

Oberwolfach Preprints



OWP 2013 - 05

FRANZ J. KIRÁLY, LOUIS THERAN, RYOTA TOMIOKA, AND
TAKEAKI UNO

The Algebraic Combinatorial Approach for Low-
Rank Matrix Completion

Mathematisches Forschungsinstitut Oberwolfach gGmbH
Oberwolfach Preprints (OWP) ISSN 1864-7596

Oberwolfach Preprints (OWP)

Starting in 2007, the MFO publishes a preprint series which mainly contains research results related to a longer stay in Oberwolfach. In particular, this concerns the Research in Pairs-Programme (RiP) and the Oberwolfach-Leibniz-Fellows (OWLF), but this can also include an Oberwolfach Lecture, for example.

A preprint can have a size from 1 - 200 pages, and the MFO will publish it on its website as well as by hard copy. Every RiP group or Oberwolfach-Leibniz-Fellow may receive on request 30 free hard copies (DIN A4, black and white copy) by surface mail.

Of course, the full copy right is left to the authors. The MFO only needs the right to publish it on its website *www.mfo.de* as a documentation of the research work done at the MFO, which you are accepting by sending us your file.

In case of interest, please send a **pdf file** of your preprint by email to *rip@mfo.de* or *owlf@mfo.de*, respectively. The file should be sent to the MFO within 12 months after your stay as RiP or OWLF at the MFO.

There are no requirements for the format of the preprint, except that the introduction should contain a short appreciation and that the paper size (respectively format) should be DIN A4, "letter" or "article".

On the front page of the hard copies, which contains the logo of the MFO, title and authors, we shall add a running number (20XX - XX).

We cordially invite the researchers within the RiP or OWLF programme to make use of this offer and would like to thank you in advance for your cooperation.

Imprint:

Mathematisches Forschungsinstitut Oberwolfach gGmbH (MFO)
Schwarzwaldstrasse 9-11
77709 Oberwolfach-Walke
Germany

Tel +49 7834 979 50
Fax +49 7834 979 55
Email admin@mfo.de
URL www.mfo.de

The Oberwolfach Preprints (OWP, ISSN 1864-7596) are published by the MFO.
Copyright of the content is held by the authors.

The Algebraic Combinatorial Approach for Low-Rank Matrix Completion

Franz J. Király^{cbd1}, Louis Theran^{b2}, Ryota Tomioka^{a3}, and Takeaki Uno^{e4}

^aDepartment of Mathematical Informatics, The University of Tokyo

^bDiscrete Geometry Group, Freie Universität Berlin

^cMachine Learning Group, Berlin Institute of Technology

^dMathematisches Forschungsinstitut Oberwolfach

^eNational Institute of Informatics, Tokyo

Abstract

We propose an algebraic combinatorial framework for the problem of completing partially observed low-rank matrices. We show that the intrinsic properties of the problem, including which entries can be reconstructed, and the degrees of freedom in the reconstruction, do not depend on the values of the observed entries, but only on their position. We associate combinatorial and algebraic objects, differentials and matroids, which are descriptors of the particular reconstruction task, to the set of observed entries, and apply them to obtain reconstruction bounds. We show how similar techniques can be used to obtain reconstruction bounds on general compressed sensing problems with algebraic compression constraints. Using the new theory, we develop several algorithms for low-rank matrix completion, which allow to determine which set of entries can be potentially reconstructed and which not, and how, and we present algorithms which apply algebraic combinatorial methods in order to reconstruct the missing entries.

¹franz.j.kiraly@tu-berlin.de

²theran@math.fu-berlin.de

³tomioka@mist.i.u-tokyo.ac.jp

⁴uno@nii.jp

Contents

1	Introduction	3
1.1	Matrix Completion	3
1.2	Results	5
1.3	Acknowledgements	7
2	Theory of Low-Rank Matrix Completion	7
2.1	What is Low-Rank Matrix Completion?	8
2.1.1	Separating Generative and Noise Models	8
2.1.2	The Parametric Model of Low-Rank Matrix Completion	10
2.2	Genericity	13
2.2.1	Sampling of the Mask	13
2.2.2	Sampling of the True Matrix	14
2.3	The Algebraic Combinatorics of Matrix Completion	16
2.3.1	The Algebra of Matrix Completion	17
2.3.2	The Combinatorics of Matrix Completion	21
2.4	Partial Completion	30
2.5	Differential Analysis of Completion Patterns	32
2.5.1	Degrees of Freedom	33
2.5.2	Differentials	34
2.5.3	The Determinantal Matroid	40
2.6	Completeness of random masks	49
2.6.1	Random graph models	49
2.6.2	Completeness of incoherent matrices	49
2.6.3	Sparser sampling and partial completeness	50
2.6.4	Denser sampling and solving minor by minor	51
2.7	Sufficient Sampling Densities for Algebraic Compressed Sensing	52
2.7.1	Finiteness of Random Maskings	52
2.7.2	Finiteness of Random Projections	60
3	Algorithms	60
3.1	Randomized Algorithms for Completeness	60
3.2	Algorithms for checking Closability	62
3.3	Algebraic Reconstruction of Matrices	62
4	Experiments	65
4.1	Randomized Algorithms for Completeness	65
4.2	Phase Transitions	67
5	Conclusion	68
A	Sheaves of Matroids	75

1. Introduction

1.1 Matrix Completion

Matrix completion is the task to reconstruct low-rank matrices from a subset of its entries and occurs naturally in many practically relevant problems, such as missing feature imputation, multi-task learning [2], transductive learning [13], or collaborative filtering and link prediction [1, 29, 45].

With the nuclear norm heuristic having been applied with increasing success in the reconstruction of low-rank matrices [5, 45], it has become increasingly important to analyze the potential and limitations of matrix completion methods.

Existing approaches can be classified by the assumptions about the sampling procedure and the low-rank matrices whose entries are measured. Candès and Recht [5] analyzed the noiseless setting, and have shown under uniform sampling that incoherent low-rank matrices can be recovered with high probability. Salakhutdinov and Srebro [40] considered the more realistic setting where the rows and columns are non-uniformly sampled. Negahban and Wainwright [33] showed under the same row/column weighted sampling that non-spiky low-rank matrices can be recovered with large probability. Foygel and Srebro [12] have shown under uniform sampling that the max-norm heuristic [44] can achieve superior reconstruction guarantee under the non-spikiness assumption on the underlying low-rank matrix.

All the above theoretical guarantees are built on some assumption on the sampling procedure, e.g., uniform sampling. In a practical setting, we always know which entries we can observe and which entries we cannot (the so-called mask). One may ask if we could obtain a stronger theoretical guarantee (of success or failure) conditioned on the mask we have.

On the other hand, all the above theories are also based on some assumptions on the underlying low-rank matrix, which are usually uncheckable. Although it is widely known that we cannot recover arbitrary low-rank matrices (see, e.g., Equation (1.1) in [5]), one may ask if there is a theory for matrix completion for almost all matrices, depending only on the mask.

Following the expository paper of Király and Tomioka [25], we view matrix completion as a problem lying in the intersection of two mathematical realms, combinatorics and algebra. Here the combinatorial structure arises from the masking pattern, which can be viewed as a bipartite graph, and the algebraic structure arises from the low-rank-ness of the underlying matrix. It is probably fair to say that previous studies (with some exceptions we mention below) have not paid enough attention to these underlying mathematical structures.

Looking into more details about the combinatorial/algebraic structures of the problem allows us to derive novel necessary and sufficient conditions for any matrix completion algorithm to succeed in recovering the underlying low-rank matrix from some of its entries. Figure 1 shows how combinatorial properties of the mask, such as r -closable, $2r$ -regular, and r -connected relate to unique/finite (up to finite number of solutions) completability. Although these combinatorial properties are implied with high probability from the sampling models (e.g., uniform) depending on the expected number of observable entries, they had hardly been discussed in the literature. We first discuss these combinatorial properties for a fixed mask in detail, and then show when these conditions are satisfied with high probability under a certain sampling model of the mask.

Another point that differentiates our work from previous studies is that we avoid making any explicit assumption on the underlying low-rank matrix. Although this may sound intractable, it is possible when we consider completability for generic low-rank matrices, i.e., for almost all low-

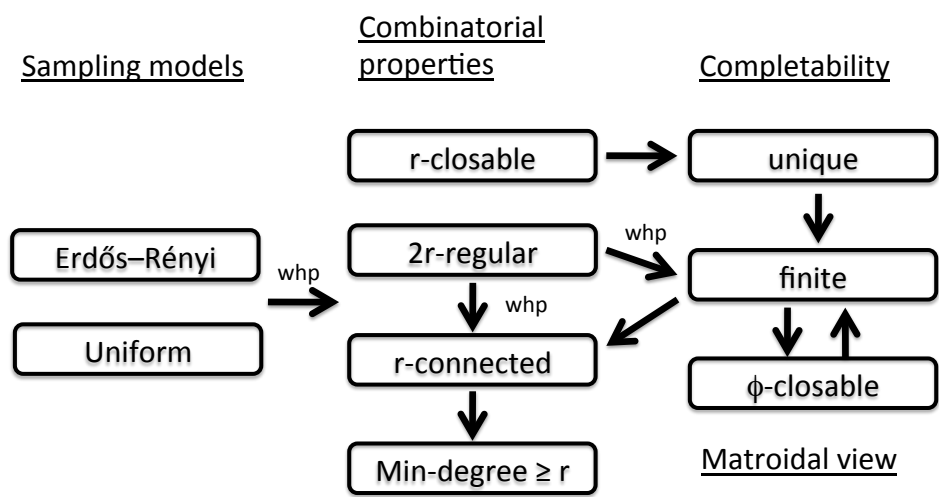


Figure 1: Combinatorial properties we discuss in this paper. “whp” means with high probability for sufficiently high sampling density (or expected number of observed entries.)

rank matrices allowing a set of exceptions of measure zero. This is illustrated in Figure 2. Since it is clear that we cannot successfully recover an arbitrary low-rank matrix, we also need to make exceptions. However, the set of exceptional cases has zero Lebesgue measure, thus also zero probability measure under any continuous sampling. Previous studies have generally used some bound on some quantities (coherence/spikiness) that characterize the goodness of the matrix, which correspond to a set of exceptional cases with positive measure.

Exploiting the algebraic/combinatorial structures of the problem, we propose the notion of *partial completability*. Precisely, our algebraic-combinatorial tool allows us to tell which entry can be imputed and which entry cannot. Since an entry is (finitely) completable if and only if that entry has some algebraic dependence on the observed entries, this notion has a connection to *matroid theory*, which capture the notion of independence, dependence, and span for subsets




Figure 2: Difference between the conditions used previously in literature (non-spiky/incoherent) and the generic assumption we use in this paper.

of a finite set, which in this case is the set of entries of a low-rank matrix, making them the “right” tool for formalizing degrees of freedom. We propose a scalable randomized algorithm to identify which entries that may be recovered from the given ones, which can be considered as a generalization of an algorithm proposed in Singer and Cucuringu [43], but we rigorously prove its correctness. Note that even when an entry is not completable, the nuclear norm heuristic would give some result. However the result may not be reliable in that case.

Furthermore, we propose a polynomial time algorithm to check for the property we call r -closable and at the same time actually perform matrix completion for a mask with such a property. We show that this approach can be superior to the well studied nuclear norm heuristic in some cases. In addition, we discuss the limitation of this approach, and how it is related to the more general notion of circuits of a matroid, which are not however as easy to compute as the r -closure.

1.2 Results

As the general overview in the previous section indicates, matrix completion has, until now, been analyzed predominantly in the context of convex optimization. Indeed, naively, one could think that the findings of Candès and Tao [6], which optimally characterize the asymptotic bounds for reconstructability of a coherent true matrix, settle, once and for all, the problem of matrix completion and all that can be known about it.

However, examining the literature more carefully, part of the theoretical and practical findings have already shown that the structural and computational phenomena in matrix completion are far from being understood. On the theoretical side, for example, Singer and Cucuringu [43] have tried to analyze the identifiability of matrix completion from a combinatorial point of view. While their work, which relates matrix completion to the Framework Realizability Problem for rigid bar-joint frameworks, remains mainly conjectural, they are able to give conjectural statements and algorithms on the completability on partially known matrices which do not rely on the convex optimization setting but only on combinatorial features which were also observed in different contexts, see Oh [34]. On the other hand, the practical findings in the existing literature are also far from being complete. While the existing results give rise to algorithms good asymptotic guarantees, they often fail for the case of small matrices or small samples.

In this paper, we will explore both of these white spots on the map by taking into account the intrinsic structure of the matrix completion problem which has, until now, only been addressed to a marginal extent. It will turn out that matrix completion does not only have deep relations to functional analysis, as it has already been observed, e.g., by Candès and Recht [5], Candès and Tao [6], but also with combinatorial commutative algebra, algebraic geometry, matroid theory, deterministic and probabilistic graph theory, and percolation theory. By combining these contributions to a closed whole, we obtain what we believe to be the right tools to investigate theoretical and practical aspects of matrix completion and a more general class of compressed sensing problems which exhibit combinatorial algebraic structure.

Here is a summary of our main contributions in this paper, which can also serve as a guide for reading:

- In sections 2.1 and 2.2, we express the problem of matrix completion in terms of noisy parametric estimation and, for the first time, explicitly separate the generative model for the true matrix and the measurement process. This central part allows us to treat the

properties of the measurement separately from the properties of the matrix; in fact, the genericity formalism introduced section 2.2 will allow to remove the influence of the true matrix almost completely in identifiability considerations.

- In section 2.3, under the assumption of generic sampling introduced in section 2.2, we apply, following some ideas from the paper Király and Tomioka [25] and some new, elementary techniques from Algebraic Geometry and Combinatorial Commutative Algebra which allow us to parameterize and characterize the measurement process by a bipartite graph. We show that all properties of the measurement, as for example degrees of freedom, as well as identifiability, are completely encoded by this graph and its algebraic properties, thus accessible to algebraic-combinatorial methods. As a practical counterpart, this implies that whether the true matrix can be reconstructed does, generically, not depend on the values at the known entries but only their position.
- In section 2.4, we introduce a problem which is analogous to matrix completion. Instead of asking for a reconstruction of all missing elements of a matrix, we ask for a reconstruction of some. In particular, we ask the reconstruction of which missing entries is in principle possible. This task, which we term partial matrix completion, has apparently not appeared in the literature yet, but is amenable to the techniques developed in the previous chapters, and, in our opinion, of high practical relevance since in general, not all entries are to be reconstructed. Our results include the fact that the set of entries which can be reconstructed from the measurements also depends only on the positions of the known entries, not their values.
- In section 2.5, we introduce, for the first time, combinatorial algebraic tools which allow us to practically characterize the degrees of freedom which are destroyed by the measurement process in terms of the graph defined in section 2.3. Matroid theory allows to further characterize the patterns which guide possible ways of reconstruction, and the theory of differentials gives a grasp on their calculation.
- The theory developed in section 2.5 gives rise to several randomized algorithms, later presented in section 3.1 which are efficiently able to determine which missing entries can in principle be reconstructed, and which not. A special case is the conjectural algorithm proposed by Singer and Cucuringu [43], but also includes more general applications including the partial matrix completion problem introduced in section 2.4. In particular, we present an algorithm which answers the question which entries can be in principle reconstructed, given the positions of the known entries.
- The analysis of a special reconstruction pattern discussed in section 2.5 motivates a novel algorithm which can perform reconstruction on the bipartite graph counterpart, which we describe in section 3.2. By adding algebraic calculations onto the purely graph theoretical foundations, we obtain in section 3.3 an first-of-its-kind-algorithm which performs matrix completion not via optimization, but via Combinatorial-Algebraic calculations.
- Since a graph parameterizes the measurement process, random measurements correspond to random graphs. In section 2.6, we formalize this correspondence and, as an application obtain bounds for the number of measurements which is necessary for reconstruction. Also, we provide theoretical and conjectural evidence for phase transition phenomena which can

be and have been observed in practical matrix completion settings, as well as explanations for why and how classical matrix completion algorithms fail on certain classes of measurements.

