
$$

we show that the maximum likelihood estimator recovers the true permutation with high probability provided $\Gamma(n, \mathrm{snr}) \gg c$, where $c$ denotes an absolute constant. Conversely, if $\Gamma(n, \mathrm{snr}) \ll c$, then exact permutation recovery is impossible. For illustration, we have plotted the behaviour of the maximum likelihood estimator for the case when $d=1$ in Figure 2 . Evidently, there is a sharp phase transition between error and exact recovery as the ratio $\Gamma(n, \mathrm{snr})$ varies from 3 to 5 .

[^0]

Figure 2. Empirical frequency of the event $\left\{\widehat{\Pi}_{M L}=\Pi^{*}\right\}$ over 1000 independent trials with $d=1$, plotted against $\Gamma(n$, snr $)$ for different values of $n$. The probability of successful permutation recovery undergoes a phase transition as $\Gamma$ ( $n, \mathrm{snr}$ ) varies from 3 to 5 . This is consistent with the prediction of Theorems 1 and 2

Let us now turn to more precise statements of our results. We first define the maximum likelihood estimator (MLE) as

$$
\begin{equation*}
\left(\widehat{\Pi}_{\mathrm{ML}}, \widehat{x}_{\mathrm{ML}}\right)=\arg \min _{\substack{\Pi \in \mathcal{P}_{n} \\ x \in \mathbb{R}^{d}}}\|y-\Pi A x\|_{2}^{2} \tag{3}
\end{equation*}
$$

The following theorem provides an upper bound on the probability of error of $\widehat{\Pi}_{\mathrm{ML}}$, with $\left(c_{1}, c_{2}\right)$ denoting absolute constants.

Theorem 1. For any $d<n$ and $\epsilon<\sqrt{n}$, if

$$
\begin{equation*}
\log \left(\frac{\left\|x^{*}\right\|_{2}^{2}}{\sigma^{2}}\right) \geq\left(c_{1} \frac{n}{n-d}+\epsilon\right) \log n \tag{4}
\end{equation*}
$$

then $\operatorname{Pr}\left\{\widehat{\Pi}_{\mathrm{ML}} \neq \Pi^{*}\right\} \leq c_{2} n^{-2 \epsilon}$.
Theorem 1 provides conditions on the signal-to-noise ratio snr $=\frac{\left\|x^{*}\right\|_{2}^{2}}{\sigma^{2}}$ that are sufficient for permutation recovery in the non-asymptotic, noisy regime. In contrast, the results of Unnikrishnan et al. [UHV15] are stated in the limit snr $\rightarrow \infty$, without an explicit characterization of the scaling behavior.

We also note that Theorem 1 holds for all values of $d<n$, whereas the results of Unnikrishnan et al. UHV15 require $n \geq 2 d$ for identifiability of $x^{*}$ in the noiseless case. Although the recovery of $\Pi^{*}$ and $x^{*}$ are not directly comparable, it is worth pointing out that the discrepancy also arises due to the difference between our fixed and unknown signal model, and the adversarial signal model assumed in the paper UHV15.

We now turn to the following converse result, which complements Theorem 1.

Theorem 2. For any $\delta \in(0,2)$, if

$$
\begin{equation*}
2+\log \left(1+\frac{\left\|x^{*}\right\|_{2}^{2}}{\sigma^{2}}\right) \leq(2-\delta) \log n \tag{5}
\end{equation*}
$$

then $\operatorname{Pr}\left\{\widehat{\Pi} \neq \Pi^{*}\right\} \geq 1-c_{3} e^{-c_{4} n \delta}$ for any estimator $\widehat{\Pi}$.
Theorem 2 serves as a "strong converse" for our problem, since it guarantees that if condition (5) is satisfied, then the probability of error of any estimator goes to 1 as $n$ goes to infinity. Indeed, it is proved using the strong converse argument for the Gaussian channel Sha59. In fact, we are also able to show the following "weak converse" in the presence of side information.

Proposition 1. If $n \geq 9$ and

$$
\log \left(1+\frac{\left\|x^{*}\right\|_{2}^{2}}{\sigma^{2}}\right) \leq \frac{8}{9} \log \left(\frac{n}{8}\right)
$$

then $\operatorname{Pr}\left\{\widehat{\Pi} \neq \Pi^{*}\right\} \geq 1 / 2$ for any estimator $\widehat{\Pi}$, even if it is known a-priori that $\mathrm{d}_{\mathrm{H}}\left(\Pi^{*}, I\right) \leq 2$.
As mentioned earlier, restriction of $\Pi^{*}$ constitutes some application-dependent prior; the strongest such prior restricts it to a Hamming ball of radius 2 around the identity. Proposition 1 asserts that even this side information does not substantially change the statistical limits of permutation recovery.

It is also worth noting that the converse results hold uniformly over $d$. In particular, if $d \leq p n$ for some fixed $p<1$, Theorems 1 and 2 together yield the threshold behavior of identifiability in terms of $\Gamma(n, \mathrm{snr})$ that was discussed above. In the next section, we find that a similar phenomenon occurs even with approximate permutation recovery.

### 3.2 Limits of approximate permutation recovery

The techniques we used to prove results for exact permutation recovery can be suitably modified to obtain results for approximate permutation recovery to within a Hamming distortion $D$. In particular, we show the following converse result for approximate recovery.

Theorem 3. For any $2<D \leq n-1$, if

$$
\begin{equation*}
\log \left(1+\frac{\left\|x^{*}\right\|_{2}^{2}}{\sigma^{2}}\right) \leq \frac{n-D+1}{n} \log \left(\frac{n-D+1}{2 e}\right) \tag{6}
\end{equation*}
$$

then $\operatorname{Pr}\left\{\mathrm{d}_{\mathrm{H}}\left(\widehat{\Pi}, \Pi^{*}\right) \geq D\right\} \geq 1 / 2$ for any estimator $\widehat{\Pi}$.
Note that for any $D \leq p n$ with $p \in(0,1)$, Theorems 1 and 3 provide a set of sufficient and necessary conditions for approximate permutation recovery that match up to constant factors. In particular, the necessary condition resembles that for exact permutation recovery, and the same SNR threshold behaviour is seen even here. We remark that a corresponding converse with side information can also be proved for approximate permutation recovery using techniques similar to the proof of Proposition 1. It is also worth mentioning the following:

Remark 1. The converse results given by Theorem 2, Proposition 1, and Theorem 3 hold even when the estimator has exact knowledge of $x^{*}$.

### 3.3 Computational aspects

In the previous sections, we considered the MLE given by equation (3) and analyzed its statistical properties. However, since equation (3) involves a combinatorial minimization over $n$ ! permutations, it is unclear if $\widehat{\Pi}_{\mathrm{ML}}$ can be computed efficiently. The following theorem addresses this question.
Theorem 4. For $d=1$, the MLE $\widehat{\Pi}_{\mathrm{ML}}$ can be computed in time $\mathcal{O}(n \log n)$ for any choice of the measurement matrix $A$. In contrast, if $d>1$, then $\widehat{\Pi}_{\mathrm{ML}}$ is NP-hard to compute.

The algorithm used to prove the first part of the theorem involves a simple sorting operation, which introduces the $\mathcal{O}(n \log n)$ complexity. We emphasize that the algorithm assumes no prior knowledge about the distribution of the data; for every given $A$ and $y$, it returns the optimal solution to problem (3).

The second part of the theorem asserts that the algorithmic simplicity enjoyed by the $d=1$ case does not extend to general $d$. The proof proceeds by a reduction from the NP-complete partition problem. We stress here that the NP-hardness claim holds over worst case input instances. In particular, it does not preclude the possibility that there exists a polynomial time algorithm that solves problem (3) with high probability when $A$ is chosen randomly as in our original setting. However, we conjecture that solving problem (3) over random $A$ is also a computationally hard problem, conditioned on an average-case hardness assumption.

## 4 Proofs of Main Results

In this section, we prove our main results. Technical details are deferred to the appendices. Throughout the proofs, we assume that $n$ is larger than some universal constant. The case where $n$ is smaller can be handled by changing the constants in our proofs appropriately. We also use the notation $c, c^{\prime}$ to denote absolute constants that can change from line to line.

We begin with the proof of Theorem1. At a high level, it involves bounding the probability that any fixed permutation is preferred to $\Pi^{*}$ by the estimator. The analysis requires precise control on the lower tails of $\chi^{2}$-random variables, and tight bounds on the norms of random projections, for which we use results derived in the context of dimensionality reduction by Dasgupta and Gupta DG03.

In order to simplify the exposition, we first consider the case when $d=1$ in Section 4.1, and later make the necessary modifications for the general case in Section 4.2.

### 4.1 Proof of Theorem 1: $d=1$ case

Recall the definition of the maximum likelihood estimator

$$
\left(\widehat{\Pi}_{\mathrm{ML}}, \widehat{x}_{\mathrm{ML}}\right)=\arg \min _{\Pi \in \mathcal{P}_{n}} \min _{x \in \mathbb{R}^{d}}\|y-\Pi A x\|_{2}^{2}
$$

For a fixed permutation matrix $\Pi$, assuming that $A$ has full column rank 2 , the minimizing argument $x$ is simply $(\Pi A)^{\dagger} y$, where $X^{\dagger}=\left(X^{\top} X\right)^{-1} X^{\top}$ represents the pseudoinverse of a matrix $X$. By computing the minimum over $x \in \mathbb{R}^{d}$ in the above equation, we find that the maximum likelihood estimate of the permutation is given by

$$
\begin{equation*}
\widehat{\Pi}_{\mathrm{ML}}=\arg \min _{\Pi \in \mathcal{P}_{n}}\left\|P_{\Pi}^{\perp} y\right\|_{2}^{2}, \tag{7}
\end{equation*}
$$

[^1]where $P_{\Pi}^{\perp}=I-\Pi A\left(A^{\top} A\right)^{-1}(\Pi A)^{\top}$ denotes the projection onto the orthogonal complement of the column space of $\Pi A$.

For a fixed $\Pi \in \mathcal{P}_{n}$, define the random variable

$$
\begin{equation*}
\Delta\left(\Pi, \Pi^{*}\right):=\left\|P_{\Pi}^{\perp} y\right\|_{2}^{2}-\left\|P_{\Pi^{*}}^{\perp} y\right\|_{2}^{2} . \tag{8}
\end{equation*}
$$

For any permutation $\Pi$, the estimator (7) prefers the permutation $\Pi$ to $\Pi^{*}$ if $\Delta\left(\Pi, \Pi^{*}\right) \leq 0$. The overall error event occurs when $\Delta\left(\Pi, \Pi^{*}\right) \leq 0$ for some $\Pi$, meaning that

$$
\begin{equation*}
\left\{\widehat{\Pi}_{\mathrm{ML}} \neq \Pi^{*}\right\}=\bigcup_{\Pi \in \mathcal{P}_{n} \backslash \Pi^{*}}\left\{\Delta\left(\Pi, \Pi^{*}\right) \leq 0\right\} \tag{9}
\end{equation*}
$$

Equation (9) holds for any value of $d$. We shortly specialize to the $d=1$ case. Our strategy for proving Theorem 1 boils down to bounding the probability of each error event in the RHS of equation (9) using the following key lemma, proved in Section 4.1.1. Recall the definition of $\mathrm{d}_{\mathrm{H}}\left(\Pi, \Pi^{\prime}\right)$, the Hamming distance between two permutation matrices.