- In section 2.1, it has been shown that the conditioning on the true mask can be removed in the analysis of identifiability. In section 2.7, we show that this is a general principle: we develop a theory for compressed sensing under algebraic compression constraints (i.e., the constraints can be expressed by polynomial equations) and prove an upper probabilistic bound on the sufficient number of samples which are needed for reconstruction, which only depends on properties of the constraints, and not on the properties of the true signal. For the special case of matrix completion, we obtain sufficient bounds which are similar to those of Candès and Recht [5] and Candès and Tao [6], but now with the conditioning on the incoherence of the true matrix completely removed.
- In the experiment section we underline our theoretical findings, conjectures, and practical claims with evidence from simulations. In section 4.1, we show how the number of reconstructible entries behaves with increasing sample size, and in section 4.2, we compare the various theoretically predicted and practically observed phase transitions of in the Matrix Completion problem, amongst those the identifiability phase transition. Moreover, we compare the performance of the various known and novel algorithms to these phase transitions.
- Appendix A contains a technical treatise on sheaves of matroids on schemes (e.g., algebraic varieties). It summarizes the genericity properties of a matroid of sections, when evaluated at different points of the scheme. While the results presented there are folklore and maybe not surprising, we decided to have them included since they seem not to be written up in the existing literature.

Summarizing, this paper contains many results of theoretical and practical type, which do not necessarily need to be read in sequence to be understood. Sections 2.1 and 2.2 are fundamental and suggested reading, while sections 2.3, 2.6 or 2.7 can serve as independent starting points. Also, the algorithms should be accessible (though not completely understandable) without having read the theory part.

1.3 Acknowledgements

We thank Andriy Bondarenko, Winfried Bruns, Mihyun Kang, Yusuke Kobayashi, Martin Kreuzer, Cris Moore, Klaus-Robert Müller, Kazuo Murota, Kiyohito Nagano, Zvi Rosen, Raman Sanyal, Bernd Sturmfels, and Günter Ziegler for valuable discussions. RT is partially supported by MEXT KAKENHI 22700138, the Global COE program “The Research and Training Center for New Development in Mathematics”, FK by Mathematisches Forschungsinstitut Oberwolfach (MFO), and LT by the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement no 247029-SDModels.

2. Theory of Low-Rank Matrix Completion

In this section, we formulate low-rank matrix completion as a parametric estimation task. Then, different approaches of sampling are discussed, arguing that the generic algebraic framework is

the most proper way of approaching the problem. Subsequently, novel Algebraic Combinatorial tools will be derived for exploiting the inherent structure of low-rank matrix completion, which allow to construct methods to solve and understand the features of matrix completion from the structural, algebraic point of view.

2.1 What is Low-Rank Matrix Completion?

Matrix completion is the task of imputing missing entries in a matrix, given other, known, but possibly noisy entries. low-rank matrix completion is doing that under assumption of a low-rank model, that is, informally:

Problem 2.1.1. *Let A be a matrix, in which some entries are missing, and some entries are known. Given some target rank r , find a matrix A' of rank r , close to A .*

From both a mathematical and procedural point of view, Problem 2.1.1 is ill-defined. The standard way of parameterizing and well-posing the matrix completion model is assuming a generative truth, i.e., that there exists a *true matrix* A , of which one *observes* some entries, plus observation noise. Thus, under the generative assumption, Problem 2.1.1 reformulates to

Problem 2.1.2. *Let $A \in \mathbb{C}^{m \times n}$ be an unknown matrix, of known rank r . Let $\varepsilon \in (\mathbb{C} \cup \{\infty\})^{m \times n}$ be a noise matrix. Given the observed matrix $A + \varepsilon$, reconstruct A .*

In this description of low-rank matrix completion, the model of the truth is well-defined, but without assumptions on A and the noise ε , it is practically useless. Thus, as it is common practice in statistics and learning theory, in order to obtain a proper, well-defined and practically applicable model, one needs to

- (i) separate the *generative model* from the *noise model*. That is, separate the fact which entries are observed from the accuracy with which they are observed, if they are observed.
- (ii) specify the *generative sampling model*. That is, introduce and specify random variables for sampling A , and the set of observed entries.
- (iii) specify the *observational noise model*. That is, introduce and specify random variables for the noise ε .

Also, measures of closeness from the observations to the putative truth need to be defined when attempting to reconstruct and give reconstruction guarantees; however these are explicitly not part of the model, but also need the above three points to be fulfilled to allow for proper evaluation.

2.1.1 Separating Generative and Noise Models

In order to perform (i) the separation of generative and noise models in a sound way, we need to introduce mathematical notation to parameterize the generative sampling process. First, we introduce notation for the set of all low rank matrices, from which the truth will be sampled:

Notations 2.1.3. *The set of all complex $(m \times n)$ -matrices of rank r or less will be denoted by*

$$\mathcal{M}(m \times n, r) = \{A \in \mathbb{C}^{m \times n}; \text{rk}A \leq r\}.$$

We will always assume that $m \leq n$; by transposing the matrices, this is no loss of generality.

It will become important later that the set $\mathcal{M}(m \times n, r)$ is the solution of a set of polynomial equation in the matrix entries - the minor equations - making $\mathcal{M}(m \times n, r)$ an algebraic variety. $\mathcal{M}(m \times n, r)$ is often called *the* determinantal variety¹, while subsets of it may be called a determinantal variety.

Next, we need to introduce notation for the process of specifying the observed entries of the matrix:

Definition 2.1.4. A map $\Omega : \mathbb{C}^{m \times n} \rightarrow \mathbb{C}^a$ which sends a matrix to a fixed tuple of its entries, i.e.

$$\Omega : (a_{ij})_{\substack{1 \leq i \leq m \\ 1 \leq j \leq n}} \mapsto (a_{i_1 j_1}, a_{i_2 j_2}, \dots, a_{i_a j_a}),$$

is called *masking in rank r* . Such a map is uniquely defined by the set of entries $i_k j_k$ in the image set. When clear from the context, we omit the qualifier “in rank r ”.

We call the unique matrix which has ones at those entries, and zeroes elsewhere, the *mask* of Ω and denote it by $M(\Omega)$. Similarly, we will call a matrix M having only ones and zeroes a *mask*, and the map Ω such that $M(\Omega) = M$ the *masking associated to M* . When no confusion is possible, we will denote it by Ω_M and implicitly assume that the rank r is fixed.

Note that Definition 2.1.4 allows for an entry to be observed several times; that may be useful if the observation is noisy. However, in the rest of the paper, we will not explicitly make use of this fact, so we will assume that no entry is observed twice, i.e., the pairs (i_ℓ, j_ℓ) are all different.

Naturally, we will be interested in the behavior of a masking Ω when its range is restricted to the low-rank matrices $\mathcal{M}(m \times n, r)$. Before proceeding to reformulating the matrix completion model, we give examples for the definitions above:

Example 2.1.5. For any $m, n \in \mathbb{N}$, one has $\mathcal{M}(m \times n, m) = \mathbb{C}^{m \times n}$.

The simplest non-trivial examples of determinantal varieties are those containing matrices of rank at most one:

$$\mathcal{M}(n \times n, n - 1) = \{A \in \mathbb{C}^{n \times n} ; \det A = 0\}.$$

For example, the co-rank one (2×2) -matrices are

$$\mathcal{M}(2 \times 2, 1) = \left\{ \begin{pmatrix} a_{11} & a_{21} \\ a_{12} & a_{22} \end{pmatrix} ; a_{11}a_{22} = a_{12}a_{21} \right\}.$$

Example 2.1.6. Consider a true (3×3) -matrix

$$A = \begin{pmatrix} a_{11} & a_{21} & a_{31} \\ a_{12} & a_{22} & a_{32} \\ a_{13} & a_{23} & a_{33} \end{pmatrix} = \begin{pmatrix} 1 & 2 & 3 \\ 2 & 4 & 6 \\ 4 & 8 & 12 \end{pmatrix}$$

which has rank one. By observing five entries (exactly, i.e., without noise), one may arrive at one of the following two partial matrices:

$$A_1 = \begin{pmatrix} 1 & 2 & 3 \\ & 4 & \\ 4 & & \end{pmatrix} \quad \text{or} \quad A_2 = \begin{pmatrix} 1 & & 3 \\ & 4 & \\ 4 & & 12 \end{pmatrix}.$$

¹To be more precise, the usual determinantal variety is the projective closure of the affine variety $\mathcal{M}(m \times n, r)$ we define. However, the generic behavior under algebraic maps, including fiber and image dimensions, does not change fundamentally when restricting the projective morphisms to the images and pre-images of the affine variety $\mathcal{M}(m \times n, r)$ which is dense in its projective closure.

The masks corresponding to the two matrices are

$$M_1 = \begin{pmatrix} 1 & 1 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{pmatrix} \quad \text{and} \quad M_2 = \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 1 \end{pmatrix},$$

so the corresponding maskings are the maps

$$\Omega_1 : \begin{pmatrix} a_{11} & a_{21} & a_{31} \\ a_{12} & a_{22} & a_{32} \\ a_{13} & a_{23} & a_{33} \end{pmatrix} \mapsto \begin{pmatrix} a_{11} & a_{21} & a_{31} \\ & a_{22} & \\ a_{13} & & \end{pmatrix}$$

and

$$\Omega_2 : \begin{pmatrix} a_{11} & a_{21} & a_{31} \\ a_{12} & a_{22} & a_{32} \\ a_{13} & a_{23} & a_{33} \end{pmatrix} \mapsto \begin{pmatrix} a_{11} & & a_{31} \\ & a_{22} & \\ a_{13} & & a_{33} \end{pmatrix}.$$

In particular, one has $M_i = M(\Omega_i)$, and $A_i = \Omega_i(A)$ for $i = 1, 2$. Also note that Ω_i could have been expressed by the map which sends A to the Hadamard product, i.e., the componentwise product, $A \circ M_i$.

Thus, the generative sampling process is modelled by applying some masking Ω to the true matrix which is in the range, while the noise acts on the image of Ω , formally separating both by the map given by Ω .

2.1.2 The Parametric Model of Low-Rank Matrix Completion

Using the notations and definitions introduced in section 2.1.1, we can now provide a complete model description for low-rank matrix completion:

Problem 2.1.7. *Let r be the true rank, let A be a $\mathcal{M}(m \times n, r)$ -valued random variable, modelling the sampling of the true matrix. Let M be a $(m \times n)$ -mask-valued random variable modelling the position of the observed entries, and $\alpha = \|M\|_1$ the integer-valued random variable which is the number of observed entries. Let ε be a \mathbb{C}^α -valued random variable, modelling the noise.*

Then, construct a good estimator for A which takes $\Omega_M(A) + \varepsilon$ as input.

What is chosen for particular sampling distributions on A and M is different throughout the literature, similarly the noise model ε . While choices for sampling A and M will be thoroughly discussed in the next section, we will not put much emphasis on the noise model ε yet, while acknowledging that it is extremely important for practical purposes. However, identifiability of the generative model is independent of the noise ε while it is well-behaved; in fact one direction of this claim is straightforward to see, and we summarize it in the following important remark:

Remark 2.1.8. *If it is impossible to identify A from $\Omega(A)$, then there can exist no consistent estimator for A which makes no use of hidden knowledge.*

Even more can be seen to be true, as the following result shows:

Theorem 2.1.9. *Keep the notations in Problem 2.1.7. There exists a consistent², deterministic estimator for the true matrix A if and only if A is identifiable from $\Omega(A)$.*

Proof. No consistent estimator for A can exist if A is not identifiable from $\Omega(A)$, since the estimator cannot know which element of $\Omega^{-1}(\Omega(A))$ to estimate, as all elements in $\Omega^{-1}(\Omega(A))$ are indistinguishable by the assumptions, as was already discussed in Remark 2.1.8.

It remains to prove that there exists an algorithmic estimator $\widehat{A}(B)$ which estimates A from $B = A + \frac{\varepsilon}{\sqrt{N}}$ consistently, assuming A is identifiable from $\Omega(A)$, and $\varepsilon \in \mathbb{C}^a$ is centered noise. Consistency means, as discussed in the footnote, that $\widehat{A}(B) \rightarrow A$ in probability, as $N \rightarrow \infty$. Note that both \widehat{A} and B depend on N , but for ease of notation, we do not make this explicit.

Instead of constructing the estimator \widehat{A} directly, we will first construct a consistent estimator \widehat{a} for a fixed missing entry a of A , which will then be modified into a consistent estimator of \widehat{A} . Note that taking \widehat{a} component-wise does not suffice here, as this matrix needs not to be of rank r . The proof will be split in three steps: (i) the construction of the consistent estimator \widehat{a} for a , and (ii) the proof that \widehat{a} estimates a consistently if A is identifiable, and (iii) the construction of the estimator \widehat{A} from \widehat{a} , and the proof of its consistency.

(i) For constructing the estimator \widehat{a} , we will make use of the fact that Ω is an algebraic map. This implies that the fiber $\Omega^{-1}(\Omega(A))$ is an algebraic variety, which is moreover finite (=zero-dimensional) since $\{A\} = \Omega^{-1}(\Omega(A))$. The latter also implies that there is a finite set of polynomials f_1, \dots, f_k such that a is the unique solution of the equations

$$f_1(\Omega(A), X) = 0, \dots, f_k(\Omega(A), X) = 0$$

in X , where the f_i have to be read as polynomials in the coordinate entries of $\Omega(A)$ and a a (and not matrix polynomials in $\Omega(A)$ and a). The polynomials f_1, \dots, f_k can be algorithmically obtained by symbolic (though highly complex) calculations, e.g., symbolic elimination of all variables from A not in $\Omega(A)$ except a . Thus, for each i , the equation

$$f_i(B, a) = 0$$

is a polynomial equation, let d_i be its degree in a . Given B , we can solve for a by numerical algorithms, yielding solutions

$$a_{i1}, \dots, a_{id_i},$$

possibly with multiplicity. We now determine $h : \mathbb{N} \rightarrow \mathbb{N}$ minimizing the value

$$L(h) = \sum_{i,j=1}^k \left\| a_{i,h(i)} - a_{j,h(j)} \right\|,$$

which is algorithmically possible as there are only finitely many choices for h . Then, we set

$$\widehat{a} = \frac{1}{k} \sum_{i=1}^k a_{i,h(i)}.$$

²here, consistent is defined by the variance convention: for the observed matrix $B = \Omega(A) + \frac{\varepsilon}{\sqrt{N}}$, some scaling factor N , and centered noise ε with finite variance, the estimator $\widehat{A}(B)$ converges in probability to A for $N \rightarrow \infty$. This is equivalent to observing each noisy entry of $\Omega(A)$ with multiplicity N and taking the number of samples convention for consistency in the number of observations.

(ii) We prove consistency of the estimator \widehat{a} constructed in (i). As $N \rightarrow \infty$, it holds that $B \rightarrow \Omega(A)$ and thus that $f_i(B, X) \rightarrow f_i(\Omega(A), X)$, in terms of the coefficients. Thus, the $a_{i_1}, \dots, a_{i_{d_i}}$, converge (with multiplicities) to the solutions $\alpha_{i_1}, \dots, \alpha_{i_{d_i}}$, of the equation $f_i(\Omega(A), X) = 0$ in the variable X , up to renumbering. Thus, it holds that

$$L(h) = \sum_{i,j=1}^k \|a_{i,h(i)} - a_{j,h(j)}\| \rightarrow \sum_{i,j=1}^k \|\alpha_{i,h(i)} - \alpha_{j,h(j)}\| =: \ell(h).$$

As discussed in (i), a is the unique solution of the equations

$$f_1(\Omega(A), X) = 0, \dots, f_k(\Omega(A), X) = 0,$$

so it holds that $\ell(h) = 0$ if and only if

$$\alpha_{i,h(i)} = a \quad \text{for all } 1 \leq i \leq k,$$

and $\ell(h) > 0$ else. Since there are only finitely many choices for h , it holds that

$$\widehat{a} = \frac{1}{k} \sum_{i=1}^k a_{i,h(i)} \rightarrow \frac{1}{k} \sum_{i=1}^k \alpha_{i,h(i)} = \frac{1}{k} \sum_{i=1}^k a = a$$

for $N \rightarrow \infty$, which proves consistency of \widehat{a} .