Lemma 1. For $d=1$ and any two permutation matrices $\Pi$ and $\Pi^{*}$, and provided $\frac{\left\|x^{*}\right\|_{2}^{2}}{\sigma^{2}}>1$, we have

$$
\operatorname{Pr}\left\{\Delta\left(\Pi, \Pi^{*}\right) \leq 0\right\} \leq c^{\prime} \exp \left(-c \mathrm{~d}_{\mathrm{H}}\left(\Pi, \Pi^{*}\right) \log \left(\frac{\left\|x^{*}\right\|_{2}^{2}}{\sigma^{2}}\right)\right)
$$

We are now ready to prove Theorem 1 .

Proof of Theorem 1 for $d=1$. Fix $\epsilon>0$ and assume that the following consequence of condition (4) holds:

$$
\begin{equation*}
c \log \left(\frac{\left\|x^{*}\right\|_{2}^{2}}{\sigma^{2}}\right) \geq(1+\epsilon) \log n \tag{10}
\end{equation*}
$$

where $c$ is the same as in Lemma 1. Now, observe that

$$
\begin{aligned}
\operatorname{Pr}\left\{\widehat{\Pi}_{\mathrm{ML}} \neq \Pi^{*}\right\} & \leq \sum_{\Pi \in \mathcal{P}_{n} \backslash \Pi^{*}} \operatorname{Pr}\left\{\Delta\left(\Pi, \Pi^{*}\right) \leq 0\right\} \\
& \stackrel{(\mathrm{i})}{\leq} \sum_{\Pi \in \mathcal{P}_{n} \backslash \Pi^{*}} c^{\prime} \exp \left(-c \mathrm{~d}_{\mathrm{H}}\left(\Pi, \Pi^{*}\right) \log \left(\frac{\left\|x^{*}\right\|_{2}^{2}}{\sigma^{2}}\right)\right) \\
& \leq c^{\prime} \sum_{2 \leq k \leq n} n^{k} \exp \left(-c k \log \left(\frac{\left\|x^{*}\right\|_{2}^{2}}{\sigma^{2}}\right)\right) \\
& \stackrel{(i i)}{\leq} c^{\prime} \sum_{2 \leq k \leq n} n^{-\epsilon k} \\
& \leq c^{\prime} \frac{1}{n^{\epsilon}\left(n^{\epsilon}-1\right)}
\end{aligned}
$$

where step (i) follows since $\#\left\{\Pi: \mathrm{d}_{\mathrm{H}}\left(\Pi, \Pi^{*}\right)=k\right\} \leq n^{k}$, and step (ii) follows from condition (10). Relabelling the constants in condition (10) proves the theorem.

### 4.1.1 Proof of Lemma 1

Before the proof, we establish notation. For each $\delta>0$, define the events

$$
\begin{align*}
& \mathcal{F}_{1}(\delta)=\left\{\left|\left\|P_{\Pi^{*}}^{\perp} y\right\|_{2}^{2}-\left\|P_{\Pi}^{\perp} w\right\|_{2}^{2}\right| \geq \delta\right\}, \text { and }  \tag{11a}\\
& \mathcal{F}_{2}(\delta)=\left\{\left\|P_{\Pi}^{\perp} y\right\|_{2}^{2}-\left\|P_{\Pi}^{\perp} w\right\|_{2}^{2} \leq 2 \delta\right\} \tag{11b}
\end{align*}
$$

Evidently,

$$
\begin{equation*}
\left\{\Delta\left(\Pi, \Pi^{*}\right) \leq 0\right\} \subseteq \mathcal{F}_{1}(\delta) \cup \mathcal{F}_{2}(\delta) \tag{12}
\end{equation*}
$$

Indeed, if neither $\mathcal{F}_{1}(\delta)$ nor $\mathcal{F}_{2}(\delta)$ occurs

$$
\Delta\left(\Pi, \Pi^{*}\right)=\left(\left\|P_{\Pi}^{\perp} y\right\|_{2}^{2}-\left\|P_{\Pi}^{\perp} w\right\|_{2}^{2}\right)-\left(\left\|P_{\Pi^{*}}^{\perp} y\right\|_{2}^{2}-\left\|P_{\Pi}^{\perp} w\right\|_{2}^{2}\right)>2 \delta-\delta=\delta
$$

Thus, to prove Lemma 1 , we shall bound the probability of the two events $\mathcal{F}_{1}(\delta)$ and $\mathcal{F}_{2}(\delta)$ individually, and then invoke the union bound. Note that inequality $\sqrt{12}$ holds for all values of $\delta>0$; it is convenient to choose $\delta^{*}:=\frac{1}{3}\left\|P_{\Pi}^{\perp} \Pi^{*} A x^{*}\right\|_{2}^{2}$. With this choice, the following lemma, proved in Appendix $B$, bounds the probabilities of the individual events over randomness in $w$ conditioned on a given $A$.
Lemma 2. For any $\delta>0$ and with $\delta^{*}=\frac{1}{3}\left\|P_{\Pi}^{\perp} \Pi^{*} A x^{*}\right\|_{2}^{2}$, we have

$$
\begin{align*}
\operatorname{Pr}_{w}\left\{\mathcal{F}_{1}(\delta)\right\} & \leq c^{\prime} \exp \left(-c \frac{\delta}{\sigma^{2}}\right), \text { and }  \tag{13a}\\
\operatorname{Pr}_{w}\left\{\mathcal{F}_{2}\left(\delta^{*}\right)\right\} & \leq c^{\prime} \exp \left(-c \frac{\delta^{*}}{\sigma^{2}}\right) \tag{13b}
\end{align*}
$$

The next lemma, proved in Section 4.1.2, is needed in order to incorporate the randomness in $A$ into the required tail bound. It is convenient to introduce the shorthand $T_{\Pi}:=\left\|P_{\Pi}^{\perp} \Pi^{*} A x^{*}\right\|_{2}^{2}$.

Lemma 3. For $d=1$ and any two permutation matrices $\Pi$ and $\Pi^{*}$ at Hamming distance $h$, we have

$$
\begin{equation*}
\operatorname{Pr}_{A}\left\{T_{\Pi} \leq t\left\|x^{*}\right\|_{2}^{2}\right\} \leq 6 \exp \left(-\frac{h}{10}\left[\log \frac{h}{t}+\frac{t}{h}-1\right]\right) \tag{14}
\end{equation*}
$$

for all $t \in[0, h]$.
We now have all the ingredients to prove Lemma 1 .
Proof of Lemma 1. Applying Lemma 2 and using the union bound then yields

$$
\begin{align*}
\operatorname{Pr}_{w}\left\{\Delta\left(\Pi, \Pi^{*}\right) \leq 0\right\} & \leq \operatorname{Pr}_{w}\left\{\mathcal{F}_{1}\left(\delta^{*}\right)\right\}+\operatorname{Pr}_{w}\left\{\mathcal{F}_{2}\left(\delta^{*}\right)\right\} \\
& \leq c^{\prime} \exp \left(-c \frac{T_{\Pi}}{\sigma^{2}}\right) \tag{15}
\end{align*}
$$

Combining bound 15 with Lemma 3 yields

$$
\begin{align*}
\operatorname{Pr}\left\{\Delta\left(\Pi, \Pi^{*}\right) \leq 0\right\} & \leq c^{\prime} \exp \left(-c \frac{t\left\|x^{*}\right\|_{2}^{2}}{\sigma^{2}}\right) \operatorname{Pr}_{A}\left\{T_{\Pi} \geq t\left\|x^{*}\right\|_{2}^{2}\right\}+\operatorname{Pr}_{A}\left\{T_{\Pi} \leq t\left\|x^{*}\right\|_{2}^{2}\right\} \\
& \leq c^{\prime} \exp \left(-c \frac{t\left\|x^{*}\right\|_{2}^{2}}{\sigma^{2}}\right)+6 \exp \left(-\frac{h}{10}\left[\log \frac{h}{t}+\frac{t}{h}-1\right]\right) \tag{16}
\end{align*}
$$

where the last inequality holds provided that $t \in[0, h]$, and the probability in the LHS is now taken over randomness in both $w$ and $A$.

Using the shorthand snr $:=\frac{\left\|x^{*}\right\|_{2}^{2}}{\sigma^{2}}$, setting $t=h \frac{\log \text { snr }}{\mathrm{snr}}$, and noting that $t \in[0, h]$ since snr $>1$, we have

$$
\begin{equation*}
\operatorname{Pr}\left\{\Delta\left(\Pi, \Pi^{*}\right) \leq 0\right\} \leq c^{\prime} \exp (-c h \log \operatorname{snr})+6 \exp \left(-\frac{h}{10}\left[\log \left(\frac{\mathrm{snr}}{\log \operatorname{snr}}\right)+\frac{\log \operatorname{snr}}{\mathrm{snr}}-1\right]\right) \tag{17}
\end{equation*}
$$

It is easily verified that for all snr $>1$, we have

$$
\begin{equation*}
\log \left(\frac{\mathrm{snr}}{\log s n r}\right)+\frac{\log \mathrm{snr}}{\mathrm{snr}}-1>\frac{\log \mathrm{snr}}{4} . \tag{18}
\end{equation*}
$$

Hence, after substituting for snr, we have

$$
\begin{equation*}
\operatorname{Pr}\left\{\Delta\left(\Pi, \Pi^{*}\right) \leq 0\right\} \leq c^{\prime} \exp \left(-c h \log \left(\frac{\left\|x^{*}\right\|_{2}^{2}}{\sigma^{2}}\right)\right) \tag{19}
\end{equation*}
$$

### 4.1.2 Proof of Lemma 3

In the case $d=1$, the matrix $A$ is composed of a single vector $a \in \mathbb{R}^{n}$. Recalling the random variable $T_{\Pi}=\left\|P_{\Pi}^{\perp} \Pi^{*} A x^{*}\right\|_{2}^{2}$, we have

$$
\begin{aligned}
T_{\Pi} & =\left(x^{*}\right)^{2}\left(\|a\|_{2}^{2}-\frac{1}{\|a\|_{2}^{2}}\left\langle a_{\Pi}, a\right\rangle^{2}\right) \\
& \stackrel{(\mathrm{i})}{\geq}\left(x^{*}\right)^{2}\left(\|a\|_{2}^{2}-\left|\left\langle a, a_{\Pi}\right\rangle\right|\right) \\
& =\frac{\left(x^{*}\right)^{2}}{2} \min \left(\left\|a-a_{\Pi}\right\|_{2}^{2},\left\|a+a_{\Pi}\right\|_{2}^{2}\right),
\end{aligned}
$$

where step (i) follows from the Cauchy Schwarz inequality. Applying the union bound then yields

$$
\operatorname{Pr}\left\{T_{\Pi} \leq t\left(x^{*}\right)^{2}\right\} \leq \operatorname{Pr}\left\{\left\|a-a_{\Pi}\right\|_{2}^{2} \leq 2 t\right\}+\operatorname{Pr}\left\{\left\|a+a_{\Pi}\right\|_{2}^{2} \leq 2 t\right\}
$$

Let $Z_{\ell}$ and $\tilde{Z}_{\ell}$ denote (not necessarily independent) $\chi^{2}$ random variables with $\ell$ degrees of freedom. We split the analysis into two cases.