(iii) We now construct the estimator \widehat{A} , first we give its explicit form. First, consider the estimator \widehat{A}_0 which is B for the known entries and \widehat{a} for each unknown entry. By (ii), it holds that $\widehat{A}_0 \rightarrow A$ as $N \rightarrow \infty$, but the estimator \widehat{A}_0 is in general not consistent as it is not contained in $\mathcal{M}(m \times n, r)$. This can be remedied by replacing \widehat{A}_0 with its rank- r -Young-Eckart approximation, which is algorithmically done as follows:

First, calculate the SVD of \widehat{A}_0 , i.e., a decomposition

$$\widehat{A}_0 = U \cdot D \cdot V^\top$$

into an orthogonal $(m \times m)$ -matrix U , an orthogonal $(n \times m)$ -matrix V , and a diagonal $(m \times m)$ -matrix D , where we assume without loss of generality that $m \leq n$, and that the diagonal entries of D are ordered descendingly by absolute value. Then, let D' be the matrix where we replace all except the r entries of D with biggest absolute value by 0. The rank- r -Young-Eckart approximation of \widehat{A}_0 is then the matrix $U \cdot D' \cdot V^\top$. We set

$$\widehat{A} := U \cdot D' \cdot V^\top.$$

The Young-Eckart theorem states that \widehat{A} minimizes the Frobenius distance

$$\|\widehat{A} - \widehat{A}_0\|_2$$

under the constraint that $\text{rk} \widehat{A} = r$. Since $\mathcal{M}(m \times n, r)$ is a closed set, this implies that $\widehat{A} \rightarrow A$ for $N \rightarrow \infty$, which proves consistency of \widehat{A} . □

In particular, Theorem 2.1.9 shows: if A cannot be identified from $\Omega(A)$, then no algorithm without hidden knowledge can reconstruct all the missing entries in $\Omega(A)$.

This stringently motivates the analysis of properties of Ω alone, since the statement is independent of the noise ε under the condition of well-behavedness of the latter.

The estimator given in the proof of Theorem 2.1.9 is in general very inefficient and is proposed only for the purpose of completing the statement which relates the exact algebraic morphism to identifiability of a statistical problem. In fact, we do not expect that there is a much better algorithm that applies to *any* true matrix. In his PhD thesis, Michael Dobbins [10]³ proved, via a reduction to the Polytope Realization Problem [37, Parts I–II] that:

Proposition 2.1.10. *Deciding if a partial matrix has a low-rank completion is as hard as deciding if any set of polynomial equations has a solution.*

This problem is known to be in PSPACE, is at *least* NP-hard. (This particular formulation of hardness, as well as a discussion from the complexity-theoretic perspective is in Shor [42].) Thus, to obtain efficient algorithms, we will need to make *some* kind of assumption on the input. Ours will be that the true matrix A is *generic*, a concept that we describe next.

2.2 Genericity

In this section, we discuss sampling assumptions for the generative model of the mask M defining the masking Ω , and the true matrix A . We introduce a new, algebraic genericity assumption for the true matrix, which will allow later to remove the influence of the sampling of the true matrix A onto the behavior of Ω . As in applications, the mask is known, while the sampling assumptions on the true matrix A and the true matrix A itself are in general unknown, we will argue that this is at the same time the most natural and weakest possible assumption on the sampling of the true matrix A .

2.2.1 Sampling of the Mask

Several ways of sampling the mask have been considered in the literature. Table 1 gives a short list of sampling distributions for the mask M . There are also different and less common sampling

fixed mask	M	the mask M is fixed
uniform	$U(m \times n, k)$	the number of observed entries $\alpha = \ M\ _1$ is fixed, otherwise the sampling is uniform
Erdős-Rényi	$G(m \times n, p)$	each entry of M is independently Bernoulli distributed with parameter p

Table 1: Possible sampling assumptions for the mask

assumptions like the CLP or the Power Law model, which will however not be discussed in this paper. In the Erdős-Rényi model, the number of edges $\alpha = \|M\|_1$ is binomially distributed as $\alpha \sim \mathcal{B}(p, mn)$, with expected value $mn \cdot p$ and variance $mn \cdot p(1 - p)$. Thus, the relative variance, i.e., the variance of $\frac{\alpha}{mn}$, is $\frac{p(1-p)}{mn}$ which approaches zero as $mn \rightarrow \infty$, so the qualitative behavior of $G(m \times n, p)$ will approach that of $U(m \times n, mn \cdot p)$ in the limit.

³We thank Günter Ziegler for reminding us about Dobbins’s results.

Note that in practical applications, it is always possible to identify the result of the mask sampling, since it is - tautologically - always known which entries of the true matrix are known, and which entries are unknown.

2.2.2 Sampling of the True Matrix

Table 2 gives a short list of definitions for sampling assumptions on the true matrix A . Note that usually, the specific distribution from which the sampling occurs is not specified, only properties are specified which hold for the sampling distribution or the sampled matrix. The reason for this is that the relevant statements hold for any sampling distribution fulfilling these assumptions.

incoherent	for the singular value decomposition $A = USV^\top$, and a global constant C , it holds that $\ U\ _0 \leq \frac{C}{\sqrt{m}}$ and $\ U\ _0 \leq \frac{C}{\sqrt{n}}$
non-spiky	there exists a global constant C bounding the spikiness ratio from above, i.e., it holds that $mn \cdot \ A\ _\infty^2 \cdot \ A\ _2^{-2} \leq C$
(Zariski-)generic	(algebraic) subsets with Lebesgue/Hausdorff measure zero have (conditional) probability zero

Table 2: Possible sampling assumptions for the mask

The most common strategy is to restrict the sampling of matrices to a subset of all matrices, such as incoherence [5] or non-spikiness [33]. We propose a weaker condition, inspired by the Zariski topology⁴ on algebraic sets: we only postulate that the sampling distribution of the true matrix assigns no positive probability to any Hausdorff measure zero subset of $\mathcal{M}(m \times n, r)$. For what concerns the following, one could also postulate that only for irreducible algebraic subsets. This Zariski-like condition is indeed weaker, as the coherent, or spiky matrices form a set of positive probability measure in general. In fact, any continuous probability distribution fulfills the postulate, in particular also any distribution supported on non-spiky or coherent matrices.

As the underlying sampling process is unknown in practical applications, as opposed to the chosen mask, we argue that the proposed sampling for true matrices is the weakest possible condition which excludes sampling degeneracies. Thus, we will term it *generic* sampling, and any random variable fulfilling this condition will be termed generic, or generic matrix. In the next section we will see in fact that under generic sampling assumptions, the behavior of the masking operation and its identifiability depends only on the properties of the mask, and is completely independent of properties of the true matrix.

Furthermore, different scenarios restrict to *symmetric/Hermitian* or *antisymmetric/anti-Hermitian* matrices, and/or *real* matrices, as opposed to general (i.e., not necessarily symmetric or Hermitian) complex matrices. We will discuss these sampling assumptions additionally when appropriate. Other sampling assumptions include definiteness or sign assumptions on the true matrix. As these conditions are semi-algebraic and completely change the flavor of the problem, they will not be discussed in this paper.

⁴In the Zariski topology, the closed subsets of \mathbb{C}^n are exactly those which can be written as set of solutions for a finite set of polynomials; i.e., $U \subseteq \mathbb{C}^n$ is open if and only if there are polynomials p_1, \dots, p_m in n variables such that $U = \{x \in \mathbb{C}^n ; p_i(x) \neq 0 \text{ for some } 1 \leq i \leq m\}$. The closed sets are called algebraic sets, and carry the inherited Zariski topology. Zariski closed sets and relatively Zariski closed sets have probability measure zero under any continuous random variable, see the appendix of [27] for more details.

For formal purposes, we want to state our definition of (algebraic) genericity:

Definition 2.2.1. *Let P be some property on the matrices $\mathbb{C}^{m \times n}$, such that the set X of P matrices is admits a Hausdorff measure (e.g., when X is an algebraic variety), let Q be any property on the matrices $\mathbb{C}^{m \times n}$, let Y be the set of matrices that are not Q . We say that*

“A generic P matrix is Q ”

if the set $X \cap Y$ is a negligible set (i.e. a subset of a null set) under the Hausdorff measure on X .

The given definition is a bit broader than the usual concept of genericity applied for moduli spaces of algebraic objects, but is, for the current setting, maybe the most intuitive one, without making any difference in the consequences, since it implies that any matrix valued continuous random variable with support on the P matrices will fulfill Q almost surely. Indeed, in the case that P and Q define algebraic sets, the definitions (the given one, and very generic/general for the moduli space) are equivalent. A more detailed comparison and relation of different concepts of genericity, and how they imply each other, can be found in the appendix of [27].

To give further intuition for this concept of genericity, and its meaning in the world of low-rank matrices, we give some examples for valid statements:

Example 2.2.2. *Recall that we have assumed $r \leq m \leq n$.*

- *A generic $(m \times n)$ -matrix has only non-zero entries.*
- *Let A be any fixed $(m \times n)$ -matrix of rank r . A generic $(m \times n)$ -matrix of rank r is not equal⁵ to A .*
- *A generic $(m \times n)$ -matrix of rank r or less has rank exactly r .*
- *A generic $(m \times n)$ -matrix of rank r has no vanishing $(r \times r)$ -minor.*
- *A generic $(m \times n)$ -matrix of rank r has no real eigenvalues.*
- *A generic positive semi-definite real $(m \times n)$ -matrix of rank r is positive definite.*

These statements can be proved using algebraic, or probabilistic methods with a proper conditioning. Note that the use of “generic” is not equivalent to the use of “in general”, since, for example, a matrix of rank r or less in general needs not to have rank r . Also, note, that a generic $(m \times n)$ -matrix is not a single, fixed matrix, as generic is not a qualifier for a single, fixed matrix; it is, to the contrary, in fact a qualifier for statements about almost all $(m \times n)$ -matrices. Similarly, it can be thought of as a $(m \times n)$ -matrix-valued random variable having the generic sampling assumption, about which the sentences make probability one statements.

⁵ This might seem as a paradox which implies that generic matrices do not exist. First, as explained below, the word “generic” is not a qualifier for a fixed matrix, but a qualifier for an almost-true statement. The reason that this paradoxically seeming statement is true is then in complete analogy to the fact that the probability to sample a fixed value from a continuous random variable being zero does not imply that this random variable has empty range of definition.

2.3 The Algebraic Combinatorics of Matrix Completion

This section is devoted to developing the basic algebraic structure of low-rank matrix completion, and to state some elementary results which come from the mere fact that low-rank matrix completion is an algebraic problem in the strict sense. Namely, the generative mapping Ω_M , as it occurs in the final problem description 2.1.7, is an algebraic map. This makes the analysis of the generative model of matrix completion, and its identifiability in view of Theorem 2.1.9 amenable to the vast range of tools from commutative algebra and algebraic geometry. A comprehensive introduction into basic algebraic geometric concepts can be found in the introductory chapters of the book by Cox et al. [9]. Note that knowledge of advanced concepts of commutative algebra, or algebraic geometry should not be necessary for understanding the current paper apart from some proof details. Part of the following exposition follows the results obtained in the paper by Király and Tomioka [26].

We motivate the theory which follows with a central question about the identifiability of matrix completion:

Question 2.3.1. *Given sampling conditions on the true matrices (including the true rank) and a fixed mask M : when is the generative model of matrix completion identifiable?*

Theorem 2.1.9 states that identifiable of the generative model is equivalent to injectivity of the masking Ω_M , under the given sampling conditions. Thus, the first question which has been asked about identifiability of the model is the following:

Question 2.3.2. *Given a fixed mask M , when is Ω_M injective?*

Injectivity, or one-to-one-ness, is by construction the necessary and sufficient condition for properly inverting a map, and thus for identifiability of the generative model.

For the community, it has been long known that the answer to Question 2.3.2 is very unsatisfactory - it is, basically: in all interesting cases, never. The following proposition, which, together with its proof, is taken from [26], gives the corresponding formal statement which has already been asserted by Candès and Recht [5].

Proposition 2.3.3. *Let $r \geq 2$, let M be a mask with $\alpha = \|M\|_1$ known entries. Then the masking $\Omega : \mathcal{M}(m \times n, r) \rightarrow \mathbb{C}^\alpha$ is injective if and only if $\alpha = mn$.*

Proof. Clearly, if $\alpha = mn$; then Ω is injective, as it is the identity map. So it suffices to prove: if $r \geq 2$ and $\alpha < mn$, there exists a matrix A such that $\{A\} \neq \Omega^{-1}(\Omega(A))$. Now since $\alpha < mn$, there exists an index ij such that $M(\Omega)_{ij} = 0$. Let A be any matrix whose columns, except the j -th, span an $(r - 1)$ -dimensional vector space. Since X is of (at most) rank r , the set $\Omega^{-1}(\Omega(A))$ contains any matrix \tilde{A} which is identical to A but has an arbitrary entry at the index ij . \square

This answer to Question 2.3.2, which seems to be strongly negative, is maybe the main reason which has led the community to believe that in order to obtain firm results on the generative model of matrix completion, the sampling of the true matrices has to be restricted. As discussed in section 2.2.2, most sampling models from literature put rather strict assumptions on the sampling of the true matrix. As such, the obtained results usually mix the generative sampling model for the true matrix and the mask, which makes in the end unclear which of the two is at the source of the observed phenomena.

2.3.1 The Algebra of Matrix Completion

The following argumentation, which is naturally and intuitively obtained from the algebraic structure of the problem, shows that the strong and unnatural conditioning on the true matrix - which can furthermore, by construction, never be verified in the real world - is not necessary to obtain identifiability results for the matrix completion problem. Assuming merely generic sampling, which means, for any sampling process of true matrices having no support at null sets, one can show that identifiability of matrix completion depends only on the mask and the true rank, and not on the true matrix, or any further sampling assumptions on the true matrix.

In order to state the result compactly, we need notation for characterizing the degrees of freedom one has in the reconstruction:

Definition 2.3.4. *Let A be an $(m \times n)$ -matrix of rank r , let Ω be an $(m \times n)$ -masking in rank r . Then, the set $\Omega^{-1}(\Omega(A))$ is called the fiber of Ω at A , and, alternatively, the fiber of Ω over $\Omega(A)$. We will call the integer*

$$\dim_A \Omega = \dim \Omega^{-1}(\Omega(A))$$

the fiber dimension (of Ω) at A . Similarly, we will call

$$\#_A \Omega = \# \Omega^{-1}(\Omega(A))$$

the fiber cardinality (of Ω) at A .

For given A , the fiber dimension $\dim_A \Omega$ is exactly the number of degrees of freedom one has in continuously deforming A without changing the observed entries $\Omega(A)$ or the rank. That is, the fiber dimension is the number degrees of freedom in the range of Ω which do not appear in its image at the element A . Even more informally, the fiber dimension is the degrees of freedom killed by Ω in a neighborhood of A . In particular, $\dim_A \Omega = 0$ is equivalent to saying that $\Omega(A)$ has a finite set of possible reconstructions. In this case, $\#_A \Omega$ is an integer, and otherwise it is ∞ .

The following series of observations implies that the invariants introduced in Definition 2.3.4 are indeed generic and characteristic invariants of the masking. They follow from the fact that Ω is an algebraic map. The most important fact is that for a generic true matrix, its behavior, in terms of fiber dimension and number of possible reconstructions, does not depend on the particular entries or the structure of the true matrix. These results were first stated and proved in [26], for completeness, we reproduce the proofs. Candès and Recht [5] have implicitly used and assumed some of those.

Theorem 2.3.5. *Let A be a generic $(m \times n)$ -matrix of rank r , let Ω be a masking in rank r . Then, the following depend only on the true rank r and the mask $M(\Omega)$:*

- (i) *The fiber dimension $\dim_A \Omega$.*
- (ii) *The fiber cardinality $\#_A \Omega$.*
- (iii) *Whether $\#_A \Omega = 1$, i.e., whether A is uniquely completable.*

Proof. (i) By definition of genericity, it suffices showing that there is a Zariski open dense set U in $\mathbb{C}^{m \times n}$, such that for all matrices $A \in U$, the set of possible completions $\Omega^{-1}(\Omega(A))$ has the same dimension and cardinality. But this is a direct consequence of the upper semicontinuity theorem (see e.g. I.8, Corollary 3 in Mumford [31]), when applied to the algebraic map

$\Omega : \mathcal{M}(r; m \times n) \longrightarrow \mathbb{C}^\alpha$, considering that $\mathcal{M}(m \times n, r)$ is irreducible.

(ii) In the case of $\dim_A \Omega > 0$, one has $\#_A \Omega = \infty$ and the statement is true. If $\dim_A \Omega = 0$, the statement is an application of the purity of the branch locus, see [30, Chapter 1, Proposition 3.8].

(iii) is a special case of (ii). □

Theorem 2.3.5 shows that identifiability properties of Ω are independent of the true matrix A as long as it is generically sampled; namely $\dim_A \Omega$ and $\#_A \Omega$ are independent from A , so we can remove the qualifier A which Theorem 2.3.5 has proved to be unnecessary:

Definition 2.3.6. *Let Ω be a $(m \times n)$ -masking in rank r , let A be a generic $(m \times n)$ -matrix. We will write $\#\Omega$ for the (generic) value of $\#_A \Omega$, and $\dim \Omega$ for the (generic) value of $\dim_A \Omega$.*

Also, Theorem 2.3.5 provides motivation for the following definitions which characterize the generative identifiability of Ω generically:

Definition 2.3.7. *Let Ω be a $(m \times n)$ -masking in rank r . We call*

- (i) Ω generically injective and $\mathcal{M}(\Omega)$ identifiable or completable (in rank r), if $\#\Omega = 1$.
- (ii) Ω (generically) finite and $\mathcal{M}(\Omega)$ finitely identifiable or finitely completable (in rank r), if $\#\Omega < \infty$.
- (iii) Ω generically k -to-one (in rank r), if $\#\Omega = k$.
- (iv) Ω infinite and $\mathcal{M}(\Omega)$ unidentifiable (in rank r), if $\dim \Omega > 0$.

Thus, generic injectivity of Ω means that Ω is 1:1 almost everywhere on its range; that is, after a restriction to the complement of a Lebesgue null set. Similarly, generic finiteness and generic k -to-one-ness mean being k :1 everywhere, with the same k . Theorem 2.3.5 ascertains that every masking Ω will be either generically injective, generically k -to-one for some $k \geq 2$, or infinite. Note that the qualifier in rank r has always to be present for well-definedness.

For intuition, we illustrate the results in Theorem 2.3.5 with the simplest non-trivial (in the sense of Proposition 2.3.3) example:

Example 2.3.8. *Consider the set $\mathcal{M}(2 \times 2, 1)$ of (2×2) -matrices of rank 0 and 1. It is the set*

$$\mathcal{M}(2 \times 2, 1) = \left\{ \left(\begin{array}{cc} a_{11} & a_{21} \\ a_{12} & a_{22} \end{array} \right) \in \mathbb{C}^{2 \times 2}; a_{11}a_{22} = a_{12}a_{21} \right\}.$$

Consider the mask

$$M = \left(\begin{array}{cc} 1 & 1 \\ 1 & 0 \end{array} \right)$$

The masking given by M is

$$\Omega_M : \left(\begin{array}{cc} a_{11} & a_{21} \\ a_{12} & a_{22} \end{array} \right) \mapsto \left(\begin{array}{cc} a_{11} & a_{21} \\ a_{12} & \end{array} \right).$$

We will write $\Omega = \Omega_M$ for convenience. Let $B \in \mathcal{M}(2 \times 2, 1)$ be some fixed matrix with

$$B = \begin{pmatrix} b_{11} & b_{21} \\ b_{12} & b_{22} \end{pmatrix}.$$

Then, the fiber at B is

$$\Omega^{-1}(\Omega(B)) = \left\{ \begin{pmatrix} b_{11} & b_{21} \\ b_{12} & a_{22} \end{pmatrix} \in \mathbb{C}^{2 \times 2} ; b_{11}a_{22} = b_{12}b_{21} \right\}.$$

Note that b_{ij} are now fixed in that set, while a_{22} is the free entry. Now, one of the following two cases has to happen:

Case 1: $b_{11} \neq 0$. Then,

$$\Omega^{-1}(\Omega(B)) = \left\{ \begin{pmatrix} b_{11} & b_{21} \\ b_{12} & a_{22} \end{pmatrix} \in \mathbb{C}^{2 \times 2} ; a_{22} = \frac{b_{12}b_{21}}{b_{11}} \right\} = \{B\}.$$

In this case, $\dim_B \Omega = 0$, and $\#_B \Omega = 1$.

Case 2: $b_{11} = 0$. Then,

$$\begin{aligned} \Omega^{-1}(\Omega(B)) &= \left\{ \begin{pmatrix} 0 & b_{21} \\ b_{12} & a_{22} \end{pmatrix} \in \mathbb{C}^{2 \times 2} ; 0 = b_{12}b_{21} \right\} \\ &= \left\{ B + \begin{pmatrix} 0 & 0 \\ 0 & \lambda \end{pmatrix} ; \lambda \in \mathbb{C} \right\}. \end{aligned}$$

In this case, $\dim_B \Omega = 1$, and $\#_B \Omega = \infty$.

Case 1 is the generic case, as b_{11} is generically non-zero, see the first bullet in Example 2.2.2. Thus, $\dim \Omega = 0$, and $\# \Omega = 1$, so Ω is generically injective, while not being injective (also compare Proposition 2.3.3).

The algebraic theory also implies some results on possible degeneracies which may occur, i.e., if A is not sampled generically (compare the two cases in Example 2.3.8):

Theorem 2.3.9. *let Ω be an $(m \times n)$ -masking in rank r . Let B be any fixed $(m \times n)$ -matrix of rank r or less (not necessarily generic). Then*

(i) $\dim_B \Omega \geq \dim \Omega$.

(ii) If $\dim_B \Omega = 0$, then $\#_B \Omega \leq \# \Omega$.

Proof. The proof is similar to that of Theorem 2.3.9. (i) follows from a more careful application of the upper semicontinuity theorem (see e.g. I.8, Corollary 3 in Mumford [31]), (ii) from purity of the branch locus [30, Chapter 1, Proposition 3.8] and the fact that $\#_B \Omega$ is upper bounded by the degree of the field extension $K(\mathcal{M}(m \times n, r))/K(\Omega(\mathcal{M}(m \times n, r)))$, which is the same as $\# \Omega$. \square

Note that if $\dim \Omega > 0$, then Theorem 2.3.9 (i) implies $\#_B \Omega = \infty$ for arbitrary B . That is, if a generic matrix can not be reconstructed from a given masking, no matrix can be reconstructed. The similar statement that if a generic matrix can be reconstructed, all matrices can be reconstructed, is false, as the proof of Proposition 2.3.3 shows. Furthermore, for a given masking, there can exist a matrix having unique reconstruction only if the masking is generically finite. We want to remark another important consequence of Theorem 2.3.9:

Remark 2.3.10. *The results presented in Theorems 2.3.5 and 2.3.9 are not specific for the case of low-rank matrix completion. They are special instances of classical results from Algebraic Geometry, which are valid for any well-behaved⁶ algebraic map. Thus, they are in similar form valid for any parametric estimation problem which can be decomposed into an exact generative part, given by an algebraic map, and a noise part.*

Theorems 2.3.5 and 2.3.9 allow to replace Question 2.3.2, which we have seen to be uninformative without further specification, with a question that is equal in spirit and excludes a null set of pathological cases.

Question 2.3.11. *Given a fixed mask M , when is M identifiable in rank r ? When is M finitely identifiable in rank r ?*

Theorems 2.3.5 and 2.3.9 show that Question 2.3.11 is well-defined, since the answer depends only on M and the true rank r . Also note that since generic injectivity implies generic finiteness, any condition necessary for generic finiteness is also for injectivity, and any condition sufficient for generic injectivity will be for sufficient generic finiteness. In general, generic injectivity and generic finiteness do not coincide, as the following example shows.

Example 2.3.12. *The mask*

$$M = \begin{pmatrix} 0 & 1 & 1 & 1 \\ 1 & 0 & 1 & 1 \\ 1 & 1 & 0 & 1 \\ 1 & 1 & 1 & 0 \end{pmatrix}.$$

is generically two-to-one in rank 2.

Proof. This can be shown by a symbolic computation, i.e., substituting random - thus generic - values (occurring in a random rank 2 matrix) for the known entries, then computing a degrevlex Gröbner basis for any term order. The initial ideal will be generated by a single quadratic monomial and three linear monomials, showing that there exist two solutions. Since the choice for the substitutions was generic, this implies the statement. \square

It is also important to note that the results of Theorems 2.3.5 and 2.3.9 do not hold (set-theoretically), when working over a field which is not algebraically closed, for example the real numbers \mathbb{R} . Thus in particular, over an algebraically non closed field, Question 2.3.11 is in general not well-defined. Here is the smallest example in which behavior over the complex and reals numbers diverges:

⁶that is, for any proper morphism of Noetherian schemes $X \rightarrow \mathcal{Y}$ over an algebraically closed base field of characteristic zero, and where X is irreducible

Example 2.3.13. Consider the mask M from Example 2.3.12, where it was shown that in rank 2, the masking Ω_M is finitely identifiable, and generically two-to-one. Denote by $\Omega_{\mathbb{R}}$ the restriction of Ω_M to real rank at most 2 matrices. When considering a generic real matrix A of rank 2 as true matrix, then the (set-theoretic) fiber $\Omega_{\mathbb{R}}^{-1}(\Omega_{\mathbb{R}}(A))$ contains, generically, two or no elements. It is not true that it contains generically two elements, nor is it true that it contains generically no element. However, even though A is generic real, the fiber over the complex numbers $\Omega_M^{-1}(\Omega_M(A))$ will generically have two elements.

Example 2.3.13 shows that over the real numbers, there may be several different generic behaviors for the identifiability of Ω , and not only one as in the complex case treated in Theorem 2.3.5. The sets where different types of generic behavior occur are semi-algebraic subsets of $\mathcal{M}(m \times n, r) \cap \mathbb{R}^{m \times n}$. That is, the sets are cut out by polynomial inequalities in the matrix entries, e.g., definiteness (or compare G_1, G_2 in Example 2.3.13). As it is in general hard to distinguish and analyze those semi-algebraic subsets properly, we refrain from doing so for the rest of the paper. However, Theorems 2.3.5 and 2.3.9 give bounds for identifiability; that is, for some masking, being generically finite over the complex numbers is, by Theorem 2.3.9 a necessary condition for any real matrix to have a finite set of possible reconstructions.

Remark 2.3.14. The behavior of the mask in Example 2.3.13 is determined by the behavior of the single degree two polynomial occurring in the Gröbner basis computed in the proof following Example 2.3.12. It is that of the well-known quadratic equation

$$x^2 + bx + c = 0, \quad \text{with } b, c \in \mathbb{R}$$

which has two real solutions if $b^2 > 4c$ and no real solution if $b^2 < 4c$. Both cases are generic in the sense that they have positive Lebesgue measure, as the sets

$$G_1 = \{(b, c); b^2 > 4c\} \quad \text{and} \quad G_2 = \{(b, c); b^2 < 4c\}$$

are not null sets in $\mathbb{R}^2 = \{(b, c); b, c \in \mathbb{R}\}$. The case $b^2 = 4c$ where a single solution occurs is degenerate (and lies in the ramification locus of the parameter map, compare the proof of Theorems 2.3.5 (ii)). Over the complex numbers, the equation has, for generic choice of b, c , always two solutions.

2.3.2 The Combinatorics of Matrix Completion

The generative model low-rank matrix completion is not only algebraic, but has also deep combinatorial features. This was first noticed by Singer and Cucuringu [43], drawing parallels to Rigidity Theory (e.g., Graver et al. [15]). We develop these connections further, studying a generic degree of freedom heuristics.

The combinatorial information in each mask is encoded in a bipartite graph associated to it. We recall the notions of bipartite graph and its adjacency matrix:

Definition 2.3.15. A labeled bipartite graph G is a tuple $G = (V, W, E)$, where $V = [m]$ is the set of red vertices, $W = [n]$ is the set of blue vertices⁷, and $E \subseteq V_1 \times V_2$. The set E is interpreted as the set of edges running between V and W . We will denote the set of edges E of a graph G by $E(G)$, and its cardinality by $e(G) = \#E(G)$.

⁷The symbol $[k]$ will be used to denote the set of natural numbers not larger than k , i.e., $[k] = \{1, 2, \dots, k\}$.

Two bipartite graphs $G_1 = (V_1, W_1, E_1)$ and $G_2 = (V_2, W_2, E_2)$ are isomorphic if there are a pair of bijections $\sigma_V : V_1 \rightarrow V_2$ and $\sigma_W : W_1 \rightarrow W_2$ inducing a bijection $(i, j) \mapsto (\sigma_V(i), \sigma_W(j))$ on the edge sets. The equivalence classes under the induced relation are isomorphism types or (unlabeled) bipartite graphs.

Given a bipartite graph $G = (V, W, E)$, its transpose G^\top is the bipartite graph $G^\top = (W, V, E^\top)$, where $(j, i) \in E^\top$ if and only if $(i, j) \in E$.

The adjacency matrix of a labeled bipartite graph $G = (V, W, E)$, is the $(\#V \times \#W)$ -matrix $M(G)$, where each row corresponds uniquely to an element of V , and each column corresponds uniquely to an element of W . The entry in row i and column j is 1 if and only if $(i, j) \in E$, else it is 0.

Note that in all of these definitions, the bipartition of a (labeled or unlabeled) bipartite graph is fixed. Informally, labelled bipartite graphs are isomorphic if one can be obtained from the other by some relabeling of each vertex class *separately* (i.e., preserving the bipartition). The reason for separating labeled and unlabeled bipartite graphs is that masks correspond to *labeled* bipartite graphs bijectively, via adjacency matrices, but generic completability will turn out to depend only on the *unlabeled* bipartite graph associated with the mask (Proposition 2.5.26).

Definition 2.3.16. Let Ω be a masking with mask $M(\Omega)$. We will call the unique unlabeled bipartite graph $G(\Omega)$ which has adjacency matrix $M(\Omega)$ the adjacency graph of Ω . We will also write $E(\Omega) = E(G(\Omega))$ and $e(\Omega) = e(G(\Omega))$. If we start with the mask $M = M(\Omega)$, we will also denote $G(\Omega)$ by $G(M)$, and similarly replace Ω by M in $E(M) = E(\Omega)$ and $e(M) = e(\Omega)$.


Also, $M(G^\top) = M(G)^\top$, but G and G^\top are in general not isomorphic.

Before continuing, we illustrate the definition of the adjacency graph of a masking by some examples:

Example 2.3.17. Consider the masks from Example 2.1.6 The masks corresponding to the two matrices are

$$M_1 = \begin{pmatrix} 1 & 1 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{pmatrix} \quad \text{and} \quad M_2 = \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 1 \end{pmatrix}.$$

We now interpret M_1 and M_2 as bipartite adjacency graphs. That is, with Definitions 2.3.15 and 2.3.16, both graphs $G(M_k)$, $k = 1, 2$ have three red and three blue vertices: the red vertices are the three rows of M_k , and the blue vertices are the three columns of M_k . An edge is drawn between red vertex/row i and blue vertex/column j if and only if M_k has the entry 1 at position (i, j) . Thus, the graphs G_k shown in figure are the adjacency graphs $G_k = G(M_k)$.



Since a mask M is uniquely represented by its associated graph $G(M)$, Question 2.3.11 can be rephrased into a question on algebraic graph combinatorics:

Question 2.3.18. *Given a bipartite graph G with adjacency matrix M , what are combinatorial conditions on G for M to be identifiable? What are the conditions for M to be finitely identifiable?*

In the following, we give some combinatorial properties on the graph which are sufficient or necessary for generic finiteness. This includes an exact characterization for the case of rank one, which was already derived by Candès and Recht [5] and Singer and Cucuringu [43] using other techniques.

First we give a sufficient condition that relies on a simple algorithmic procedure, which we call r -closure. It corresponds algebraically to subsequently solving minor equations. For ease of definition, we first formally define a replacement operation on graphs, which recursively adds edges to existing sub-patterns.

Definition 2.3.19. *Let H', H be labeled (bipartite) graphs with the same vertex set (where H', H may have isolated vertices). Let M', M be the adjacency matrices of H', H . If $M - M'$ is non-negative (i.e., it is a mask or adjacency matrix), then the induced injection of (unlabeled) graphs $\phi : H' \rightarrow H$ is called a closing.⁸*

The concept of closing will become important in an algorithmic context, where we search for subgraphs H' and add all edges which are missing in H . The map ϕ is needed to describe exactly where the edges are missing, since specification of H' and H does in general not suffice:

Example 2.3.20. *Consider the masks/adjacency matrices*

$$M' = \begin{pmatrix} 1 & 1 & 0 \\ 1 & 0 & 0 \end{pmatrix} \quad \text{and} \quad M = \begin{pmatrix} 1 & 1 & 0 \\ 1 & 1 & 1 \end{pmatrix}.$$

This induces an injection of graphs $H' \rightarrow H$ which is different from the injection one gets when replacing M by

$$M'' = \begin{pmatrix} 1 & 1 & 1 \\ 1 & 1 & 0 \end{pmatrix},$$

since in one case two edges in the same row are added, in the other two edges in two different rows and two different columns are added.