Case $h \geq 3$ : Lemma 8 from Appendix A. 1 guarantees that

$$
\begin{align*}
& \frac{\left\|a-a_{\Pi}\right\|_{2}^{2}}{2} \stackrel{d}{=} Z_{h_{1}}+Z_{h_{2}}+Z_{h_{3}}, \text { and }  \tag{20a}\\
& \frac{\left\|a+a_{\Pi}\right\|_{2}^{2}}{2} \stackrel{d}{=} \widetilde{Z}_{h_{1}}+\widetilde{Z}_{h_{2}}+\widetilde{Z}_{h_{3}}+\widetilde{Z}_{n-h} \tag{20b}
\end{align*}
$$

where $\stackrel{d}{=}$ denotes equality in distribution and $h_{1}, h_{2}, h_{3} \geq \frac{h}{5}$ with $h_{1}+h_{2}+h_{3}=h$. An application of the union bound then yields

$$
\operatorname{Pr}\left\{\left\|a-a_{\Pi}\right\|_{2}^{2} \leq 2 t\right\} \leq \sum_{i=1}^{3} \operatorname{Pr}\left\{Z_{h_{i}} \leq t \frac{h_{i}}{h}\right\} .
$$

Similarly, provided that $h \geq 3$, we have

$$
\begin{aligned}
\operatorname{Pr}\left\{\left\|a+a_{\Pi}\right\|_{2}^{2} \leq 2 t\right\} & \leq \operatorname{Pr}\left\{\widetilde{Z}_{h_{1}}+\widetilde{Z}_{h_{2}}+\widetilde{Z}_{h_{3}}+\widetilde{Z}_{n-h} \leq t\right\} \\
& \stackrel{\text { (ii) }}{\leq} \operatorname{Pr}\left\{\widetilde{Z}_{h_{1}}+\widetilde{Z}_{h_{2}}+\widetilde{Z}_{h_{3}} \leq t\right\} \\
& (\text { (iii) } \\
\leq & \sum_{i=1}^{3} \operatorname{Pr}\left\{\widetilde{Z}_{h_{i}} \leq t \frac{h_{i}}{h}\right\}
\end{aligned}
$$

where inequality (ii) follows from the non-negativity of $Z_{n-h}$, and the monotonicity of the CDF; and inequality (iii) from the union bound. Finally, bounds on the lower tails of $\chi^{2}$ random variables (see Lemma 9 in Appendix A.2) yield

$$
\begin{aligned}
\operatorname{Pr}\left\{Z_{h_{i}} \leq t \frac{h_{i}}{h}\right\}=\operatorname{Pr}\left\{\widetilde{Z}_{h_{i}} \leq t \frac{h_{i}}{h}\right\} & \stackrel{\text { (iv) }}{\leq}\left(\frac{t}{h} \exp \left(1-\frac{t}{h}\right)\right)^{h_{i} / 2} \\
& \stackrel{(\mathrm{v})}{\leq}\left(\frac{t}{h} \exp \left(1-\frac{t}{h}\right)\right)^{h / 10}
\end{aligned}
$$

Here, inequality (iv) is valid provided $\frac{t h_{i}}{h} \leq h_{i}$, or equivalently, if $t \leq h$, whereas inequality (v) follows since $h_{i} \geq h / 5$ and the function $x e^{1-x} \in[0,1]$ for all $x \in[0,1]$. Combining the pieces proves Lemma 3 for $h \geq 3$.

Case $h=2$ : In this case, we have

$$
\frac{\left\|a-a_{\Pi}\right\|_{2}^{2}}{2} \stackrel{d}{=} 2 Z_{1}, \quad \text { and } \quad \frac{\left\|a+a_{\Pi}\right\|_{2}^{2}}{2} \stackrel{d}{=} 2 \widetilde{Z}_{1}+\widetilde{Z}_{n-2}
$$

Proceeding as before by applying the union bound and Lemma 9, we have that for $t \leq 2$, the random variable $T_{\Pi}$ obeys the tail bound

$$
\operatorname{Pr}\left\{T_{\Pi} \leq t\left(x^{*}\right)^{2}\right\} \leq 2\left(\frac{t}{2} \exp \left(1-\frac{t}{2}\right)\right)^{1 / 2} \leq 6\left(\frac{t}{h} \exp \left(1-\frac{t}{h}\right)\right)^{h / 10}, \text { for } h=2
$$

In the next section, we prove Theorem 1 for the general case.

### 4.2 Proof of Theorem 1: Case $d \in\{2,3, \ldots, n-1\}$

In order to be consistent, we follow the same proof structure as for the $d=1$ case. Recall the definition of $\Delta\left(\Pi, \Pi^{*}\right)$ from equation (8). We begin with an equivalent of the key lemma to bound the probability of the event $\left\{\Delta\left(\Pi, \Pi^{*}\right) \leq 0\right\}$.

Lemma 4. For any $1<d<n$, any two permutation matrices $\Pi$ and $\Pi^{*}$ at Hamming distance $h$, and provided $\left(\frac{\left\|x^{*}\right\|_{2}^{2}}{\sigma^{2}}\right) n^{-\frac{2 n}{n-d}}>\frac{5}{4}$, we have

$$
\begin{equation*}
\operatorname{Pr}\left\{\Delta\left(\Pi, \Pi^{*}\right) \leq 0\right\} \leq c^{\prime} \max \left[\exp \left(-n \log \frac{n}{2}\right), \exp \left(\operatorname{ch}\left(\log \left(\frac{\left\|x^{*}\right\|_{2}^{2}}{\sigma^{2}}\right)-\frac{2 n}{n-d} \log n\right)\right)\right] \tag{21}
\end{equation*}
$$

We prove Lemma 4 in Section 4.2.1. Taking it as given, we are ready to prove Theorem 1 for the general case.

Proof of Theorem 1, general case. As before, we use the union bound to prove the theorem. We begin by fixing some $\epsilon \in(0, \sqrt{n})$ and assuming that the following consequence of condition (4) holds:

$$
\begin{equation*}
c \log \left(\frac{\left\|x^{*}\right\|_{2}^{2}}{\sigma^{2}}\right) \geq\left(1+\epsilon+c \frac{2 n}{n-d}\right) \log n \tag{22}
\end{equation*}
$$

Now define $b(k):=\sum_{\Pi: \mathrm{d}_{\mathrm{H}}\left(\Pi, \Pi^{*}\right)=k} \operatorname{Pr}\left\{\Delta\left(\Pi, \Pi^{*}\right) \leq 0\right\}$. Applying Lemma 4 then yields

$$
\begin{equation*}
b(k) \leq \frac{n!}{(n-k)!} c^{\prime} \max \left\{\exp \left(-n \log \frac{n}{2}\right), \exp \left(-c k\left(\log \left(\frac{\left\|x^{*}\right\|_{2}^{2}}{\sigma^{2}}\right)-\frac{2 n}{n-d} \log n\right)\right)\right\} \tag{23}
\end{equation*}
$$

We upper bound $b(k)$ by splitting the analysis into two cases.
Case 1: If the first term attains the maximum in the RHS of inequality (23), then for all $2 \leq k \leq n$, we have

$$
\begin{aligned}
& b(k) \leq c^{\prime} n!\exp (-n \log n+n \log 2) \\
& \quad \text { (i) } \\
& \leq c^{\prime} e \sqrt{n} \exp (-n \log n+n \log 2+-n+n \log n) \\
& \text { (ii) } \frac{c^{\prime}}{\leq} n^{2 \epsilon+1}
\end{aligned}
$$

where inequality (i) follows from the well-known upper bound $n!\leq e \sqrt{n}\left(\frac{n}{e}\right)^{n}$, and inequality (ii) holds since $\epsilon \in(0, \sqrt{n})$.

Case 2: Alternatively, if the maximum is attained by the second term in the RHS of inequality (23), then we have

$$
\begin{aligned}
& b(k) \leq n^{k} c^{\prime} \exp \left(-c k\left(\log \left(\frac{\left\|x^{*}\right\|_{2}^{2}}{\sigma^{2}}\right)-\frac{2 n}{n-d} \log n\right)\right) \\
& \quad \stackrel{\text { (iii) }}{\leq} c^{\prime} n^{-\epsilon k}
\end{aligned}
$$

where step (iii) follows from condition (22).
Combining the two cases, we have

$$
b(k) \leq \max \left\{c^{\prime} n^{-\epsilon h}, c n^{-2 \epsilon-1}\right\} \leq\left(c^{\prime} n^{-\epsilon h}+c n^{-2 \epsilon-1}\right)
$$

The last step is to use the union bound to obtain

$$
\begin{align*}
\operatorname{Pr}\left\{\widehat{\Pi}_{\mathrm{ML}} \neq \Pi^{*}\right\} & \leq \sum_{2 \leq k \leq n} b(k)  \tag{24}\\
& \leq \sum_{2 \leq k \leq n}\left(c^{\prime} n^{-\epsilon h}+c n^{-2 \epsilon-1}\right) \\
& (\mathrm{iv}) c n^{-2 \epsilon},
\end{align*}
$$

where step (iv) follows by a calculation similar to the one carried out for the $d=1$ case. Relabelling the constants in condition (22) completes the proof.

### 4.2.1 Proof of Lemma 4

The first part of the proof is exactly the same as that of Lemma 1. In particular, Lemma 2 applies without modification to yield a bound identical to the inequality (15), given by

$$
\begin{equation*}
\operatorname{Pr}_{w}\left\{\Delta\left(\Pi, \Pi^{*}\right) \leq 0\right\} \leq c^{\prime} \exp \left(-c \frac{T_{\Pi}}{\sigma^{2}}\right) \tag{25}
\end{equation*}
$$

where $T_{\Pi}=\left\|P_{\Pi}^{\perp} \Pi^{*} A x^{*}\right\|_{2}^{2}$, as before.
The major difference from the $d=1$ case is in the random variable $T_{\Pi}$. Accordingly, we state the following parallel lemma to Lemma 3 .

Lemma 5. For $1<d<n$, any two permutation matrices $\Pi$ and $\Pi^{*}$ at Hamming distance $h$, and $t \leq h n^{-\frac{2 n}{n-d}}$, we have

$$
\begin{equation*}
\operatorname{Pr}_{A}\left\{T_{\Pi} \leq t\left\|x^{*}\right\|_{2}^{2}\right\} \leq 2 \max \left\{\exp \left(-n \log \frac{n}{2}\right), 6 \exp \left(-\frac{h}{10}\left[\log \left(\frac{h}{t n^{\frac{2 n}{n-d}}}\right)+\frac{t n^{\frac{2 n}{n-d}}}{h}-1\right]\right)\right\} \tag{26}
\end{equation*}
$$

The proof of Lemma 5 appears in Section 4.2 .2 . We are now ready to prove Lemma 4 .

Proof of Lemma 4. We prove Lemma 4 from Lemma 5 and equation 25 by an argument similar to the one before. In particular, in a similar vein to the steps leading up to equation (16), we have

$$
\begin{equation*}
\operatorname{Pr}\left\{\Delta\left(\Pi, \Pi^{*}\right) \leq 0\right\} \leq c^{\prime} \exp \left(-c \frac{t\left\|x^{*}\right\|_{2}^{2}}{\sigma^{2}}\right)+\operatorname{Pr}_{A}\left\{T_{\Pi} \leq t\left\|x^{*}\right\|_{2}^{2}\right\} \tag{27}
\end{equation*}
$$

We now use the shorthand snr $:=\frac{\left\|x^{*}\right\|_{2}^{2}}{\sigma^{2}}$ and let $t^{*}=h \frac{\log \left(\operatorname{snr} \cdot n^{-\frac{2 n}{n-d}}\right)}{\mathrm{snr}}$. Noting that snr $\cdot n^{-\frac{2 n}{n-d}}>$ $5 / 4$ yields $t^{*} \leq h n^{-\frac{2 n}{n-d}}$, we set $t=t^{*}$ in inequality (27) to obtain

$$
\begin{equation*}
\operatorname{Pr}\left\{\Delta\left(\Pi, \Pi^{*}\right) \leq 0\right\} \leq c^{\prime} \exp \left(-c h \log s n^{-\frac{2 n}{n-d}}\right)+\operatorname{Pr}_{A}\left\{T_{\Pi} \leq t^{*}\left\|x^{*}\right\|_{2}^{2}\right\} \tag{28}
\end{equation*}
$$

Since $\operatorname{Pr}_{A}\left\{T_{\Pi} \leq t^{*}\left\|x^{*}\right\|_{2}^{2}\right\}$ can be bounded by a maximum of two terms 26 , we now split the analysis into two cases depending on which term attains the maximum.