Definition 2.3.21. *Let G be a (bipartite) graph with vertex set V , let $\phi : H' \rightarrow H$ be a closing. We define⁹ a graph $\text{cl}[\phi](G)$, having vertex set V and containing G , together with a closing $\text{cl}[\phi] : G \rightarrow \text{cl}[\phi](G)$, by the following properties:*

- (i) *Every edge e in $\text{cl}[\phi](G)$ is an image of an edge e' in G under $\text{cl}[\phi]$, or there exists a subgraph $F \subseteq G$ isomorphic to H' such that e' connects vertices in the image $\text{cl}[\phi](F)$.*
- (ii) *For each subgraph $F \subseteq G$ isomorphic to H' , there is a closure ρ such that the restriction of $\text{cl}[\phi]$ to F factors as $\rho \circ \phi$.*

⁸An alternative (shorter but more technical) definition of a closing is: a graph monomorphism $\phi : H' \rightarrow H$ with $V(H') = V(H)$ is called closing.

⁹an alternative definition of $\text{cl}[\phi](G)$, is as follows: Take a closing ϕ which corresponds to two $(m' \times n')$ -masks M, M' with $M - M'$ positive. Let G be a bipartite graph with $(m \times n)$ adjacency matrix A' . Let A be the unique $(m \times n)$ -mask with the least number of non-zero entries such that for all row selection matrices $P_m \in \mathbb{C}^{m' \times m}$ and $P_n \in \mathbb{C}^{n' \times n}$ (i.e., P_m, P_n are the first m' resp. n' rows of a permutation matrix), the matrix $P_m A P_n^T - M$ is positive if $P_m A' P_n^T - M'$ is positive. Then, $\text{cl}[\phi](G)$ is defined as the graph $G(A)$.

Let ϕ_1, \dots, ϕ_k be closings. Then we write $\text{cl}^0[\phi_1, \dots, \phi_k](G) = G$, and, by induction,

$$\text{cl}^n[\phi](G) = \text{cl}[\phi_1] \dots \text{cl}[\phi_k] \left(\text{cl}^{n-1}[\phi_1, \dots, \phi_k](G) \right)$$

for all integers $n \geq 1$.

If it is clear from the context, we let the operations cl^n also act on the adjacency matrices instead of the graphs.

Intuitively, the operation $\text{cl}[\phi]$ takes all instances of H' in G , and adds all missing edges in H , in the way the map ϕ prescribes it. $\text{cl}[\phi_1, \dots, \phi_k]$ does the same for several closure patterns ϕ_1, \dots, ϕ_k . Thus, it can be seen that the closing operation cl^n does not add any new edges for big enough n , which makes the following definition well-posed.

Example 2.3.22. Let $H = K_{2,2}$ be the complete bipartite graph with two red and two blue vertices (often called biclique), and let $H' = K_{2,2}^-$ be $K_{2,2}$ minus one edge. The graph H' , which we call the almost biclique, has the same isomorphism type for any choice of the missing edge, so the notation $H' = K_{2,2}^-$ has no ambiguity, and the induced map $\phi : H' \rightarrow H$ that adds the missing edge from H' to H is canonical.

Let G be the bipartite graph with adjacency matrix

$$M = \begin{pmatrix} 0 & 1 & 1 \\ 1 & 1 & 0 \\ 1 & 0 & 0 \end{pmatrix}.$$

Then,

$$\text{cl}[\phi](M) = \begin{pmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 0 \end{pmatrix} \quad \text{and} \quad \text{cl}^2[\phi](M) = \begin{pmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{pmatrix}.$$

Definition 2.3.23. Let ϕ_1, \dots, ϕ_k be closings. Let $N \in \mathbb{N}$ be any integer such that

$$\text{cl}^N[\phi_1, \dots, \phi_k](G) = \text{cl}^{N-1}[\phi_1, \dots, \phi_k](G).$$

Then, we write $\text{cl}^\infty[\phi_1, \dots, \phi_k](G) = \text{cl}^N[\phi_1, \dots, \phi_k](G)$. The graph $\text{cl}^\infty[\phi_1, \dots, \phi_k](G)$ is called $[\phi_1, \dots, \phi_k]$ -closure of G . If $k = 1$, we also write ϕ -closure instead of $[\phi]$ -closure.

For matrix completion, the most obvious closure operation is of same type as in Example 2.3.22, i.e., adding missing edges to almost complete bicliques:

Definition 2.3.24. Denote by $K_{r+1, r+1}^-$ the complete bipartite graph $K_{r+1, r+1}$ minus one edge. Let $\phi : K_{r+1, r+1}^- \rightarrow K_{r+1, r+1}$ be the canonical inclusion map.

Let G be a bipartite graph with m red and n blue vertices. Instead of ϕ -closure of G , we will also say r -closure of G . If $\text{cl}^\infty[\phi](G) = K_{m, n}$, then we call G an r -closable (bipartite) graph.

Intuitively, the r -closure of a graph G is obtained by repeatedly adding single missing edges which complete a subgraph G to the complete subgraph $K_{r+1, r+1}$. It generalizes transitive closure, as the following lemma shows:

Lemma 2.3.25. The 1-closure of a bipartite graph G is the transitive bipartite closure of G . A bipartite graph is 1-closable if and only if it is connected.

Proof. It suffices to prove equivalence of 1-closure and transitive bipartite closure, as the second statement follows from that.

First, we show that the 1-closure is contained in the transitive closure. This is equivalent to showing that every edge contained in the 1-closure is contained in the transitive closure. But that follows from the fact that any edge added in the closure process connects vertices in the same connected component, since $K_{2,2}^-$ is connected, and the new edge is added between two vertices of an already existing $K_{r+1,r+1}^-$.

Now we show that transitive closure is contained in the 1-closure; i.e., any edge contained in the transitive closure is contained in the 1-closure. As closure is defined by adding edges at positions given by subgraphs, it suffices to prove that for bipartite trees, by choosing a spanning forest of G .

We show this by induction on the number of edges of the tree. The statement is true for edge numbers 1 and 2 since the trees are bipartite, thus transitive closure will add no edges. We now complete the induction by assuming validity for trees with n edges and proving it for those with $n + 1$ edges.

Let T be a tree with $n + 1$ edges. Thus, we can decompose $T = T' \cup G$, where T' is a tree with n edges and G is a single-edge graph. By induction assumption, the 1-closure of T' is a complete bipartite graph K , thus the 1-closure of T is equal to the 1-closure of $K \cup G$. Let v be the unique vertex in G not in K , let w be the vertex in $G \cup K$. By construction, every vertex of different color in $K \cup G$ is connected to w in K . Moreover, the distance of v to any differently colored vertex in $G \cup K$ is at most 3, since K is complete. Thus, the 1-closure of $G \cup K$ will contain the edges of v to any other edge, implying that the 1-closure of $G \cup K$ is complete.

This completes the induction, thus the proof of the statement. \square

The graph theoretical concept of r -closure now allows to formulate a sufficient graph theoretical condition for finite completability, which was already obtained in [26]. We reproduce the Proposition and the proof here.

Proposition 2.3.26. *A masking Ω is generically injective in rank r if $G(\Omega)$ is r -closable.*

Proof. If $G(\Omega)$ contains a subgraph isomorphic to $K_{r+1,r+1}^-$, this means that for a generic matrix A , some $(r+1 \times r+1)$ -sub-matrix A' is known in $\Omega(A)$, except for one matrix element - corresponding to the missing edge in $K_{r+1,r+1}^-$. Since A has rank r , the determinant of A' vanishes. As all entries but one are known, the vanishing minor condition gives a linear equation on the missing entry. The linear equation is non-trivial, i.e., not the trivial equation $0 = 0$, since due to the genericity of A , the linear and constant coefficients are all non-zero. Thus the linear equation allows to uniquely reconstruct A' and thus uniquely determine an unknown entry of A . Now r -closability translates to the fact that this process can be repeated until the whole of A is uniquely reconstructed. As we have assumed that A is generic, this implies generic injectivity for Ω . \square

We want to mention that r -closability of the associated graph is neither a necessary condition for generic injectivity, nor for generic finiteness. The mask from Example 2.3.13 proves the latter, as it is not 2-closable. We will now prove by example that r -closability is not necessary for generic injectivity:

Example 2.3.27. *The mask*

$$M = \begin{pmatrix} 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 0 & 1 & 1 \\ 1 & 1 & 0 & 1 & 1 \\ 0 & 1 & 1 & 1 & 1 \\ 0 & 1 & 1 & 1 & 1 \end{pmatrix}$$

is uniquely completable in rank 3, but not 3-closable.

However, note: For $m, n \leq 3$, generic injectivity and r -closability coincide. For $m, n \leq 5$, generic finiteness and r -closability coincide.

Proof. We can see that M is uniquely completable in rank 3 by first observing that the two missing entries in the middle appear in two 4×4 vanishing minor equations, one in the top left corner (rows 1 to 4 and columns 1 to 4) and another in the bottom right corner (rows 3 to 6 and columns 2 to 5); note that both of these equations are linear with respect to the two missing entries, because they appear in the same column. Therefore, the two missing entries lie in the solution of a system of two linear equations, which are linearly independent since the top left entry and the bottom right - which occur in the coefficients of the equations - are algebraically independent. Thus, the two missing entries can be uniquely determined. After filling these in, the mask becomes 3-closable and therefore unique. On the other hand, M is not 3-closable because for each missing entry (i, j) , the bipartite subgraph $(N(j), N(i), E')$, where $N(j)$ is the set of neighbors of j and $E' = E(M) \cap (N(j) \times N(i))$, does not contain a 3×3 biclique. Due to symmetry, we need to check only 5 bipartite subgraphs to see this.

The statements on finiteness and injectivity follow from an exhaustive search using the algorithms in Section 3 □

In Section 2.5.3, we will see, that a necessary condition on generic finiteness can be formulated in terms of some closure, which in general is not equivalent to some r -closure.

Now, we prove some necessary conditions for generic finiteness of a masking. First, recall the definition of r -connectedness:

Definition 2.3.28. *We say a bipartite graph $G = (V, W, E)$ is r -edge-connected, or r -connected, if for any non-trivial vertex partition¹⁰ of G into two graphs G_1, G_2 , the set the number of edges running between G_1 and G_2 is lower bounded by r , i.e.,*

$$e(G) - e(G_1) - e(G_2) \geq r.$$

That means, G stays connected after removing an arbitrary set of $r - 1$ edges.

We also define an abbreviation for the number of degrees of freedom of $(m \times n)$ -matrices of rank r , compare Remark 2.3.30:

Definition 2.3.29. *For $m, n \in \mathbb{N}$, we will write*

$$d_r(m, n) = mn - \max(0, m - r) \max(0, n - r).$$

If G is a graph with m red and n blue vertices, we will also write $d_r(G) = d_r(m, n)$.

¹⁰one has $V = V_1 \cup V_2, V_1 \cap V_2 = \emptyset$ and $W = W_1 \cup W_2, W_1 \cap W_2 = \emptyset$ and $E(G_1) \cup E(G_2) \subseteq E(G)$; non-trivial means that each of G_1, G_2 contains at least one vertex

Note that $d_r(m, n) = mn$ if $m \leq r$ or $n \leq r$, else $d_r(m, n) = r \cdot (m + n - r)$.

Remark 2.3.30. The number $d_r(m, n)$ is the dimension of the determinantal variety $\mathcal{M}(m \times n, r)$, which is classically known. Intuitively, this is the number of (algebraic) degrees of freedom in the set of $(m \times n)$ -matrices of rank r (or less). If $n \leq r$ or $m \leq r$ holds, then $\mathcal{M}(m \times n, r) \cong \mathbb{C}^{mn}$, and it directly follows that $d_r(m, n) = \dim \mathcal{M}(m \times n, r)$ in that case. In all other cases, $d_r(m, n) = r \cdot (m + n - r)$. There are several way to prove that this number is the same as $\dim \mathcal{M}(m \times n, r)$, we want to give two heuristic arguments (not proofs) why this should be the case.

First, consider a rank r matrix A of size $m \times n$. We can choose A by first choosing the column span, which is a r -dimensional vector subspace V of n -space. Classically, this choice is known to have $r(n - r)$ degrees of freedom, and is parameterized by the Grassmannian $\text{Gr}(r, n)$ (the latter is also an algebraic variety, and its dimension is $r(n - r)$). Then, one can choose each column of A from V ; since V is r -dimensional, this is only r degrees of freedom for each column, so in total

$$r(n - r) + mr = r \cdot (m + n - r)$$

degrees of freedom.

Alternatively, one can write $A = UV^\top$ with U a $(m \times r)$, and V a $(n \times r)$ -matrix. There are a total of $r \cdot (m + n)$ entries in both U and V , but one can replace a particular choice of U, V by $U \cdot B$ and $V \cdot (B^{-1})^\top$, where B is any full rank $(r \times r)$ -matrix. There are r^2 degrees of freedom to choose B . Since the choice of B are degrees of freedom which do not appear in the choice of A ; one has to subtract them from those in the choice of U, V , giving a total of

$$r \cdot (m + n) - r^2 = r \cdot (m + n - r)$$

degrees of freedom.

Note that both arguments do not constitute proofs, as it has to be shown that the degrees of freedom added together are not redundant in the first argument, and that there are no other degrees of freedom which do not appear in A that could be subtracted. Both arguments give $r \cdot (m + n - r)$ as an upper bound on the degrees of freedom though.

Now we state some necessary conditions on generic finiteness:

Proposition 2.3.31. *If a masking Ω is generically finite in rank r , then:*

- (i) $e(\Omega) \geq d_r(G(\Omega))$
- (ii) Each vertex of $G(\Omega)$ has degree at least r
- (iii) $G(\Omega)$ is r -connected

Proof. (i) By the dimension formula, it holds that

$$\dim \Omega = \dim \mathcal{M}(m \times n, r) - \dim \Omega(\mathcal{M}(m \times n, r)).$$

By definition, Ω is generically finite if and only if $\dim \Omega = 0$, thus Ω is generically finite if and only if

$$\dim \mathcal{M}(m \times n, r) = \dim \Omega(\mathcal{M}(m \times n, r)).$$

Again, by definition, one has $\dim \Omega(\mathcal{M}(m \times n, r)) \leq e(\Omega)$. Thus, if Ω is generically finite, then

$$d_r(G(\Omega)) = \dim \mathcal{M}(m \times n, r) = \dim \Omega(\mathcal{M}(m \times n, r)) \leq e(\Omega).$$

(iii) implies (ii) in the special case of the partition in the graph in one vertex and the rest. We will show that the statement (iii) follows from the more general Proposition 2.3.34. Since Proposition 2.3.34 is proved later, note that there are no loops in the proof structure. So assume that Ω is generically finite, Proposition 2.3.34 (iii) below then shows that for any vertex partition of $G(\Omega)$ into two graphs G_1, G_2 , it holds that

$$e(G) - e(G_1) - e(G_2) \geq d_r(G) - d_r(G_1) - d_r(G_2).$$

An elementary calculation shows that the right hand side is always r or greater if the G_i are non-trivial, thus

$$e(G) - e(G_1) - e(G_2) \geq r$$

for any vertex partition of $G(\Omega)$ into G_1, G_2 , which is the definition of r -connectedness. □

The following example proves that the conditions given in Proposition 2.3.31 are not sufficient:

Example 2.3.32. *Consider the mask*

$$M = \begin{pmatrix} 0 & 0 & 1 & 1 & 1 \\ 0 & 1 & 1 & 1 & 1 \\ 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 \end{pmatrix}.$$

M is not finitely completable in rank 2, but 2-connected. In particular, each vertex in $G(M)$ has degree at least 2. This is equivalent to the fact that each row and each column of M has at least 2 non-zero entries. Also, $e(\Omega) = 14 \geq 14 = r \cdot (m + n - r)$. Thus, no single of the conditions in Proposition 2.3.31 is sufficient for finite completability, nor is their conjunction.

Example 2.3.32 also shows that r -connectedness is too weak to describe finite completability. Namely, if the graphs in a vertex partition, as in the Definition 2.3.28 of r -connectedness, are similarly large, the number of edges running between them has to be bigger than r ; also, the more balanced the partition is, the more edges have to run between the partition components.

We now introduce a concept of rank-related sparsity, which in its dual notion, will be equivalently reflecting that fact. Singer and Cucuringu [43] have already conjectured that some sparsity concept might play a role in describing the completable masks.