Case 1: First, suppose that the second term attains the maximum in inequality (26), i.e., $\operatorname{Pr}_{A}\left\{T_{\Pi} \leq t^{*}\left\|x^{*}\right\|_{2}^{2}\right\} \leq 12 \exp \left(-\frac{h}{10}\left[\log \left(\frac{h}{t^{*} n^{\frac{2 n}{n-d}}}\right)+\frac{t^{*} n^{\frac{2 n}{n-d}}}{h}-1\right]\right)$. Substituting for $t^{*}$, we have

We have snr $\cdot n^{-\frac{2 n}{n-d}}>\frac{5}{4}$, a condition which leads to the following pair of easily verifiable inequalities:

$$
\begin{align*}
& \log \left(\frac{\mathrm{snr} \cdot n^{-\frac{2 n}{n-d}}}{\log \left(\mathrm{snr} \cdot n^{-\frac{2 n}{n-d}}\right)}\right)+\frac{\log \left(\mathrm{snr} \cdot n^{-\frac{2 n}{n-d}}\right)}{\mathrm{snr} \cdot n^{-\frac{2 n}{n-d}}}-1 \geq \frac{\log \mathrm{snr} \cdot n^{-\frac{2 n}{n-d}}}{4}, \text { and }  \tag{30a}\\
& \log \left(\frac{\mathrm{snr} \cdot n^{-\frac{2 n}{n-d}}}{\log \left(\mathrm{snr} \cdot n^{-\frac{2 n}{n-d}}\right)}\right)+\frac{\log \left(\mathrm{snr} \cdot n^{-\frac{2 n}{n-d}}\right)}{\mathrm{snr} \cdot n^{-\frac{2 n}{n-d}}}-1 \leq 5 \log \left(\mathrm{snr} \cdot n^{-\frac{2 n}{n-d}}\right) . \tag{30b}
\end{align*}
$$

Using inequality (30a), we have

$$
\begin{equation*}
\operatorname{Pr}_{A}\left\{T_{\Pi} \leq t^{*}\left\|x^{*}\right\|_{2}^{2}\right\} \leq 12 \exp \left(-\operatorname{ch} \log \left(\operatorname{snr} \cdot n^{-\frac{2 n}{n-d}}\right)\right) . \tag{31}
\end{equation*}
$$

Now using inequalities (31) and (28) together yields

$$
\begin{equation*}
\operatorname{Pr}\left\{\Delta\left(\Pi, \Pi^{*}\right) \leq 0\right\} \leq c^{\prime} \exp \left(-c h \log \left(\operatorname{snr} \cdot n^{-\frac{2 n}{n-d}}\right)\right) \tag{32}
\end{equation*}
$$

It remains to handle the second case.

Case 2: Suppose now that $\operatorname{Pr}_{A}\left\{T_{\Pi} \leq t^{*}\left\|x^{*}\right\|_{2}^{2}\right\} \leq 2 \exp \left(-n \log \frac{n}{2}\right)$, i.e., that the first term in RHS of inequality (26) attains the maximum when $t=t^{*}$. In this case, we have

$$
\begin{aligned}
\exp \left(-n \log \frac{n}{2}\right) & \geq 6 \exp \left(-\frac{h}{10}\left[\log \left(\frac{h}{t^{*} n^{\frac{2 n}{n-d}}}\right)+\frac{t^{*} n^{\frac{2 n}{n-d}}}{h}-1\right]\right) \\
& \geq c^{\prime} \exp \left(-c h \log \left(\operatorname{snr} \cdot n^{-\frac{2 n}{n-d}}\right)\right),
\end{aligned}
$$

where step (i) follows from the right inequality (30b). Now substituting into inequality (28), we have

$$
\begin{align*}
\operatorname{Pr}\left\{\Delta\left(\Pi, \Pi^{*}\right) \leq 0\right\} & \leq c^{\prime} \exp \left(-c h \log \left(\operatorname{snr} \cdot n^{-\frac{2 n}{n-d}}\right)\right)+2 \exp \left(-n \log \frac{n}{2}\right) \\
& \leq c^{\prime} \exp \left(-n \log \frac{n}{2}\right) . \tag{33}
\end{align*}
$$

Combining equations (32) and (33) completes the proof of Lemma 4.

### 4.2.2 Proof of Lemma 5

We begin by reducing the problem to the case $x^{*}=e_{1}\left\|x^{*}\right\|_{2}$. In particular, if $W x^{*}=e_{1}\left\|x^{*}\right\|_{2}$ for a $d \times d$ unitary matrix $W$ and writing $A=\widetilde{A} W$, we have by rotation invariance of the Gaussian distribution that the the entries of $\widetilde{A}$ are distributed as i.i.d. standard Gaussians. It can be verified that $T_{\Pi}=\left\|I-\Pi \widetilde{A}\left(\widetilde{A}^{\top} \widetilde{A}\right)^{-1}(\Pi \widetilde{A})^{\top} \Pi^{*} \widetilde{A} e_{1}\right\|_{2}^{2}\left\|x^{*}\right\|_{2}^{2}$. Since $\widetilde{A} \stackrel{d}{=} A$, the reduction is complete.

In order to keep the notation uncluttered, we denote the first column of $A$ by $a$. We also denote the span of the first column of $\Pi A$ by $S_{1}$ and that of the last $d-1$ columns of $\Pi A$ by $S_{-1}$. Denote their respective orthogonal complements by $S_{1}^{\perp}$ and $S_{-1}^{\perp}$. We then have

$$
\begin{aligned}
T_{\Pi} & =\left\|x^{*}\right\|_{2}^{2}\left\|P_{\Pi}^{\perp} a\right\|_{2}^{2} \\
& =\left\|x^{*}\right\|_{2}^{2}\left\|P_{S_{-1} \cap S_{1}^{\perp}} a\right\|_{2}^{2} \\
& =\left\|x^{*}\right\|_{2}^{2}\left\|P_{S_{-1} \cap S_{1}^{\perp}}^{\perp} P_{S_{1}^{\perp}} a\right\|_{2}^{2} .
\end{aligned}
$$

We now condition on $a$. Consequently, the subspace $S_{1}^{\perp}$ is a fixed ( $n-1$ )-dimensional subspace. Additionally, $S_{-1}^{\perp} \cap S_{1}^{\perp}$ is the intersection of a uniformly random $(n-(d-1))$-dimensional subspace with a fixed ( $n-1$ )-dimensional subspace, and is therefore a uniformly random $(n-d)$-dimensional subspace within $S_{1}^{\perp}$. Writing $u=\frac{P_{S_{1}^{\perp}} a}{\left\|P_{S_{1}^{\perp}} a\right\|_{2}}$, we have

$$
T_{\Pi} \stackrel{d}{=}\left\|x^{*}\right\|_{2}^{2}\left\|P_{S_{-1}^{\perp} \cap S_{1}^{\perp}} u\right\|_{2}^{2}\left\|P_{S_{1}^{\perp}} a\right\|_{2}^{2} .
$$

Now since $u \in S_{1}^{\perp}$, note that $\left\|P_{S_{-1} \cap S_{1}^{\perp}} u\right\|_{2}^{2}$ is the squared length of a projection of an $(n-1)$ dimensional unit vector onto a uniformly chosen $(n-d)$-dimensional subspace. In other words, denoting a uniformly random projection from $m$ dimensions to $k$ dimensions by $P_{k}^{m}$, we have

$$
\left\|P_{S_{-1}^{\perp} \cap S_{1}^{\perp}} u\right\|_{2}^{2} \stackrel{d}{=}\left\|P_{n-d}^{n-1} v_{1}\right\|_{2}^{2} \stackrel{(\mathrm{i})}{=} 1-\left\|P_{d-1}^{n-1} v_{1}\right\|_{2}^{2},
$$

where $v_{1}$ represents a fixed standard basis vector in $n-1$ dimensions. The quantities $P_{n-d}^{n-1}$ and $P_{d-1}^{n-1}$ are projections onto orthogonal subspaces, and step (i) is a consequence of the Pythagorean theorem.

Now removing the conditioning on $a$, we see that the term for $d>1$ can be lower bounded by the corresponding $T_{\Pi}$ for $d=1$, but scaled by a random factor - the norm of a random projection. Using $T_{\Pi}^{1}:=\left\|P_{S_{1}}^{\perp} a\right\|_{2}^{2}\left\|x^{*}\right\|_{2}^{2}$ to denote $T_{\Pi}$ when $d=1$, we have

$$
\begin{equation*}
T_{\Pi}=\left(1-X_{d-1}\right) T_{\Pi}^{1} \tag{34}
\end{equation*}
$$

where we have introduced the shorthand $X_{d-1}=\left\|P_{d-1}^{n-1} e_{1}\right\|_{2}^{2}$.
We first handle the random projection term in equation (34) using Lemma 11 in Appendix A.2. In particular, substituting $\beta=(1-z) \frac{n-1}{d-1}$ in inequality 555 yields

$$
\begin{aligned}
\operatorname{Pr}\left\{1-X_{d-1} \leq z\right\} & \leq\left(\frac{n-1}{d-1}\right)^{(d-1) / 2}\left(\frac{z(n-1)}{n-d}\right)^{(n-d) / 2} \\
& \stackrel{\text { (i) }}{\leq} \sqrt{\binom{n-1}{d-1}} \sqrt{\binom{n-1}{n-d}} z^{\frac{n-d}{2}} \\
& =\binom{n-1}{d-1} z^{\frac{n-d}{2}} \\
& \text { (ii) } 2^{n-1} z^{\frac{n-d}{2}}
\end{aligned}
$$

where in steps (i) and (ii), we have used the standard inequality $2^{n} \geq\binom{ n}{r} \geq\binom{ n}{r}^{r}$. Now setting $z=n^{\frac{-2 n}{n-d}}$, which ensures that $(1-z) \frac{n-1}{d-1}>1$ for all $d<n$ and large enough $n$, we have

$$
\begin{equation*}
\operatorname{Pr}\left\{1-X_{d-1} \leq n^{\frac{-2 n}{n-d}}\right\} \leq \exp \left(-n \log \frac{n}{2}\right) . \tag{35}
\end{equation*}
$$

Applying the union bound then yields

$$
\begin{equation*}
\operatorname{Pr}\left\{T_{\Pi} \leq t\left\|x^{*}\right\|_{2}^{2}\right\} \leq \operatorname{Pr}\left\{1-X_{d-1} \leq n^{\frac{-2 n}{n-d}}\right\}+\operatorname{Pr}\left\{T_{\Pi}^{1} \leq t n^{\frac{2 n}{n-d}}\left\|x^{*}\right\|_{2}^{2}\right\} \tag{36}
\end{equation*}
$$

We have already computed an upper bound on $\operatorname{Pr}\left\{T_{\Pi}^{1} \leq t n^{\frac{2 n}{n-d}}\left\|x^{*}\right\|_{2}^{2}\right\}$ in Lemma 3. Applying it yields that provided $t \leq h n^{-\frac{2 n}{n-d}}$, we have

$$
\begin{equation*}
\operatorname{Pr}\left\{T_{\Pi}^{1} \leq t n^{\frac{2 n}{n-d}}\left\|x^{*}\right\|_{2}^{2}\right\} \leq 6\left(\frac{t n^{\frac{2 n}{n-d}}}{h} \exp \left(1-\frac{t n^{\frac{2 n}{n-d}}}{h}\right)\right)^{h / 10} . \tag{37}
\end{equation*}
$$

Combining equations (37) and (35) with the union bound (36) and performing some algebraic manipulation then completes the proof of Lemma 5 .