Definition 2.3.33. *A bipartite graph G is called rank- r -sparse, if for all subgraphs $G' \subseteq G$ it holds that $e(G') \leq d_r(G')$.*

If, additionally, G has exactly $d_r(G)$ edges, the graph G is called rank- r -tight.

We say G is spanned by a rank- r -tight graph if G contains a rank- r -tight graph with m red and n blue vertices. Abbreviatingly, we also say that G is spanned in rank r .

Rank r -sparsity is closely related to combinatorial properties defined using partitions of the vertices and edges (cf. the Nash-Williams-Tutte Arboricity Theorem [32, 49]):

Proposition 2.3.34. *Let G be a bipartite graph with least $d_r(G)$ edges. Consider the following statements:*

(i) G is spanned in rank r .

(ii) For all partitions of the edges¹¹ of G into subgraphs G_1, \dots, G_N ,

$$d_r(G) \leq \sum_{i=1}^N d_r(G_i). \quad (1)$$

(iii) For all partitions of the vertices inducing subgraphs¹² of G into graphs G_1, \dots, G_N ,

$$e(G) - d_r(G) \geq \sum_{i=1}^N (e(G_i) - d_r(G_i)). \quad (2)$$

Then, the implications (i) \Rightarrow (ii) \Rightarrow (iii) hold, and each of the three conditions (i), (ii), (iii) is necessary for Ω to be generically finite.

Proof. That (i) is necessary for Ω to be generically finite follows from Theorem 2.5.31; since Theorem 2.5.31 will be proved later, it is important to note that there are no loops in the proof structure. Necessity of (ii) and (iii) follow directly once the implications are proved.

(i) \Rightarrow (ii): We show that a graph G which violates the inequality in (ii) cannot be spanned in rank r . Let G_1, \dots, G_N some edge partition such that

$$d_r(G) > \sum_{i=1}^N d_r(G_i).$$

Let H be any rank- r -sparse graph contained in G . Denote by h the number of edges of H , and by h_i the number of edges of H , lying in G_i . By definition, one has $h_i \leq d_r(G_i)$, and $h = h_1 + \dots + h_N$. Inserting into the inequality above gives

$$d_r(H) = d_r(G) > \sum_{i=1}^N d_r(G_i) \geq \sum_{i=1}^N h_i = h,$$

which shows, by contradiction, that H is not rank- r -tight.

(ii) \Rightarrow (iii): Let G_1, \dots, G_N be a vertex partition of G , as in (iii). Let H_1, \dots, H_M be single-edge graphs for all the edges not contained in any of the G_i . Thus, $G_1, \dots, G_N, H_1, \dots, H_M$ is an edge partition of G . Then, by (ii), one has

$$d_r(G) \leq M + \sum_{i=1}^N d_r(G_i).$$

The condition in (iii) follows from

$$M = e(G) - \sum_{i=1}^N e(G_i)$$

and elementary arithmetic. □

¹¹i.e., if $G = (V, W, E)$ and $G_i = (V_i, W_i, E_i)$, then $E = E_1 \cup E_2 \cup \dots \cup E_N$, and $E_i \cap E_j = \emptyset$ for all i, j

¹²i.e., if $G = (V, W, E)$ and $G_i = (V_i, W_i, E_i)$, then $V = V_1 \cup V_2 \cup \dots \cup V_N$, and $V_i \cap V_j = \emptyset$ for all i, j ; same for W

Remark 2.3.35. *The conditions (i), (ii), and (iii) can be derived from various heuristics for finite completability of a mask:*

- (i) *To be minimally completable, a mask needs $r \cdot (m + n - r)$ total edges by a dimension count, and no subgraph should be “overloaded”.*
- (ii) *For any partition, the sum on the r.h.s. of (1) is an upper bound for the maximum size of a rank r -sparse subgraph.*
- (iii) *In a completable mask, if we replace the pieces of any partition of the rows and columns with completely observed $r \times r$ sub-matrices, the associated contracted graph is also completable. The inequality of (2) asserts that there are enough edges for this property to hold.*

Condition (iii) was also proved directly in Király and Tomioka [25]; the proof path presented here can be specialized to that one.

After we develop the machinery of determinantal matroids, we will be able to show that (i) is indeed necessary for generic finite completability of a masking (Theorem 2.5.31), implying the same thing for (ii) and (iii).

2.4 Partial Completion

Matrix completion, as defined so far, asks whether a low-rank matrix can be completely reconstructed from a set of its entries. However, in practical scenarios, e.g., recommendation, or prediction, it is more common that one is only interested in reconstructing some of the missing entries, not all. Most approaches overlook this fact as they rely on reconstructing the complete matrix first.

We will call this task partial low-rank matrix completion, or just Partial Completion. Most important observations made in section 2.3 still hold for partial completion, for example the existence of a zero-measure set of exceptions or that reconstructability only depends on the pattern of observed entries. To prove this, we will use the tools from Algebraic Combinatorics from section 2.3. One can also decompose the generative sampling model into an algebraic part, and a noise part - we leave that to the reader, as it is very similar to what is presented in Section 2.2.

First, we want to formally state the problem of Partial Completion. As the problem generalizes matrix completion, we already start with a refined formulation that includes generic sampling:

Question 2.4.1. *Given an $(m \times n)$ -mask M of rank r , and a generic matrix A : Which entries of A can be reconstructed from the masked matrix $\Omega_M(A)$?*

In order to get a formal grasp on Question 2.4.1, we define the analogues of masking for Partial Completion:

Definition 2.4.2. *Let N, M be $(m \times n)$ -masks such that $N - M$ is a mask. Let Ω_M, Ω_N be the corresponding maskings in rank r . The unique map $\Omega_{N/M} : \Omega_N(\mathcal{M}(m \times n, r)) \rightarrow \mathbb{C}^a$ such that*

$$\Omega_M = \Omega_{N/M} \circ \Omega_N$$

is called a partial masking in rank r . The pair (N, M) is called partial mask and denoted by $M(\Omega_{N/M})$. If $N - M$ contains exactly one non-zero entry, in the i -th row and j -th column, we will also write $((ij), M)$ for the partial mask. The inclusion map of graphs $(G \hookrightarrow H)$ such that H has adjacency matrix N and G has adjacency matrix M is called graph map of Ω and denoted by $G(\Omega)$.

We also will look at one special fiber and define what fiber dimension and cardinality are in this case:

Definition 2.4.3. Let A be an $(m \times n)$ -matrix of rank r , let (N, M) be a partial $(m \times n)$ -mask. Let Ω_N, Ω_M be the corresponding maskings in rank r , let $\Omega_{N/M}$ be the partial masking in rank r . We will call the integer

$$\dim_A \Omega_{N/M} = \dim \Omega_{N/M}^{-1}(\Omega_M(A))$$

the fiber dimension (of $\Omega_{N/M}$) at A . Similarly, we will call

$$\#_A \Omega_{N/M} = \#\Omega_{N/M}^{-1}(\Omega_M(A))$$

the fiber cardinality (of Ω) at A .

The analogue of Theorem 2.3.5 for the Partial Completion setting is:

Theorem 2.4.4. Let A be a generic $(m \times n)$ -matrix of rank r , let Ω be a partial masking in rank r . Then, the following depend only on the true rank r and the mask $M(\Omega)$:

- (i) The fiber dimension $\dim_A \Omega$.
- (ii) The fiber cardinality $\#_A \Omega$.
- (iii) Whether $\#_A \Omega = 1$, i.e., whether the entries of A masked by Ω are uniquely completable.

Proof. The proof is completely analogous to that of Theorem 2.3.5. The only additional thing to note is that the range of $\Omega_{N/M}$, which is $\Omega_N(\mathcal{M}(m \times n, r))$, is irreducible. But that is true since $\Omega_N(\mathcal{M}(m \times n, r))$ is a projection of an irreducible variety. \square

The following are the generalized definitions concerning identifiability and generic behavior to the Partial Completion setting:

Definition 2.4.5. Let Ω be a partial $(m \times n)$ -masking in rank r , let A be a generic $(m \times n)$ -matrix. We will write $\#\Omega$ for the (generic) value of $\#_A \Omega$, and $\dim \Omega$ for the (generic) value of $\dim_A \Omega$.

Again, the generic values of dimension and cardinality bound all possible values:

Theorem 2.4.6. Let Ω be an $(m \times n)$ -masking in rank r . Let B be any fixed $(m \times n)$ -matrix of rank r or less (not necessarily generic). Then

- (i) $\dim_B \Omega \geq \dim \Omega$.
- (ii) If $\dim_B \Omega = 0$, then $\#_B \Omega \leq \#\Omega$.

Proof. The proof is completely analogous to that of Theorem 2.3.9. \square

Now we introduce the analogues characterizing the generic behavior for Partial Completion:

Definition 2.4.7. Let Ω be a partial $(m \times n)$ -masking in rank r . We call

- (i) Ω generically injective and $M(\Omega)$ identifiable or completable (in rank r), if $\#\Omega = 1$.
- (ii) Ω (generically) finite and $M(\Omega)$ finitely identifiable or finitely completable (in rank r), if $\#\Omega < \infty$.

(iii) Ω generically k -to-one (in rank r), if $\#\Omega = k$.

(iv) Ω infinite and $M(\Omega)$ unidentifiable (in rank r), if $\dim \Omega > 0$.

Remark 2.4.8. Let Ω be a partial masking with partial mask (N, M) . Whether Ω is generically injective, finite, etc. can be checked separately for single entries of $N - M$. That is, write

$$N - M = \sum_{i=1}^n (N_i - M) \quad \text{with masks } N_i$$

such that $\|N_i - M\|_1 = 1$, i.e., $N_i - M$ contains only one non-zero entry. Then it can be shown:

- (i) $M(\Omega)$ is identifiable if and only if (N_i, M) is identifiable for all $1 \leq i \leq n$.
- (ii) $M(\Omega)$ is finitely identifiable if and only if (N_i, M) is finitely identifiable for all $1 \leq i \leq n$.
- (iii) $M(\Omega)$ is unidentifiable if and only if there exists an i such that (N_i, M) is unidentifiable.
- (iv) If Ω is generically k -to-one, and the prime factorization of k is $k = p_1 \cdot p_2 \cdot \dots \cdot p_\ell$, then for all j there exists an i such that (N_i, M) is k_i -to-one and p_j divides k_i .
- (v) If $\Omega_{N_i/M}$ is generically k_i -to-one, then Ω is at least $\text{lcm}(k_i)$ -to-one, and at most $(\prod_{i=1}^n k_i)$ -to-one.

Also, note that it may happen that some masks (N_i, M) are identifiable, while some other masks (N_i, M) are finitely identifiable but not identifiable.

The statements can be proved by applying Galois theory to the fact that $\dim \Omega_{N/M}$ is the same as the transcendence degree of the field extension

$$\mathbb{C}(\Omega_N(\mathcal{M}(m \times n, r))) / \Omega_M(\mathcal{M}(m \times n, r)).$$

To a partial mask, one can also associate a graph structure, namely, an injective graph morphism:

Definition 2.4.9. Let Ω be a partial masking with partial mask $M(\Omega) = (N, M)$. We will call the unique injective map of bipartite graph $G(\Omega) = (G \hookrightarrow H)$, where G is the adjacency graph of M and H is the adjacency graph of N and the injection identifies vertices, the adjacency map of (N, M) .

Similarly, the adjacency map of a partial masking uniquely characterizes its completion properties; however, for a more thorough discussion, the matroidal structure of matrix completion is needed which will be developed in the following sections.

2.5 Differential Analysis of Completion Patterns

This section is devoted to the analysis of the degrees of freedom in the matrix entries and how they interact. In particular, we want to develop tools which allow us to see which entries of a mask can be chosen independently, omitted without loss of information, or reconstructed from the known ones.

2.5.1 Degrees of Freedom

For a generically sampled true matrix, a natural question is which degrees of freedom remain after observing a specific set of entries. That is, how many degrees of freedom get killed by the masking operation, and how can we combinatorially or algebraically quantify and qualify them. Formally phrased, the question we want to answer in this and the following sections is:

Question 2.5.1. *Let Ω be a masking in rank r . How many degrees of freedom are in its image $\Omega(\mathcal{M}(m \times n, r))$, depending on the mask $M(\Omega)$?*

Finite completability of a mask can be then rephrased as the fact that the image and the range of Ω do have the same numbers of degrees of freedom, namely $d_r(G(\Omega)) = r \cdot (m + n - r)$. As we have already seen in the previous section, Algebraic Geometry provides tools to bound degrees of freedom - the formal concept for that is the (Krull) dimension - and in particular, the number of the degrees of freedom which are lost by applying Ω (to a generic matrix) is exactly the generic fiber dimension $\dim \Omega$. In this and the following section, we will go a step further and develop tools which in the end will allow to combinatorially study and algorithmically determine the exact number of degrees of freedom, in terms of the structure of the map Ω , namely the mask $M(\Omega)$, and the true rank r . Singer and Cucuringu [43] have already proposed a probabilistic algorithm for checking finite completability of a mask based on differentials, without giving a proof for its correctness. The results of this and the following sections will allow to later fill that gap, and provide more general algorithms for the mentioned purposes.

The main ingredient in the following is a refinement of a classical instrument from calculus, the Jacobian matrix. The Jacobian and its generalizations are also classic tools in Algebraic Geometry to describe fiber dimension of an algebraic map. We will now describe in which specific instance it occurs in the context of matrix completion.

The masking Ω is - as it was defined in Definition 2.1.4 - a map

$$\Omega : \mathcal{M}(m \times n, r) \rightarrow \mathbb{C}^\alpha.$$

If $A \in \mathcal{M}(m \times n, r)$, then there exist matrices $U \in \mathbb{C}^{m \times r}$ and $V \in \mathbb{C}^{n \times r}$ such that $A = U \cdot V^\top$. Conversely, given matrices $U \in \mathbb{C}^{m \times r}$ and $V \in \mathbb{C}^{n \times r}$, the matrix $A = U \cdot V^\top$ has rank at most r . This means, reformulated, that the set $\mathcal{M}(m \times n, r)$ can be parameterized (non-uniquely) by $\mathbb{C}^{(m+n) \times r}$, via the surjective map

$$\begin{aligned} \mu : \mathbb{C}^{(m+n) \times r} = \mathbb{C}^{m \times r} \times \mathbb{C}^{n \times r} &\rightarrow \mathcal{M}(m \times n, r) \\ (U, V) &\rightarrow U \cdot V^\top. \end{aligned}$$

So the composition of maps $\Omega \circ \mu$ is a map from complex $r \times (m + n)$ -space into complex α -space, and its fiber dimension can be computed by the Jacobian matrix. That is, write

$$\Omega \circ \mu = (f_1, \dots, f_\alpha)$$

with algebraic maps f_i . Consider the Jacobian matrix

$$J(U, V) = \begin{pmatrix} \frac{\partial f_1}{\partial U_{11}}(U, V) & \dots & \frac{\partial f_1}{\partial U_{mr}}(U, V) & \frac{\partial f_1}{\partial V_{11}}(U, V) & \dots & \frac{\partial f_1}{\partial V_{nr}}(U, V) \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ \frac{\partial f_\alpha}{\partial U_{11}}(U, V) & \dots & \frac{\partial f_\alpha}{\partial U_{mr}}(U, V) & \frac{\partial f_\alpha}{\partial V_{11}}(U, V) & \dots & \frac{\partial f_\alpha}{\partial V_{nr}}(U, V) \end{pmatrix},$$

where the derivatives have to be taken over all entries of the matrices U, V , i.e., all U_{ij} with $1 \leq i \leq m$ and $1 \leq j \leq r$ and V_{ij} with $1 \leq i \leq n$ and $1 \leq j \leq r$. It is possible to show with tools from Algebraic Geometry¹³ that for $A = U \cdot V^\top$, the fiber dimension $\dim_A \Omega = \dim \Omega^{-1}(\Omega(A))$, at some fixed matrix $A = U \cdot V^\top$, equals $d_r(m, n) - \text{rk}J(U, V)$. Thus, the matrix A can be reconstructed, up to finite choice, from $\Omega(A)$ if and only if $J(U, V)$ has rank $d_r(G(\Omega)) = r \cdot (m + n - r)$. This means that finite completability is determined by the rank of the Jacobian at generic U, V or at a generic A of rank r .

Arguments of this type can be refined to yield degrees-of-freedom-statments on any set of entries of the matrix A . Namely, to each entry of A , one can associate one row of the Jacobian; if one has more than one entry of A , one can calculate the degrees of freedom lying in those by the dimension of the span of the corresponding row vectors. The following subsections will be devoted to giving an exposition on the proper mathematical tools from Combinatorial Commutative Algebra to do so, and how to apply them.

2.5.2 Differentials

In this section, we will develop the basic theory of derivatives and differentials, which is a tool from commutative algebra that allows us to quantify and qualify degrees of freedom and dependencies between objects.

The basic idea which characterizes dependence of algebraic quantities, e.g., the entries of a low-rank matrix, is that of algebraic dependence. In the end, we will see that it is also the right concept to count the degrees of freedom, as it exposes analogies to linear dependence in Linear Algebra.

Definition 2.5.2. *Let K be a field over \mathbb{C} , (e.g. the field of rational functions¹⁴ $\mathbb{C}(T_1, \dots, T_m)$) let $\alpha_1, \dots, \alpha_n \in K$. Then, $\alpha_1, \dots, \alpha_n$ are called algebraically dependent (over \mathbb{C}) if there is a non-zero polynomial¹⁵ $f \in \mathbb{C}[X_1, \dots, X_n]$ such that*

$$f(\alpha_1, \dots, \alpha_n) = 0.$$

Else we call the α_i algebraically independent (over \mathbb{C}).

Intuitively, this means that if $\alpha_1, \dots, \alpha_n$ are algebraically dependent, then knowing some $n - 1$ of the α_i implies knowing that the remaining α_j must be one of finitely many values. Alternatively, one can think of an algebraically independent set of n elements carrying one degree of freedom each, in total n , while an algebraically dependent set of elements has redundancies and strictly less than n degrees of freedom¹⁶. We will explain this by the next example and give a more precise statement in Proposition 2.5.4.

Example 2.5.3. *Consider the (formal) polynomials $\alpha_1 = X^2, \alpha_2 = XY, \alpha_3 = Y^2$. The three α_i are algebraically dependent, since for*

$$f(X_1, X_2, X_3) = X_1 X_3 - X_2^2,$$

¹³One has $\dim_{(U,V)}(\Omega \circ \mu) = \dim_{(U,V)} \mu + \dim_A \Omega$, and $\dim_{(U,V)} \mu = r^2$, since one can show that the representation $U \cdot V^\top$ is unique up to multiplication $U = UB$ and $V = V(B^{-1})^\top$ with an invertible $(r \times r)$ -matrix B . On the other hand, one has $\dim_{(U,V)}(\Omega \circ \mu) = r(m + n) - \text{rk}J(U, V)$ by the Jacobian criterion.

¹⁴ $\mathbb{C}(T_1, \dots, T_m)$ is the set of all formal fractional functions f/g , where f and g are in $\mathbb{C}[T_1, \dots, T_k]$, see next footnote

¹⁵ $\mathbb{C}[X_1, \dots, X_n]$ is the set of polynomials in the n variables X_1, \dots, X_n and with coefficients in \mathbb{C}

¹⁶In Algebra, this is formalized by the transcendence degree of the field extension

one calculates that $f(\alpha_1, \alpha_2, \alpha_3) = 0$. On the other hand, there cannot be a non-zero polynomial $g(X_1, X_2)$ which evaluates to zero when substituting any two of the α_i , as there is always one of the two α_i which contains a variable (i.e., X or Y) which the other does not.

Now assume there is some truth (X, Y) and we measure some of the α_i . For a generic truth (X, Y) , knowing any two of the α_i will allow us to predict the third via $f(\alpha_1, \alpha_2, \alpha_3) = 0$, up to a finite choice. For example, knowing α_1 and α_2 , one can recover $\alpha_3 = \frac{\alpha_2^2}{\alpha_1}$ exactly. On the other hand, when one knows α_1 and α_3 , the recovery of $\alpha_2 = \pm\sqrt{\alpha_1\alpha_3}$ is only possible up to sign, i.e., one has two choices here. Moreover, knowing only one of the α_i allows no prediction whatsoever on the other α_j , since one degree of freedom remains, either in choosing (X, Y) or any second of the α_j .

The behavior in Exercise 2.5.3 occurs in fact in all similar situations:

Proposition 2.5.4. *Let $K = \mathbb{C}(X_1, \dots, X_n)$, let $\alpha_1, \dots, \alpha_k \in K$. Then, $\alpha_1, \dots, \alpha_k$ are algebraically dependent over \mathbb{C} if and only if for generic $x \in \mathbb{C}^n$ and possible reordering of the indices of the α , the values $\alpha_1(x), \dots, \alpha_{n-1}(x)$ determine the value $\alpha_n(x)$ up to finite choice.*

Proof. We prove both directions by a series of equivalences. The fact that $\alpha_1, \dots, \alpha_k$ are algebraically dependent over \mathbb{C} , by definition, is equivalent to the fact that there exists a non-zero complex polynomial

$$P : \mathbb{C}^k \rightarrow \mathbb{C}$$

such that $P(\alpha_1, \dots, \alpha_k) = 0$. Since P is non-zero, we can reorder the indices of the α such that the polynomial

$$P(\alpha_1, \dots, \alpha_{n-1}, T)$$

is non-trivially dependent on the variable T . So, equivalently, for generic x ,

$$P(\alpha_1(x), \dots, \alpha_{n-1}(x), T) = 0$$

where P is not a constant polynomial in T . This is equivalent, for $T = \alpha_n(x)$, that $\alpha_n(x)$ is determined up to finite choice for generic x . \square

Proposition 2.5.4 shows that algebraic dependency is the proper concept to treat degrees of freedom in the matrix completion setting; however, as it can be seen from Example 2.5.3 or more complicated examples as matrix completion itself, it is not always straightforward how to determine the existence of algebraic dependencies, or how to prove their non-existence, when given the measurement polynomials.

The central idea which makes the latter theoretically and also algorithmically feasible is the differential study of the polynomials, as already explained in section 2.5.1. Namely, the existence of dependencies and their degrees of freedom can be studied by the formal, or by the evaluated derivatives of the polynomials.

For a more concise description, we need to introduce the concept of formal differentials and their evaluations first. Here, we adopt an ad-hoc definition of differentials; more natural definitions and further results can be found in any introductory book on Commutative Algebra.

Definition 2.5.5. *Let K be a field over \mathbb{C} (i.e., $\mathbb{C} \subseteq K$), with multiplication \cdot_K . The set of formal differentials of K over \mathbb{C} is the set*

$$\Omega^1(K/\mathbb{C}) = \{f \cdot dg ; f, g \in K\}/\{\sim\},$$

where \sim is the equivalence relation given by

- (i) $d\alpha = 0$ for all $\alpha \in \mathbb{C}$
- (ii) $d(\alpha f) = \alpha df$ for all $\alpha \in \mathbb{C}$ and $f \in K$
- (iii) $d(f + g) = df + dg$ for all $f, g \in K$
- (iv) $d(f \cdot_K g) = g \cdot df + f \cdot dg$ for all $f, g \in K$

where we write $0 = d0$ and $df = 1 \cdot df$ for all $f \in K$.

A multiplication $\cdot : \Omega_{K/\mathbb{C}}^1 \times K \rightarrow \Omega_{K/\mathbb{C}}^1$ is defined as

$$(f, g \cdot dh) \mapsto (f \cdot_K g) \cdot dh,$$

making $\Omega_{K/\mathbb{C}}^1$ a vector space over \mathbb{C} . When clear from the context, we will omit \cdot and/or \cdot_K .

Intuitively, the equivalence relation \sim imposes all usual differentiation rules which are commonly known, e.g., for rational functions:

Example 2.5.6. Let K be the field $K = \mathbb{C}(X_1, \dots, X_n)$, i.e., the set of all rational functions in the formal variables X_1, \dots, X_n with addition and multiplication. Then,

$$\Omega_{K/\mathbb{C}}^1 = \left\{ \sum_{i=1}^n f_i dX_i ; f_i \in K \right\}.$$

That means, if $f = f(X_1, \dots, X_n)$ is any rational function in the X_i , we can always write df in the form

$$df = \sum_{i=1}^n f_i dX_i \quad \text{for some } f_i \in K.$$

It is also known from basic calculus what the f_i are, and that, given f , they are unique. Namely,

$$f_i = \frac{\partial f}{\partial X_i}.$$

For example,

$$d\left(\frac{X_1}{X_2}\right) = \left(\frac{1}{X_2}\right) dX_1 - \left(\frac{X_1}{X_2^2}\right) dX_2.$$

To the applied community, the formal operator d may also be known as the so-called total derivative.

If K is a rational function field, this behavior always occurs:

Proposition 2.5.7. Let $K = \mathbb{C}(X_1, \dots, X_n)$. Then, $\Omega_{K/\mathbb{C}}^1$ is the n -dimensional K -vector space, spanned by the formal differentials dX_i , $1 \leq i \leq n$. Given $f \in K$, there exist unique $f_1, \dots, f_n \in K$ such that

$$df = \sum_{i=1}^n f_i dX_i.$$

Proof. This follows from the uniqueness of the partial derivative of a complex rational function. □

Definition 2.5.8. The rational functions $f_i \in K$ from Proposition 2.5.7 are called formal partial derivative (of f with respect to X_i) and denoted by

$$\frac{\partial f}{\partial X_i} = f_i.$$

If K is a rational function field, the differentials can also be evaluated with respect to some point:

Definition 2.5.9. Let $K = \mathbb{C}(X_1, \dots, X_n)$. Let $f \in K$. For $P \in \mathbb{C}^n$, we define the evaluation of df at the point P as

$$df|_P = \sum_{i=1}^n \frac{\partial f}{\partial X_i}(P) dX_i.$$

Note that any evaluation yields a vector in the n -dimensional \mathbb{C} -vector space, spanned by the formal differentials $dX_i, 1 \leq i \leq n$.

The following classical result relates algebraic dependence to linear dependence:

Proposition 2.5.10. Let K be a field over \mathbb{C} , let $\alpha_1, \dots, \alpha_n \in K$. Then, $\alpha_1, \dots, \alpha_n$ are algebraically dependent if and only if $d\alpha_1, \dots, d\alpha_n$ are linearly dependent¹⁷ in $\Omega_{K/\mathbb{C}}^1$ (considered as a K -module).

Proof. Since K contains \mathbb{C} , the extension K/\mathbb{C} is always separable. The statement is, for example, implied by Proposition 16.14 in [11]. \square

Example 2.5.11. Let us consider the entries of a (2×2) -matrix of rank 1

$$A = \begin{pmatrix} a_{11} & a_{21} \\ a_{12} & a_{22} \end{pmatrix},$$

where we consider the entries as - possibly dependent - indeterminates subject to the equation

$$f(a_{11}, a_{12}, a_{21}, a_{22}) = a_{11}a_{22} - a_{12}a_{21} = 0.$$

Formally, the indeterminates live in the field $K = \mathbb{C}(a_{11}, a_{12}, a_{21}, a_{22})$. The equation above, by differentiating, gives a linear equation

$$df = a_{11} da_{22} + a_{22} da_{11} - a_{12} da_{21} - a_{21} da_{12} = 0.$$

Since the coefficients of all da_{ij} in this equation are non-zero polynomials, we see that any three of the da_{ij} are linearly independent. By Proposition 2.5.10, any three of the a_{ij} are algebraically independent. Indeed, if A is a generic matrix of rank 1, then any three of the four a_{ij} can be fixed independently, determining the remaining one up to a finite choice.

We now present the central result which will allow us to algorithmically test algebraic independence of the entries in matrix completion:

¹⁷Cave: the definition of linearly dependent in $\Omega_{K/\mathbb{C}}^1$, as a K -module, allows for coefficients in K , as opposed to coefficients on \mathbb{C} . That is, $d\alpha_1, \dots, d\alpha_n$ are linearly dependent, if and only if there exist $\lambda_1, \dots, \lambda_n \in K$, not all zero, such that $\sum_{i=1}^n \lambda_i d\alpha_i = 0$. Again note that this is different from linear dependence over \mathbb{C} .

Proposition 2.5.12. *Let $K = \mathbb{C}(X_1, \dots, X_m)$, let $f_1, \dots, f_n \in K$. Let $P \in \mathbb{C}^m$ be generic. Then, f_1, \dots, f_n are algebraically dependent if and only if $df_1|_P, \dots, df_n|_P$ are linearly dependent vectors in the m -dimensional \mathbb{C} -vector space spanned by the formal differentials dX_1, \dots, dX_m .*

Proof. By Proposition 2.5.10, it suffices to prove: $df_1|_P, \dots, df_n|_P$ are linearly dependent (over \mathbb{C}) if and only if df_1, \dots, df_n are linearly dependent (over K).

First we prove the if-direction. df_1, \dots, df_n are linearly dependent if and only if there exist $\lambda_1, \dots, \lambda_n \in K$, not all zero, such that

$$\sum_{i=1}^n \lambda_i df_i = 0.$$

Thus, we also have

$$\sum_{i=1}^n \lambda_i|_P df_i|_P = 0,$$

where the $\lambda_i|_P$ are not all zero due to the genericity of P . Thus we have proved that df_1, \dots, df_n are linearly dependent.

Now we prove the only if-direction. We assume df_1, \dots, df_n are linearly independent (thus, $n \leq m$) and show that $df_1|_P, \dots, df_n|_P$ are also linearly independent. Since we assumed df_1, \dots, df_n to be linearly independent, it follows, possibly after some reordering of the indices of f_i and X_i , that the Jacobi polynomial

$$J := \det \begin{pmatrix} \frac{\partial f_1}{\partial X_1} & \cdots & \frac{\partial f_1}{\partial X_n} \\ \vdots & \ddots & \vdots \\ \frac{\partial f_n}{\partial X_1} & \cdots & \frac{\partial f_n}{\partial X_n} \end{pmatrix}$$

is not the zero polynomial. Thus, the evaluation $J(P)$ will be non-zero due to the genericity¹⁸ of P . Thus,

$$0 \neq J(P) = \det \begin{pmatrix} \frac{\partial f_1}{\partial X_1}|_P & \cdots & \frac{\partial f_1}{\partial X_n}|_P \\ \vdots & \ddots & \vdots \\ \frac{\partial f_n}{\partial X_1}|_P & \cdots & \frac{\partial f_n}{\partial X_n}|_P \end{pmatrix},$$

and linear independence of $df_1|_P, \dots, df_n|_P$ follows. \square

Example 2.5.13. *Keep the notations of Example 2.5.11. Substituting any rank 1 matrix A with non-zero coefficients gives*

$$df = a_{11} da_{22} + a_{22} da_{11} - a_{12} da_{21} - \frac{a_{11}a_{22}}{a_{12}} da_{21} = 0.$$

If a_{11}, a_{12}, a_{22} are generically sampled, f will always give rise to a non-zero dependence between the da_{ij} . For example, if

$$A = \begin{pmatrix} 2 & 4 \\ 1 & 2 \end{pmatrix},$$

¹⁸a complex polynomial is non-zero if and only if it evaluates non-zero almost everywhere; follows, e.g., from the Schwarz-Zippel-Lemma, or the fact that algebraic sets have Lebesgue-measure zero

one obtains the evaluated equation

$$df = 2da_{22} + 2da_{11} - 4da_{21} - da_{21} = 0.$$

We reformulate the results stated so far by collecting the relevant consequences for matrix completion:

Theorem 2.5.14. *Let (N, M) be a partial mask. For $1 \leq i \leq m, 1 \leq j \leq n$, let a_{ij} be the formal variable for the (ij) -th entry of an $(m \times n)$ rank r -matrix, i.e., we present the ring of the determinantal variety $\mathcal{M}(m \times n, r)$ as*

$$\mathbb{C}[\mathcal{M}(m \times n, r)] = \mathbb{C}[a_{11}, \dots, a_{ij}, \dots, a_{mn}] / I(\mathcal{M}(m \times n, r)),$$

where $I(\mathcal{M}(m \times n, r))$ is the determinantal ideal of rank r . Then, the following are equivalent:

- (i) (N, M) is finitely identifiable in rank r .
- (ii) There exists a subset $S \subset E(M)$ such that the set $\mathcal{A} = \{a_{ij} ; (ij) \in S\}$ is algebraically independent over \mathbb{C} , and for any $(kl) \in E(N)$, the set $\{a_{kl}\} \cup \mathcal{A}$ is algebraically dependent.
- (iii) The field extension $\mathbb{C}(a_{ij}, (ij) \in E(N)) / \mathbb{C}(a_{ij}, (ij) \in E(M))$ is finite.
- (iv) There exists a subset $S \subset E(M)$ such that the set $\mathcal{A} = \{da_{ij} ; (ij) \in S\}$ is linearly independent, and for any $(kl) \in E(N)$, the set $\{da_{kl}\} \cup \mathcal{A}$ is linearly dependent (over $K = \mathbb{C}(a_{ij}, 1 \leq i \leq m, 1 \leq j \leq n)$).
- (v) Let $U \in \mathbb{C}^{m \times r}, V \in \mathbb{C}^{n \times r}$ be generic, let $A = U \cdot V^\top$. Let $V_M = \text{span}\{da_{ij} \mid_A ; (ij) \in E(M)\}$, and $V_N = \text{span}\{da_{ij} \mid_A ; (ij) \in E(N)\}$. One has $V_N \subseteq V_M$.