### 4.3 Information theoretic lower bounds on $\operatorname{Pr}\left\{\widehat{\Pi} \neq \Pi^{*}\right\}$

We prove Theorem 2 in Section 4.3.1 via the strong converse for the Gaussian channel. We prove the weak converse of Proposition 1 in Section 4.3 .2 by employing Fano's method. We note that the latter is a standard technique for proving minimax bounds in statistical inference problems Yu97, YB99.

### 4.3.1 Proof of Theorem 2

We begin by assuming that the design matrix $A$ is fixed, and that the estimator has knowledge of $x^{*}$ a-priori. Note that the latter cannot make the estimation task any easier. In proving this lower bound, we can also assume that the entries of $A x^{*}$ are distinct, since otherwise, perfect permutation recovery is impossible.

Given this setup, we now cast the problem as one of coding over a Gaussian channel. Toward this end, consider the codebook

$$
\mathcal{C}=\left\{\Pi A x^{*} \mid \Pi \in \mathcal{P}_{n}\right\}
$$

We may view $\Pi A x^{*}$ as the codeword corresponding to the permutation $\Pi$, where each permutation is associated to one of $n$ ! equally likely messages. Note that each codeword has power $\left\|A x^{*}\right\|_{2}^{2}$.

The codeword is then sent over a Gaussian channel with noise power equal to $\sum_{i=1}^{n} \sigma^{2}=n \sigma^{2}$. The decoding problem is to ascertain from the noisy observations which message was sent, or in other words, to identify the correct permutation.

We now use the non-asymptotic strong converse for the Gaussian channel Yos64. In particular, using Lemma 12 (see Appendix A.3) with $R=\frac{\log n!}{n}$ then yields that for any $\delta^{\prime}>0$, if

$$
\frac{\log n!}{n}>\frac{1+\delta^{\prime}}{2} \log \left(1+\frac{\left\|A x^{*}\right\|_{2}^{2}}{n \sigma^{2}}\right)
$$

then for any estimator $\widehat{\Pi}$, we have $\operatorname{Pr}\{\widehat{\Pi} \neq \Pi\} \geq 1-2 \cdot 2^{-n \delta^{\prime}}$. For the choice $\delta^{\prime}=\delta /(2-\delta)$, we have that if

$$
\begin{equation*}
(2-\delta) \log \left(\frac{n}{e}\right)>\log \left(1+\frac{\left\|A x^{*}\right\|_{2}^{2}}{n \sigma^{2}}\right) \tag{38}
\end{equation*}
$$

then $\operatorname{Pr}\{\widehat{\Pi} \neq \Pi\} \geq 1-2 \cdot 2^{-n \delta / 2}$. Note that the only randomness assumed so far was in the noise $w$ and the random choice of $\Pi$.

We now specialize the result for the case when $A$ is Gaussian. Toward that end, define the event

$$
\mathcal{E}(\delta)=\left\{1+\delta \geq \frac{\left\|A x^{*}\right\|_{2}^{2}}{n\left\|x^{*}\right\|_{2}^{2}}\right\}
$$

Conditioned on the event $\mathcal{E}(\delta)$, it can be verified that condition (5) implies condition (38). We also have

$$
\begin{aligned}
\operatorname{Pr}\{\mathcal{E}(\delta)\} & =1-\operatorname{Pr}\left\{\frac{\left\|A x^{*}\right\|_{2}^{2}}{n\left\|x^{*}\right\|_{2}^{2}}>1+\delta\right\} \\
& \stackrel{(\mathrm{i})}{\geq} 1-c^{\prime} e^{-c n \delta}
\end{aligned}
$$

where step (i) follows by using the sub-exponential tail bound (see Lemma 10 in Appendix A.2), since $\frac{\left\|A x^{*}\right\|_{2}^{2}}{\left\|x^{*}\right\|_{2}^{2}} \sim \chi_{n}^{2}$.

Putting together the pieces, we have that provided condition (5) holds,

$$
\begin{aligned}
\operatorname{Pr}\left\{\widehat{\Pi} \neq \Pi^{*}\right\} & \geq \operatorname{Pr}\left\{\widehat{\Pi} \neq \Pi^{*} \mid \mathcal{E}(\delta)\right\} \operatorname{Pr}\{\mathcal{E}(\delta)\} \\
& =\left(1-2 \cdot 2^{-n \delta / 2}\right)\left(1-c^{\prime} e^{-c n \delta}\right) \\
& \geq 1-c^{\prime} e^{-c n \delta}
\end{aligned}
$$

We now move on to the proof of Proposition 1

### 4.3.2 Proof of Proposition 1

Proposition 1 corresponds to a weak converse in the scenario where the estimator is also provided with the side information that $\Pi^{*}$ lies within a Hamming ball of radius 2 around the identity. We carry out the proof for the more general case where $\mathrm{d}_{\mathrm{H}}\left(\Pi^{*}, I\right) \leq \bar{h}$ for any $\bar{h} \geq 2$, later specializing to the case where $\bar{h}=2$. We denote such a Hamming ball by $\mathbb{B}_{\mathrm{H}}(\bar{h}):=\left\{\Pi \in \mathcal{P}_{n} \mid \mathrm{d}_{\mathrm{H}}(\Pi, I) \leq \bar{h}\right\}$.

For the sake of the lower bound, assume that our observation model takes the form

$$
\begin{equation*}
y=\Pi^{*} A x^{*}+\Pi^{*} w \tag{39}
\end{equation*}
$$

Since $w \in \mathbb{R}^{n}$ is i.i.d. standard normal, model (39) has a distribution equivalent to that of model (1).

Notice that the ML estimation problem is essentially a multi-way hypothesis testing problem among permutations in $\mathbb{B}_{\mathbf{H}}(\bar{h})$, and so Fano's method is directly applicable. As before, we assume that the estimator knows $x^{*}$, and consider a uniformly random choice of $\Pi^{*} \in \mathbb{B}_{\mathrm{H}}(\bar{h})$.

Now note that the observation vector $y$ is drawn from the mixture distribution

$$
\begin{equation*}
\mathbb{M}(\bar{h})=\frac{1}{\left|\mathbb{B}_{\mathbf{H}}(\bar{h})\right|} \sum_{\Pi \in \mathcal{P}_{n}} \mathbb{P}_{\Pi}, \tag{40}
\end{equation*}
$$

where $\mathbb{P}_{\Pi}$ denotes the Gaussian distribution $\mathcal{N}\left(\Pi A x^{*}, \sigma^{2} I_{n}\right)$. The following lemma provides a crucial statistic for our bounds.

Lemma 6. For $y$ drawn according to the mixture distribution $\mathbb{M}(\bar{h})$, we have

$$
\begin{equation*}
\operatorname{det} \mathbb{E}\left[y y^{\top}\right] \leq\left(\sigma^{2}+\left\|x^{*}\right\|_{2}^{2}\right)^{n}(1+n)\left(\frac{\bar{h}}{n}\right)^{n-1} \tag{41}
\end{equation*}
$$

We prove Lemma 6 in Section 4.3.3. Taking it as given, we are now in a position to prove Proposition 1 .

Proof of Proposition 1. We begin by using Fano's inequality to bound the probability of error of any tester random ensemble of $\Pi^{*}$. In particular, for any estimator $\widehat{\Pi}$ that is a measurable function of the pair $(y, A)$, we have

$$
\begin{equation*}
\operatorname{Pr}\left\{\widehat{\Pi} \neq \Pi^{*}\right\} \geq 1-\frac{I\left(\Pi^{*} ; y, A\right)+\log 2}{\log \left|\mathbb{B}_{\mathbf{H}}(\bar{h})\right|} \tag{42}
\end{equation*}
$$

Applying the chain rule for mutual information yields

$$
\begin{align*}
I\left(\Pi^{*} ; y, A\right) & =I\left(\Pi^{*} ; y \mid A\right)+I\left(\Pi^{*}, A\right) \\
& \stackrel{(\mathrm{i})}{=} I\left(\Pi^{*} ; y \mid A\right) \\
& =\mathbb{E}_{A}\left[I\left(\Pi^{*} ; y \mid A=\alpha\right)\right] \tag{43}
\end{align*}
$$

where step (i) follows since $\Pi^{*}$ is chosen independently of $A$. We now evaluate the mutual information term $I\left(\Pi^{*} ; y \mid A=\alpha\right)$, which we denote by $I_{\alpha}\left(\Pi^{*} ; y\right)$. Letting $H_{\alpha}(y):=H(y \mid A=\alpha)$ denote the conditional entropy of $y$ given a fixed realization of $A$, we have

$$
\begin{aligned}
I_{\alpha}\left(\Pi^{*} ; y\right) & =H_{\alpha}(y)-H_{\alpha}\left(y \mid \Pi^{*}\right) \\
& \text { (ii) } \frac{1}{\leq} \log \operatorname{det} \operatorname{cov} y y^{\top}-\frac{n}{2} \log \sigma^{2}
\end{aligned}
$$

where the covariance is evaluated with $A=\alpha$, and in step (ii), we have used two key facts:
(a) Gaussians maximize entropy for a fixed covariance, which bounds the first term, and
(b) For a fixed realization of $\Pi^{*}$, the vector $y$ is composed of $n$ uncorrelated Gaussians. This leads to an explicit evaluation of the second term.

Now taking expectations over $A$ and noting that $\operatorname{cov} y y^{\top} \preceq \mathbb{E}_{w}\left[y y^{\top}\right]$, we have from the concavity of the log determinant function and Jensen's inequality that

$$
\begin{align*}
I\left(\Pi^{*} ; y \mid A\right) & =\mathbb{E}_{A}\left[I_{\alpha}\left(\Pi^{*} ; y\right)\right] \\
& \leq \frac{1}{2} \log \operatorname{det} \mathbb{E}\left[y y^{\top}\right]-\frac{n}{2} \log \sigma^{2} \tag{44}
\end{align*}
$$

where the expectation in the last line is now taken over randomness in both $A$ and $w$.
Applying Lemma 6, we can then substitute inequality (41) into bound (44), which yields

$$
\mathbb{E}_{A}\left[I_{\alpha}\left(\Pi^{*} ; y\right)\right] \leq \frac{n}{2} \log \left(1+\frac{\left\|x^{*}\right\|_{2}^{2}}{\sigma^{2}}\right)+\frac{n-1}{2} \log \frac{\bar{h}}{n}+\frac{1}{2} \log (1+n)
$$

Finally, substituting into the Fano bound (42) yields the lower bound

$$
\begin{aligned}
\operatorname{Pr}\left\{\widehat{\Pi} \neq \Pi^{*}\right\} & \geq 1-\frac{\frac{n}{2} \log \left(1+\frac{\left\|x^{*}\right\|_{2}^{2}}{\sigma^{2}}\right)+\frac{n-1}{2} \log \frac{\bar{h}}{n}+\frac{1}{2} \log (1+n)+\log 2}{\log \left|\mathbb{B}_{\mathbf{H}}(\bar{h})\right|} \\
& \quad \text { (ii) } 1-\frac{\frac{n}{2} \log \left(1+\frac{\left\|x^{*}\right\|_{2}^{2}}{\sigma^{2}}\right)+\frac{n-1}{2} \log \frac{\bar{h}}{n}+\frac{1}{2} \log (1+n)+\log 2}{\bar{h} \log (n / e)}
\end{aligned}
$$

where in step (ii), we have used the fact that

$$
\left|\mathbb{B}_{\mathbf{H}}(\bar{h})\right|=\binom{n}{\bar{h}} \cdot \bar{h}!\geq(n / \bar{h})^{\bar{h}}\left(\frac{\bar{h}}{e}\right)^{\bar{h}}
$$

In other words, if

$$
\begin{equation*}
\log \left[4\left(1+\frac{\left\|x^{*}\right\|_{2}^{2}}{\sigma^{2}}\right)\right] \leq \frac{\bar{h}}{n} \log (n / e)+\frac{n-1}{n} \log (n / \bar{h})-\frac{\log (1+n)}{n} \tag{45}
\end{equation*}
$$

then $\operatorname{Pr}\left\{\widehat{\Pi} \neq \Pi^{*}\right\} \geq 1 / 2$ for any estimator $\widehat{\Pi}$. Evaluating condition (45) for $\bar{h}=2$ yields the required result.

### 4.3.3 Proof of Lemma 6

We first explicitly calculate the matrix $Y:=\mathbb{E}\left[y y^{\top}\right]$. Note that the diagonal entries take the form

$$
Y_{i i}=\left(x^{*}\right)^{\top} \mathbb{E}\left[a_{\pi_{i}} a_{\pi_{i}}^{\top}\right] x^{*}+\mathbb{E}\left[w_{\pi_{i}}^{2}\right]=\left\|x^{*}\right\|_{2}^{2}+\sigma^{2}
$$

The off-diagonal entries can be evaluated as

$$
Y_{i j}=\left(x^{*}\right)^{\top} \mathbb{E}\left[a_{\pi_{i}} a_{\pi_{j}}^{\top}\right] x^{*}+\mathbb{E}\left[w_{\pi_{i}} w_{\pi_{j}}\right] \stackrel{(\mathrm{i})}{=}\left(\frac{n-\bar{h}}{n}+\frac{\bar{h}}{n^{2}}\right)\left(\left\|x^{*}\right\|_{2}^{2}+\sigma^{2}\right), \text { for } i \neq j
$$

where step (i) follows since

$$
\begin{equation*}
\pi_{i}=\pi_{j} \text { with probability } \frac{n-\bar{h}}{n}+\frac{\bar{h}}{n^{2}} \tag{46}
\end{equation*}
$$

Equation (46) is a consequence of the fact that a uniform permutation over $\mathbb{B}_{\mathrm{H}}(\bar{h})$ can be generated by first picking $\bar{h}$ positions (the permutation set) out of $[n]$ uniformly at random, and then uniformly permuting those $\bar{h}$ positions. The probability that $\pi_{i}=\pi_{j}$ is equal to the probability that $\pi_{i}=i$, an event that occurs if:
(a) position $i$ is not chosen in the permutation set, which happens with probability $\frac{n-\bar{h}}{n}$, or if (b) position $i$ is in the permutation set but the permutation maps $i$ to itself, which happens with probability $\frac{\bar{h}}{n} \frac{1}{n}$.

Hence, the determinant of $Y$ is given by $\operatorname{det} Y=\left(\left\|x^{*}\right\|_{2}^{2}+\sigma^{2}\right)^{n} \operatorname{det} \bar{Y}$, where we have defined $\bar{Y}:=\frac{1}{\left\|x^{*}\right\|_{2}^{2}+\sigma^{2}} Y$. Note that $\bar{Y}$ is a highly structured matrix, and so its determinant can be computed exactly. In particular, letting the scalar $\beta=1$ denote the identical diagonal entries of $\bar{Y}$ and the scalar $\gamma$ denote its identical off-diagonal entries, it is easy to verify that the all ones vector $\mathbf{1}$ is an eigenvector of $\bar{Y}$, with corresponding eigenvalue $\beta+(n-1) \gamma$. Additionally, for any vector $v$ that obeys $\mathbf{1}^{\top} v=0$, we have

$$
\bar{Y} v=(\beta-\gamma) v+\gamma\left(v^{\top} \mathbf{1}\right) \mathbf{1}=(\beta-\gamma) v
$$

and so the remaining $n-1$ eigenvalues are identically $\beta-\gamma$.
Substituting for $\beta$ and $\gamma$, the eigenvalues of $\bar{Y}$ are given by

$$
\lambda_{1}(\bar{Y})=1+\frac{(n-\bar{h})(n-1)}{n}+\frac{\bar{h}(n-1)}{n^{2}}, \text { and }
$$

$$
\lambda_{2}(\bar{Y})=\lambda_{3}(\bar{Y})=\cdots=\lambda_{n}(\bar{Y})=\frac{\bar{h}}{n}-\frac{\bar{h}}{n^{2}}
$$

Hence, we have

$$
\begin{aligned}
\operatorname{det} \bar{Y} & =\left(1+\frac{(n-\bar{h})(n-1)}{n}+\frac{\bar{h}(n-1)}{n^{2}}\right)\left(\frac{\bar{h}}{n}-\frac{\bar{h}}{n^{2}}\right)^{n-1} \\
& \leq\left(1+n-\bar{h}+\frac{\bar{h}}{n}\right)\left(\frac{\bar{h}}{n}\right)^{n-1} \\
& \leq(1+n)\left(\frac{\bar{h}}{n}\right)^{n-1}
\end{aligned}
$$

where in the last step, we have used the fact that $0<\bar{h} \leq n$. This completes the proof.

### 4.4 Proof of Theorem 3

We now prove Theorem 3 for approximate permutation recovery. For any estimator $\widehat{\Pi}$, we denote by the indicator random variable $E(\widehat{\Pi}, D)$ whether or not the $\widehat{\Pi}$ has acceptable distortion, i.e., $E(\widehat{\Pi}, D)=\mathbb{I}\left[\mathrm{d}_{\mathrm{H}}\left(\widehat{\Pi}, \Pi^{*}\right) \geq D\right]$, with $E=1$ representing the error event. For $\Pi^{*}$ picked uniformly at random in $\mathcal{P}_{n}$, we have the following variant of Fano's inequality.

Lemma 7. The probability of error is lower bounded as

$$
\begin{equation*}
\operatorname{Pr}\{E(\widehat{\Pi}, D)=1\} \geq 1-\frac{I\left(\Pi^{*} ; y, A\right)+\log 2}{\log n!-\log \frac{n!}{(n-D+1)!}} \tag{47}
\end{equation*}
$$

Taking the lemma as given for the moment, we are now ready to prove Theorem 3 .

Proof of Theorem [3. The proof of the theorem follows by upper bounding the mutual information term. In particular, we have

$$
\begin{aligned}
I\left(\Pi^{*} ; y, A\right) & \leq \mathbb{E}_{A}\left[I_{a}\left(\Pi^{*} ; y\right)\right] \\
& \leq \frac{1}{2} \log \operatorname{det} \mathbb{E}\left[y y^{\top}\right]-\frac{n}{2} \log \sigma^{2} \\
& \stackrel{\text { (i) }}{\leq} \frac{n}{2} \log \left(1+\frac{\left\|x^{*}\right\|_{2}^{2}}{\sigma^{2}}\right)
\end{aligned}
$$

where the expectation on the RHS is taken over both $\Pi^{*}$ and $A$. Also, step (i) follows from the AM-GM inequality for PSD matrices $\operatorname{det} X \leq\left(\frac{1}{n} \operatorname{trace} X\right)^{n}$, and by noting that the diagonal entries of the matrix $\mathbb{E}\left[y y^{\top}\right]$ are all equal to $\left\|x^{*}\right\|_{2}^{2}+\sigma^{2}$.

Combining the pieces, we now have that $\operatorname{Pr}\left\{\widehat{\Pi} \neq \Pi^{*}\right\} \geq 1 / 2$ if

$$
\begin{equation*}
n \log \left(1+\frac{\left\|x^{*}\right\|_{2}^{2}}{\sigma^{2}}\right) \leq(n-D+1) \log \left(\frac{n-D+1}{2 e}\right) \tag{48}
\end{equation*}
$$

which completes the proof.

### 4.4.1 Proof of Lemma 7

We use the shorthand $E:=E(\widehat{\Pi}, D)$ in this proof to simplify notation. Proceeding by the usual proof of Fano's inequality, we begin by expanding $H\left(E, \Pi^{*} \mid y, A=a, \widehat{\Pi}\right)$ in two ways:

$$
\begin{align*}
H\left(E, \Pi^{*} \mid y, A, \widehat{\Pi}\right) & =H\left(\Pi^{*} \mid y, A, \widehat{\Pi}\right)+H\left(E \mid \Pi^{*}, y, A, \widehat{\Pi}\right)  \tag{49a}\\
& =H(E \mid y, A, \widehat{\Pi})+H\left(\Pi^{*} \mid E, y, A, \widehat{\Pi}\right) \tag{49b}
\end{align*}
$$

Since $\Pi^{*} \rightarrow(y, A) \rightarrow \widehat{\Pi}$ forms a Markov chain, we have $H\left(\Pi^{*} \mid y, A, \widehat{\Pi}\right)=H\left(\Pi^{*} \mid y, A\right)$. Nonnegativity of entropy yields $H\left(E \mid \Pi^{*}, y, A, \widehat{\Pi}\right) \geq 0$. Since conditioning cannot increase entropy, we have $H(E \mid y, A, \widehat{\Pi}) \leq H(E) \leq \log 2$, and $H\left(\Pi^{*} \mid E, y, A, \widehat{\Pi}\right) \leq H\left(\Pi^{*} \mid E, \widehat{\Pi}\right)$. Combining all of this with equations (49) yields

$$
\begin{align*}
H\left(\Pi^{*} \mid y, A\right) & \leq H\left(\Pi^{*} \mid E, \widehat{\Pi}\right)+\log 2 \\
& =\operatorname{Pr}\{E=1\} H\left(\Pi^{*} \mid E=1, \widehat{\Pi}\right)+(1-\operatorname{Pr}\{E=1\}) H\left(\Pi^{*} \mid E=0, \widehat{\Pi}\right)+\log 2 \tag{50}
\end{align*}
$$

We now use the fact that uniform distributions maximize entropy to bound the two terms as $H\left(\Pi^{*} \mid E=1, \widehat{\Pi}\right) \leq H\left(\Pi^{*}\right)=\log n!$, and $H\left(\Pi^{*} \mid E=0, \widehat{\Pi}\right) \leq \log \frac{n!}{(n-D+1)!}$, where the last inequality follows since $E=0$ reveals that $\Pi^{*}$ is within a Hamming ball of radius $D-1$ around $\widehat{\Pi}$, and the cardinality of that Hamming ball is $\frac{n!}{(n-D+1)!}$.

Substituting back into inequality (50) yields

$$
\begin{aligned}
\operatorname{Pr}\{E & =1\}\left(\log n!-\log \frac{n!}{(n-D+1)!}\right)+H\left(\Pi^{*}\right) \\
& \geq H\left(\Pi^{*} \mid y, A\right)-\log 2-\log \frac{n!}{(n-D+1)!}+\log n!
\end{aligned}
$$

where we have added the term $H\left(\Pi^{*}\right)=\log n$ ! to both sides. Simplifying then yields inequality (47).

### 4.5 Proofs of computational aspects

In this section, we prove Theorem 4 by providing an efficient algorithm for the $d=1$ case and showing NP-hardness for the $d>1$ case.

### 4.5.1 Proof of Theorem 4; $d=1$ case

In order to prove the theorem, we need to show an algorithm that performs the optimization (7) efficiently. Accordingly, note that for the case when $d=1$, equation (7) can be rewritten as

$$
\begin{align*}
\widehat{\Pi}_{\mathrm{ML}} & =\arg \max _{\Pi}\left\|a_{\Pi}^{\top} y\right\|^{2} \\
& =\arg \max _{\Pi} \max \left\{a_{\Pi}^{\top} y,-a_{\Pi}^{\top} y\right\} \\
& =\arg \min _{\Pi} \max \left\{\left\|a_{\Pi}-y\right\|_{2}^{2},\left\|a_{\Pi}+y\right\|_{2}^{2}\right\} \tag{51}
\end{align*}
$$

where the last step follows since $2 a_{\Pi}^{\top} y=\|a\|^{2}+\|y\|^{2}-\left\|a_{\Pi}-y\right\|_{2}^{2}$, and the first two terms do not involve optimizing over $\Pi$.

```
Algorithm 1: Exact algorithm for implementing equation (7) for the case when \(d=1\).
    Input: design matrix (vector) \(a\), observation vector \(y\)
    \(1 \Pi_{1} \leftarrow\) permutation that sorts \(a\) according to \(y\)
    \(2 \Pi_{2} \leftarrow\) permutation that sorts \(-a\) according to \(y\)
    \(3 \widehat{\Pi}_{\mathrm{ML}} \leftarrow \arg \max \left\{\left|a_{\Pi_{1}}^{\top} y\right|,\left|a_{\Pi_{2}}^{\top} y\right|\right\}\)
    Output: \(\widehat{\Pi}_{M L}\)
```

Once the optimization problem has been written in the form (51), it is easy to see that it can be solved in polynomial time. In particular, using the fact that for fixed vectors $p$ and $q$, $\left\|p_{\Pi}-q\right\|$ is minimized for $\Pi$ that sorts $a$ according to the order of $b$, we see that Algorithm 1 computes $\widehat{\Pi}_{M L}$ exactly.

The procedure defined by Algorithm 1 is clearly the correct thing to do in the noiseless case: in this case, $x^{*}$ is a scalar value that scales the entries of $a$, and so the correct permutation can be identified by a simple sorting operation. Two such operations suffice, one to account for when $x^{*}$ is positive and one more for when it is negative. Since each sort operation takes $\mathcal{O}(n \log n)$ steps, Algorithm 1 can be executed in nearly linear time.

### 4.5.2 Proof of Theorem 4; NP-hardness

In this section, we show that given a vector matrix pair $(y, A) \in \mathbb{R}^{n} \times \mathbb{R}^{n \times d}$, it is NP-hard to determine whether the equation $y=\Pi A x$ has a solution for a permutation matrix $\Pi \in \mathcal{P}_{n}$ and vector $x \in \mathbb{R}^{d}$. Clearly, this is sufficient to show that the problem (3) is NP-hard to solve in the case when $A$ and $y$ are arbitrary.

Our proof involves a reduction from the PARTITION problem, the decision version of which is defined as the following.

Definition 1 (PARTITION). Given $d$ integers $b_{1}, b_{2}, \cdots, b_{d}$, does there exist a subset $S \subset[d]$ such that

$$
\sum_{i \in S} b_{i}=\sum_{i \in[d] \backslash S} b_{i} ?
$$

It is well known PS98 that PARTITION is NP-complete. Also note that asking whether or not equation (1) has a solution ( $\Pi, x$ ) is equivalent to determining whether or not there exists a permutation $\pi$ and a vector $x$ such that $y_{\pi}=A x$ has a solution. We are now ready to prove the theorem.

Given an instance $b_{1}, \cdots b_{d}$ of PARTITION, define a vector $y \in \mathbb{Z}^{2 d+1}$ with entries

$$
y_{i}:=\left\{\begin{array}{l}
b_{i}, \text { if } i \in[d] \\
0, \text { otherwise } .
\end{array}\right.
$$

Also define the $2 d+1 \times 2 d$ matrix

$$
A:=\left[\begin{array}{cc}
I_{2 d} & \\
1_{d}^{\top} & -1_{d}^{\top}
\end{array}\right] .
$$

Clearly, the pair $(y, A)$ can be constructed from $b_{1}, \cdots b_{d}$ in time polynomial in $n=2 d+1$. We now claim that $y_{\pi}=A x$ has a solution $(\Pi, x)$ if and only if there exists a subset $S \subset[d]$ such that $\sum_{i \in S} b_{i}=\sum_{i \in[d] \backslash S} b_{i}$.

By converting to row echelon form, we see that $y_{\pi}=A x$ if and only if

$$
\begin{equation*}
\sum_{i \mid \pi(i) \leq d} y_{\pi(i)}=\sum_{i \mid \pi(i)>d} y_{\pi(i)}, \tag{52}
\end{equation*}
$$

and equation (52) holds, by construction, if and only if for $S=\{i \mid \pi(i) \leq d\} \cap[d]$, we have

$$
\sum_{i \in S} b_{i}=\sum_{i \in[d] \backslash S} b_{i} .
$$

This completes the proof.

## 5 Discussion

We analyzed the problem of exact permutation recovery in the linear regression model, and provided necessary and sufficient conditions that are tight in most regimes of $n$ and $d$. We also provided a converse for the problem of approximate permutation recovery to within some Hamming distortion. It is still an open problem to characterize the fundamental limits of exact and approximate permutation recovery for all regimes of $n, d$ and the allowable distortion $D$. In the context of exact permutation recovery, we believe that the limit suggested by Theorem 1 is tight for all regimes of $n$ and $d$, but showing this will likely require a different technique. In particular, as pointed out in Remark 1, all of our lower bounds assume that the estimator is provided with $x^{*}$ as side information; it is an interesting question as to whether stronger lower bounds can be obtained without this side information.

On the computational front, many open questions remain. The primary question concerns the design of computationally efficient estimators that succeed in similar SNR regimes. We have already shown that the maximum likelihood estimator, while being statistically optimal for moderate $d$, is computationally hard to compute in the worst case. Showing a corresponding hardness result for random $A$ is also an open problem. Finally, while this paper mainly addresses the problem of permutation recovery, the complementary problem of recovering $x^{*}$ is also interesting, and we plan to investigate its fundamental limits in future work.

## Acknowledgements

This work was partially supported by NSF Grants CCF-1528132 and CCF-0939370 (Center for Science of Information), Office of Naval Research MURI grant DOD-002888, Air Force Office of Scientific Research Grant AFOSR-FA9550-14-1-001, Office of Naval Research grant ONR-N00014, as well as National Science Foundation Grant CIF-31712-23800.

## A Auxiliary results

In this section, we prove a preliminary lemma about permutations that is useful in many of our proofs. We also derive tight bounds on the lower tails of $\chi^{2}$-random variables and state an existing result on tail bounds for random projections.

## A. 1 Independent sets of permutations

In this section, we prove a combinatorial lemma about permutations, for which we need to set up some additional notation. For a permutation $\pi$ on $k$ objects, let $G_{\pi}$ denote the
corresponding undirected incidence graph, i.e., $V\left(G_{\pi}\right)=[k]$, and $(i, j) \in E\left(G_{\pi}\right)$ iff $j=\pi(i)$ or $i=\pi(j)$.

Lemma 8. Let $\pi$ be a permutation on $k \geq 3$ objects such that $\mathrm{d}_{\mathrm{H}}(\pi, I)=k$. Then the vertices of $G_{\pi}$ can be partitioned into three sets $V_{1}, V_{2}, V_{3}$ such that each is an independent set, and $\left|V_{1}\right|,\left|V_{2}\right|,\left|V_{3}\right| \geq\left\lfloor\frac{k}{3}\right\rfloor \geq \frac{k}{5}$.

Proof. Note that for any permutation $\pi$, the corresponding graph $G_{\pi}$ is composed of cycles, and the vertices in each cycle together form an independent set. Consider one such cycle. We can go through the vertices in the order induced by the cycle, and alternate placing them in each of the 3 partitions. Clearly, this produces independent sets, and furthermore, having 3 partitions ensures that the last vertex in the cycle has some partition with which it does not share edges. If the cycle length $C \equiv 0(\bmod 3)$, then each partition gets $C / 3$ vertices, otherwise the smallest partition has $\lfloor C / 3\rfloor$ vertices. The partitions generated from the different cycles can then be combined (with relabelling, if required) to ensure that the largest partition has cardinality at most 1 more than that of the smallest partition.

## A. 2 Tails bounds on $\chi^{2}$ random variables and random projections

In our analysis, we require tight control on lower tails of $\chi^{2}$ random variables. The following lemma provides one such bound.

Lemma 9. Let $Z_{\ell}$ denote a $\chi^{2}$ random variable with $\ell$ degrees of freedom. Then for all $p \in[0, \ell]$, we have

$$
\begin{equation*}
\operatorname{Pr}\left\{Z_{\ell} \leq p\right\} \leq\left(\frac{p}{\ell} \exp \left(1-\frac{p}{\ell}\right)\right)^{\ell / 2}=\exp \left(-\frac{\ell}{2}\left[\log \frac{\ell}{p}+\frac{p}{\ell}-1\right]\right) \tag{53}
\end{equation*}
$$

Proof. The lemma is a simple consequence of the Chernoff bound. In particular, we have for all $\lambda>0$ that

$$
\begin{align*}
\operatorname{Pr}\left\{Z_{\ell} \leq p\right\} & =\operatorname{Pr}\left\{\exp \left(-\lambda Z_{\ell}\right) \geq \exp (-\lambda p)\right\} \\
& \leq \exp (\lambda p) \mathbb{E}\left[\exp \left(-\lambda Z_{\ell}\right)\right] \\
& =\exp (\lambda p)(1+2 \lambda)^{-\frac{\ell}{2}} \tag{54}
\end{align*}
$$

where in the last step, we have used $\mathbb{E}\left[\exp \left(-\lambda Z_{\ell}\right)\right]=(1+2 \lambda)^{-\frac{\ell}{2}}$, which is valid for all $\lambda>-1 / 2$. Minimizing the last expression over $\lambda>0$ then yields the choice $\lambda^{*}=\frac{1}{2}\left(\frac{\ell}{p}-1\right)$, which is greater than 0 for all $0 \leq p \leq \ell$. Substituting this choice back into equation (54) proves the lemma.

We also state the following lemma for general sub-exponential random variables (see, for example, Wainwright [Wai15, Theorem 2.2]). We use it in the context of $\chi^{2}$ random variables.

Lemma 10. Let $X$ be a sub-exponential random variable. Then for all $t>0$, we have

$$
\operatorname{Pr}\{|X-\mathbb{E}[X]| \geq t\} \leq c^{\prime} e^{-c t}
$$

Lastly, we require tail bounds on the norms of random projections, a problem that has been studied extensively in the literature on dimensionality reduction. The following lemma, a consequence of the Chernoff bound, is taken from Dasgupta and Gupta [DG03, Lemma $2.2 \mathrm{~b}]$.

Lemma 11 ([DG03]). Let $x$ be a fixed $n$-dimensional vector, and let $P_{d}^{n}$ be a projection matrix from n-dimensional space to a uniformly randomly chosen d-dimensional subspace. Then we have for every $\beta>1$ that

$$
\begin{equation*}
\operatorname{Pr}\left\{\left\|P_{d}^{n} x\right\|_{2}^{2} \geq \frac{\beta d}{n}\|x\|_{2}^{2}\right\} \leq \beta^{d / 2}\left(1+\frac{(1-\beta) d}{n-d}\right)^{(n-d) / 2} \tag{55}
\end{equation*}
$$

## A. 3 Strong converse for Gaussian channel capacity

The following result due to Shannon [Sha59] provides a strong converse for the Gaussian channel. The non-asymptotic version as stated here was also derived by Yoshihara Yos64.
Lemma 12 ([Yos64]). Consider a vector Gaussian channel on $n$ coordinates with message power $P$ and noise power $\sigma^{2}$, whose capacity is given by $\bar{R}=\log \left(1+\frac{P}{\sigma^{2}}\right)$. For any codebook $\mathcal{C}$ with $|\mathcal{C}|=2^{\text {nR }}$, if for some $\epsilon>0$ we have

$$
R>(1+\epsilon) \bar{R},
$$

then the probability of error $p_{e} \geq 1-2 \cdot 2^{-n \epsilon}$ for $n$ large enough.

## B Proof of Lemma 2

We prove each claim of the lemma separately.

## B. 1 Proof of claim 13a

To start, note that by definition of the linear model, we have $\left\|P_{\Pi^{*}}^{\perp} *\right\|_{2}^{2}=\left\|P_{\Pi^{*}}^{\perp} w\right\|_{2}^{2}$. Letting $Z_{\ell}$ denote a $\chi^{2}$ random variable with $\ell$ degrees of freedom, we claim that

$$
\left\|P_{\Pi^{*}}^{\perp} w\right\|_{2}^{2}-\left\|P_{\Pi}^{\perp} w\right\|_{2}^{2}=Z_{k}-\tilde{Z}_{k},
$$

where $k:=\min \left(d, \mathrm{~d}_{\mathrm{H}}\left(\Pi, \Pi^{*}\right)\right)$.
For the rest of the proof, we adopt the shorthand $\Pi \backslash \Pi^{\prime}:=\operatorname{range}(\Pi A) \backslash \operatorname{range}\left(\Pi^{\prime} A\right)$, and $\Pi \cap \Pi^{\prime}:=\operatorname{range}(\Pi A) \cap \operatorname{range}\left(\Pi^{\prime} A\right)$. Now, by the Pythagorean theorem, we have

$$
\left\|P_{\Pi^{*}}^{\perp} w\right\|_{2}^{2}-\left\|P_{\Pi}^{\perp} w\right\|_{2}^{2}=\left\|P_{\Pi} w\right\|_{2}^{2}-\left\|P_{\Pi^{*}} w\right\|_{2}^{2} .
$$

Splitting it up further, we can then write

$$
\left\|P_{\Pi} w\right\|_{2}^{2}=\left\|P_{\Pi \cap \Pi^{*}} w\right\|_{2}^{2}+\left\|\left(P_{\Pi}-P_{\Pi \cap \Pi^{*}}\right) w\right\|_{2}^{2},
$$

where we have used the fact that $P_{\Pi \cap \Pi^{*}} P_{\Pi}=P_{\Pi \cap \Pi^{*}}=P_{\Pi \cap \Pi^{*}} P_{\Pi^{*}}$.
Similarly for the second term, we have $\left\|P_{\Pi^{*}} w\right\|_{2}^{2}=\left\|P_{\Pi \cap \Pi^{*}} w\right\|_{2}^{2}+\left\|\left(P_{\Pi^{*}}-P_{\Pi \cap \Pi^{*}}\right) w\right\|_{2}^{2}$, and hence,

$$
\left\|P_{\Pi} w\right\|_{2}^{2}-\left\|P_{\Pi^{*}} w\right\|_{2}^{2}=\left\|\left(P_{\Pi}-P_{\Pi \cap \Pi^{*}}\right) w\right\|_{2}^{2}-\left\|\left(P_{\Pi^{*}}-P_{\Pi \cap \Pi^{*}}\right) w\right\|_{2}^{2}
$$

Now each of the two projection matrices above has rank ${ }^{3} \operatorname{dim}\left(\Pi \backslash \Pi^{*}\right)=k$, which completes the proof of the claim. To prove the lemma, note that for any $\delta>0$, we can write

$$
\operatorname{Pr}\left\{\mathcal{F}_{1}(\delta)\right\} \leq \operatorname{Pr}\left\{\left|Z_{k}-k\right| \geq \delta / 2\right\}+\operatorname{Pr}\left\{\left|\tilde{Z}_{k}-k\right| \geq \delta / 2\right\} .
$$

Using the sub-exponential tail-bound on $\chi^{2}$ random variables (see Lemma 10 in Appendix A.2) completes the proof.

[^2]
## B. 2 Proof of claim (13b)

We begin by writing

$$
\operatorname{Pr}_{w}\left\{\mathcal{F}_{2}(\delta)\right\}=\operatorname{Pr}_{w}\{\underbrace{\left\|P_{\Pi}^{\perp} \Pi^{*} A x^{*}\right\|_{2}^{2}+2\left\langle P_{\Pi}^{\perp} \Pi^{*} A x^{*}, P_{\Pi}^{\perp} w\right\rangle}_{R(A, w)} \leq 2 \delta\} .
$$

We see that conditioned on $A$, the random variable $R(A, w)$ is distributed as $\mathcal{N}\left(T_{\Pi}, 4 \sigma^{2} T_{\Pi}\right)$, where we have used the shorthand $T_{\Pi}:=\left\|P_{\Pi}^{\perp} \Pi^{*} A x^{*}\right\|_{2}^{2}$.

So applying standard Gaussian tail bounds (see, for example, Wainwright Wai15, Example 2.1]), we have

$$
\operatorname{Pr}_{w}\left\{\mathcal{F}_{2}(\delta)\right\} \leq \exp \left(-\frac{\left(T_{\Pi}-2 \delta\right)^{2}}{8 \sigma^{2} T_{\Pi}}\right) .
$$

Setting $\delta=\delta^{*}:=\frac{1}{3} T_{\Pi}$ completes the proof.

## References

[Bal62] A. V. Balakrishnan. On the problem of time jitter in sampling. IRE Transactions on Information Theory, 8(3):226-236, 1962.
[Blu11] T. Blumensath. Sampling and reconstructing signals from a union of linear subspaces. IEEE Transactions on Information Theory, 57(7):4660-4671, 2011.
[CD16] O. Collier and A. S. Dalalyan. Minimax rates in permutation estimation for feature matching. Journal of Machine Learning Research, 17(6):1-31, 2016.
[Cha15] S. Chatterjee. Matrix estimation by universal singular value thresholding. The Annals of Statistics, 43(1):177-214, 2015.
[CT06] E. J. Candes and T. Tao. Near-optimal signal recovery from random projections: Universal encoding strategies? Information Theory, IEEE Transactions on, 52(12):5406-5425, 2006.
[DDDS04] P. David, D. Dementhon, R. Duraiswami, and H. Samet. Softposit: Simultaneous pose and correspondence determination. International Journal of Computer Vision, 59(3):259-284, 2004.
[DG03] S. Dasgupta and A. Gupta. An elementary proof of a theorem of Johnson and Lindenstrauss. Random Structures \& Algorithms, 22(1):60-65, 2003.
[EBDG14] V. Emiya, A. Bonnefoy, L. Daudet, and R. Gribonval. Compressed sensing with unknown sensor permutation. In Acoustics, Speech and Signal Processing (ICASSP), 2014 IEEE International Conference on, pages 1040-1044. IEEE, 2014.
[FJBd13] F. Fogel, R. Jenatton, F. Bach, and A. d'Aspremont. Convex relaxations for permutation problems. In Advances in Neural Information Processing Systems, pages 1016-1024, 2013.
[FMR16] N. Flammarion, C. Mao, and P. Rigollet. Optimal rates of statistical seriation. arXiv preprint arXiv:1607.02435, 2016.
[HGG09] J. Huang, C. Guestrin, and L. Guibas. Fourier theoretic probabilistic inference over permutations. The Journal of Machine Learning Research, 10:997-1070, 2009.
[HM99] X. Huang and A. Madan. CAP3: A DNA sequence assembly program. Genome Research, 9(9):868-877, 1999.
[KSF ${ }^{+}$09] L. Keller, M. J. Siavoshani, C. Fragouli, K. Argyraki, and S. Diggavi. Identity aware sensor networks. In INFOCOM 2009, IEEE, pages 2177-2185. IEEE, 2009.
[LD08] Y. M. Lu and M. N. Do. A theory for sampling signals from a union of subspaces. IEEE Transactions on Signal Processing, 56(6):2334-2345, 2008.
$\left[\operatorname{LdABN}^{+} 07\right]$ E. M. Loiola, N. M. M. de Abreu, P. O. Boaventura-Netto, P. Hahn, and T. Querido. A survey for the quadratic assignment problem. European Journal of Operational Research, 176(2):657-690, 2007.
[LR14] R. J. A. Little and D. B. Rubin. Statistical analysis with missing data. John Wiley \& Sons, 2014.
[LW12] P. Loh and M. J. Wainwright. Corrupted and missing predictors: Minimax bounds for high-dimensional linear regression. In Information Theory Proceedings (ISIT), 2012 IEEE International Symposium on, pages 2601-2605. IEEE, 2012.
[MSC09] M. Marques, M. Stošić, and J. Costeira. Subspace matching: Unique solution to point matching with geometric constraints. In Computer Vision, IEEE 12th International Conference on, pages 1288-1294. IEEE, 2009.
[NS08] A. Narayanan and V. Shmatikov. Robust de-anonymization of large sparse datasets. In Security and Privacy, 2008. SP 2008. IEEE Symposium on, pages 111-125. IEEE, 2008.
[PG06] A. B. Poore and S. Gadaleta. Some assignment problems arising from multiple target tracking. Mathematical and Computer Modelling, 43(9):1074-1091, 2006.
[PS98] C. H. Papadimitriou and K. Steiglitz. Combinatorial optimization: algorithms and complexity. Courier Corporation, 1998.
[RMS12] C. Rose, I. S. Mian, and R. Song. Timing channels with multiple identical quanta. arXiv preprint arXiv:1208.1070, 2012.
[Rob51] W. S. Robinson. A method for chronologically ordering archaeological deposits. American Antiquity, pages 293-301, 1951.
[SBGW15] N. B. Shah, S. Balakrishnan, A. Guntuboyina, and M. J. Wainright. Stochastically transitive models for pairwise comparisons: Statistical and computational issues. arXiv preprint arXiv:1510.05610, 2015.
[Sha59] C. E. Shannon. Probability of error for optimal codes in a Gaussian channel. Bell System Technical Journal, 38(3):611-656, 1959.
[SZ99] L. J. Schulman and D. Zuckerman. Asymptotically good codes correcting insertions, deletions, and transpositions. IEEE Transactions on Information Theory, 45(7):2552-2557, 1999.
[TL08] S. Thrun and J. J. Leonard. Simultaneous localization and mapping. In Springer Handbook of Robotics, pages 871-889. Springer, 2008.
[UHV15] J. Unnikrishnan, S. Haghighatshoar, and M. Vetterli. Unlabeled sensing with random linear measurements. preprint arXiv:1512.00115, 2015.
[Wai15] M. J. Wainwright. High-dimensional statistics: A non-asymptotic viewpoint. in preparation. University of California, Berkeley, 2015.
[YB99] Y. Yang and A. Barron. Information-theoretic determination of minimax rates of convergence. Annals of Statistics, pages 1564-1599, 1999.
[Yos64] K. Yoshihara. Simple proofs for the strong converse theorems in some channels. In Kodai Mathematical Seminar Reports, volume 16, pages 213-222. Dept. of Mathematics, Tokyo Institute of Technology, 1964.
[Yu97] B. Yu. Assouad, Fano, and Le Cam. In Festschrift for Lucien Le Cam, pages 423-435. Springer, 1997.


[^0]:    ${ }^{1}$ Unnikrishnan et al. UHV15 proposed that $e_{i}$ be repeated $i$ times, but it is easy to see that this does not ensure recovery of an adversarially chosen $x^{*}$.

[^1]:    ${ }^{2}$ An $n \times d$ i.i.d. Gaussian random matrix has full column rank with probability 1 as long as $d \leq n$

[^2]:    ${ }^{3}$ With probability 1