Proof. Let Ω be the partial masking defined by the mask (N, M) , i.e.,

$$\Omega : \Omega_N(\mathcal{M}(m \times n, r)) \rightarrow \Omega_M(\mathcal{M}(m \times n, r)).$$

The generic fiber dimension of Ω is exactly the transcendence degree of the field extension,

$$\begin{aligned} \dim \Omega &= \dim \Omega_N(\mathcal{M}(m \times n, r)) - \dim \Omega_M(\mathcal{M}(m \times n, r)) \\ &= \text{trdeg } \mathbb{C}(\Omega_N(\mathcal{M}(m \times n, r))) / \mathbb{C}(\Omega_M(\mathcal{M}(m \times n, r))). \end{aligned}$$

Also, since Ω_N and Ω_M are projections onto the variables a_{ij} in the respective edge sets, one has

$$\begin{aligned} \mathbb{C}(\Omega_N(\mathcal{M}(m \times n, r))) &= \mathbb{C}(a_{ij}, (ij) \in E(N)) \\ \mathbb{C}(\Omega_M(\mathcal{M}(m \times n, r))) &= \mathbb{C}(a_{ij}, (ij) \in E(M)). \end{aligned}$$

The equivalence of (i), (ii) and (iii) follows from the above equalities. The equivalence of (iii) and (iv) follows from Proposition 2.5.10, the equivalence of (iii) and (v) from Proposition 2.5.12. \square

Theorem 2.5.14 is in particular also a statement for masks as originally defined in Definition 2.1.4, by taking for N the matrix containing only ones. As another important consequence, one can immediately state the following:

Proposition 2.5.15. *Let M be an $(m \times n)$ -mask. Then, there is a unique biggest (w.r.t. number of ones) mask N such that (N, M) is finitely identifiable.*

Proof. Keeping the notations of Theorem 2.5.14(i) and (v), there is a unique biggest set S such that the vector space $V_S = \text{span}\{da_{ij} \mid (ij) \in S\}$ is contained in V_M . Taking N with edge set $E(N) := S$, and using the equivalence of (i) and (v), this proves the proposition. \square

Proposition 2.5.15 motivates the following definition:

Definition 2.5.16. *Let M be an $(m \times n)$ -mask. The unique biggest (w.r.t. number of ones) mask N such that (N, M) is finitely identifiable in rank r is called (finite) completable closure of M in rank r .*

The following remark shows how Theorem 2.5.14 (v) can be made into an algorithmic rule to determine finite completability:

Remark 2.5.17. *Generic dependencies between entries of the matrix A can be determined using the affine parameterization of low-rank matrices:*

Let $A \in \mathbb{C}^{m \times n}$ be a matrix of rank r or less; then, there exist matrices $U \in \mathbb{C}^{m \times r}$ and $V \in \mathbb{C}^{n \times r}$ such that $A = UV^\top$, and conversely, any matrix of the form $A = UV^\top$ has rank at most r . Writing $u_i, 1 \leq i \leq m$ for the rows of U and $v_j, 1 \leq j \leq n$ for the rows of V , one obtains the equation

$$a_{ij} = u_i v_j^\top \quad \text{for all } 1 \leq i \leq m, 1 \leq j \leq n.$$

Thus, one can consider the elements a_{ij} to be contained in the set of rational functions

$$K = \mathbb{C}(\dots, U_{ij}, \dots, V_{k\ell}, \dots)$$

in the $r \cdot (m+n)$ variables $U_{ij}, 1 \leq i \leq m, 1 \leq j \leq r$ and $V_{k\ell}, 1 \leq k \leq n, 1 \leq \ell \leq r$, which correspond to the (formal) entries of U and V . Thus, the equation

$$a_{ij} = u_i v_j^\top = \sum_{k=1}^r U_{ik} V_{jk}$$

gives rise to the differential expansion

$$da_{ij} = du_i \cdot v_j^\top + u_i \cdot dv_j^\top = \sum_{k=1}^r (V_{jk} dU_{ik} + U_{ik} dV_{jk}).$$

Thus, using Proposition 2.5.12, algebraic dependency for some set of a_{ij} can be evaluated by choosing some random generic value U_0, V_0 for U, V , and then testing for linear dependency of the $r \cdot (m+n)$ -dimensional vectors $da_{ij} \mid_{(U_0, V_0)}$ which live in the \mathbb{C} -vector space generated by all the formal differentials dU_{ij} and dV_{ij} .

Remark 2.5.17, together with Theorem 2.5.14 for the case of a non-partial masking, also shows correctness of Algorithm 3 by Singer and Cucuringu [43].

2.5.3 The Determinantal Matroid

In section 2.5.2, we have seen that dependence of entries in any low-rank matrix may be checked by calculating the vector space spanned by the tangent vectors at a generic low-rank matrix. More specific, Theorem 2.5.14 shows that dependent sets of entries of the matrix expose exactly the

same properties as basis elements of a vector space; algebraically independent entries behave like linearly independent vectors in a vector space. In the following, we could in principle use this established link to prove dependence and degrees of freedom properties of maskings and partial maskings. Instead, we will introduce an abstract concept which bundles the relevant combinatorial properties for both linear and algebraic dependencies, the matroid, which will then allow to derive and state the results in a concise and more readable manner.

We begin by axiomatically defining what a matroid is. A matroid generalizes properties of independent sets of vectors in a vector space:

Definition 2.5.18. *Let E be a set. A collection \mathcal{J} of subsets of E is called matroid (over E), if it fulfills the following condition:*

- (i) $\emptyset \in \mathcal{J}$.
- (ii) Let $J \in \mathcal{J}$, and $I \subseteq J$. Then $I \in \mathcal{J}$.
- (iii) Let $I, J \in \mathcal{J}$, let $\#I < \#J$. Then there is $e \in J$ and $e \notin I$ such that $I \cup \{e\} \in \mathcal{J}$.

The elements of \mathcal{J} are called the independent sets (of the matroid).

Intuitively, a matroid is the collection of all subsets of E which are independent in the domain of E (e.g., linearly or algebraically independent). Definition 2.5.18 (ii) states that subsets of independent sets are also independent, Definition 2.5.18 (iii) is (resp. implies) the fact that any set of linearly independent vectors in a finite-dimensional vector space can be extended to a basis.

Example 2.5.19. *As stated, sets and subsets of vectors or algebraic elements give rise to matroids:*

- (i) Let $V \cong \mathbb{C}^n$ be a vector space, and $v_1, \dots, v_k \in V$. Let $E = \{v_1, \dots, v_k\}$. Then a basic fact from Linear Algebra is that

$$\mathcal{J} := \{I \subseteq E ; I \text{ is linearly independent}\},$$

the collection of linearly independent subsets of S , is a matroid (over E).

- (ii) Let K/\mathbb{C} be a field extension, e.g. $K = \mathbb{C}(X_1, \dots, X_n)$. Let $\alpha_1, \dots, \alpha_k \in K$, let $E = \{\alpha_1, \dots, \alpha_k\}$. Then one can prove that

$$\mathcal{J} := \{I \subseteq E ; I \text{ is algebraically independent}\},$$

is a matroid (over E). This is a special instance of algebraic matroids, see [35, chapter 5].

That (ii) is indeed a matroid can also be seen by Theorem 2.5.14 which in fact gives a one-to-one correspondence to a linear matroid as in (i).

For reading convenience, we will introduce some of the usual matroid terminology:

Definition 2.5.20. *Let \mathcal{J} be a matroid over E , let $S \subseteq E$. Then we call*

- (i) S independent (w.r.t. \mathcal{J}) if $S \in \mathcal{J}$, else dependent
- (ii) S a circuit if it is minimally dependent¹⁹.

¹⁹i.e., $C \subseteq E$ is called a circuit if C is dependent and there does not exist $C' \subsetneq C$ such that C' is dependent (w.r.t. \mathcal{J})

(iii) $B \subseteq S$ a basis of S if it is an inclusion-wise maximal independent subset²⁰ of S . Bases of E are called bases of the matroid \mathcal{J} .

(iv) the cardinality of a basis²¹ of S the rank of S and denote it by $\text{rk}(S)$.

If \mathcal{J} is an algebraic, or linear matroid, we will at times add the qualifiers “algebraic” or “linear” to avoid confusion, e.g., algebraically dependent set, or algebraic circuit.

Matroids capture combinatorially the following facts which are well-known for finite vector configurations:

Proposition 2.5.21. *Let \mathcal{J} be a matroid over E . Then,*

(i) given $S \subseteq E$, one has $\text{rk}(S) \leq \#S$.

(ii) given $S \subseteq E$, one has $\text{rk}(S) = \#S$ if and only if S is independent.

(iii) $S \subseteq E$ is a circuit if and only if $\text{rk}(S) = \#S - 1$ and for all proper subsets $S' \subsetneq S$, it holds that $\text{rk}(S') = \#S'$.

(iv) given $S \subseteq E$, a subset $B \subseteq S$ is a basis of S if and only if $\#B = \text{rk}(B) = \text{rk}(S)$.

(v) a subset $B \subseteq E$ is a basis of \mathcal{J} if and only if for all $e \in E \setminus B$, $B \cup \{e\}$ contains a unique circuit.

Proof. The proofs of the statements are elementary and can be found in [35]: (i) to (iii) can be found the beginning of section 1.3, (iv) is Lemma 1.2.4., and (v) is Proposition 1.1.6. \square

Proposition 2.5.21 (i) and (ii) generalize the basis elimination principle from Linear Algebra.

One of the most important facts for our algebraic situation is that rank of a set of algebraic elements is exactly the number of degrees of freedom it contains:

Proposition 2.5.22. *Let \mathcal{J} be the algebraic matroid corresponding to some collection of elements $E = \{\alpha_1, \dots, \alpha_k\}$ over \mathbb{C} . For $S \subseteq E$, the rank $\text{rk} E$ is exactly the transcendence degree $\text{trdeg}(K/\mathbb{C})$, where K denotes the extension field $\mathbb{C}(\alpha; \alpha \in S)$ of \mathbb{C} .*

Proof. This is implied by the discussion between Examples 6.7.8 and 6.7.9 in Oxley [35] and the fact that an algebraic matroid is a matroid. \square

We will now introduce some matroid-related concepts which are unique for the problem of matrix completion, due to its inherent structures and symmetries:

Definition 2.5.23. *We will denote by $\mathcal{E}(m \times n, r)$ the set of entries a_{ij} , $1 \leq i \leq m$, $1 \leq j \leq n$ of a matrix with rank at most r , interpreted as possibly algebraically dependent variables over \mathbb{C} .*

We will denote by $\mathcal{D}(m \times n, r)$ the matroid over $\mathcal{E}(m \times n, r)$ consisting of algebraically independent subsets of $\mathcal{E}(m \times n, r)$. It is called the algebraic independence matroid of $\mathcal{E}(m \times n, r)$, or determinantal matroid.

To the elements of both $\mathcal{E}(m \times n, r)$ and $\mathcal{D}(m \times n, r)$, we will also refer by their respective indices. That is, we will simultaneously consider $\mathcal{E}(m \times n, r)$ to be the set $\{(ij) \in \mathbb{N}^2; 1 \leq i \leq m, 1 \leq j \leq n\}$, and we will simultaneously consider elements of $\mathcal{D}(m \times n, r)$ to be sets of pairs.

For $E \subseteq \mathcal{E}(m \times n, r)$, we will denote by $\text{rk}_r(E)$ the rank $\text{rk}(E)$ of E with respect to $\mathcal{D}(m \times n, r)$. If M is a $(m \times n)$ mask, we will also write $\text{rk}_r(M) = \text{rk}_r(E(M))$.

²⁰i.e., $B \subseteq S$ is called a basis of S if B is independent and there does not exist $S \supset B' \supsetneq B$ such that B' is independent (w.r.t. \mathcal{J})

²¹The matroid axioms imply that all bases of S are the same size [35, Lemma 1.2.4].

In the following, our goal is to make use of the fact that dependence structure of a low-rank matrix does not depend on the ordering of rows and columns; this implies additional structure for the determinantal matroid. In the proof of this, we need two lemmata:

Lemma 2.5.24. *Let $\Omega : \mathcal{M}(m \times n, r) \rightarrow \mathbb{C}^\alpha$ be a masking. Then,*

$$\text{rk}_r(E(\Omega)) = d_r(m, n) - \dim \Omega.$$

Proof. This follows directly from the fact that $\dim \Omega$ is the same as the transcendence degree of the field extension

$$\mathbb{C}(\mathcal{M}(m \times n, r))/\mathbb{C}(\Omega(\mathcal{M}(m \times n, r))).$$

Moreover, one has

$$\begin{aligned} \dim(\Omega) &= \text{trdeg}(\mathbb{C}(\mathcal{M}(m \times n, r))/\mathbb{C}(\Omega(\mathcal{M}(m \times n, r)))) \\ &= \text{trdeg}(\mathbb{C}(\mathcal{M}(m \times n, r))/\mathbb{C}) - \text{trdeg}(\mathbb{C}(\Omega(\mathcal{M}(m \times n, r)))/\mathbb{C}) \\ &= \dim(\mathcal{M}(m \times n, r)) - \text{rk}(E(\Omega)), \end{aligned}$$

from which the statement follows. □

Lemma 2.5.25. *For arbitrary permutation matrices $P \in \mathbb{C}^{m \times m}$ and $Q \in \mathbb{C}^{n \times n}$, define a map*

$$\begin{aligned} \mu(P, Q) : \mathcal{M}(m \times n, r) &\rightarrow \mathcal{M}(m \times n, r) \\ A &\mapsto P \cdot A \cdot Q. \end{aligned}$$

Also, (for each m, n) define a map

$$\begin{aligned} \top : \mathcal{M}(m \times n, r) &\rightarrow \mathcal{M}(n \times m, r) \\ A &\mapsto A^\top. \end{aligned}$$

The maps $\mu(P, Q)$ and \top are well-defined, algebraic morphisms which are isomorphisms.

Proof. Well-definedness follows from the fact that the rank of a matrix cannot increase when multiplying with another matrix or transposing. The maps $M(P, Q)$ and \top are algebraic morphisms because the defining rules are algebraic. The maps are isomorphisms, since $M(P, Q) \circ M(P^{-1}, Q^{-1}) = \text{id}$ and $\top \circ \top = \text{id}$. □

Proposition 2.5.26. *Let M be an $(m \times n)$ -mask. Then the rank $\text{rk}_r(M)$ depends only²² on $G(M)$. Furthermore, $\text{rk}_r(M)$ is equal to $\text{rk}_r(M^\top)$.*

Proof. Let M and N be $(m \times n)$ -masks with maskings Ω_M and Ω_N . For the first statement, it suffices to prove that if there are permutation matrices $P \in \mathbb{C}^{m \times m}$ and $Q \in \mathbb{C}^{n \times n}$, such that $M = PNQ$, then $\text{rk}_r(M) = \text{rk}_r(N)$. Consider

$$\Omega_M = \Omega_N \circ \mu(P, Q),$$

where $\mu(P, Q)$ is defined as in Lemma 2.5.25. By Lemma 2.5.25, $\mu(P, Q)$ is an isomorphism, so $\dim \mu(P, Q) = 0$, thus it holds that

$$\dim \Omega_M = \dim \mu(P, Q) + \dim \Omega_N = \dim \Omega_N.$$

²²I.e., $\text{rk}_r(M)$ does not depend on m, n , or the numbering of the vertices induced by the presentation in M , only on the unlabelled graph structure given by $G(M)$.

Applying Lemma 2.5.24 shows that $\text{rk}_r(M) = \text{rk}_r(N)$. Similarly,

$$\Omega = \Omega_{M^\top} \circ \top.$$

By Lemma 2.5.25, \top is an isomorphism, so $\dim \top = 0$, thus it holds that

$$\dim \Omega_M = \dim \top + \dim \Omega_{M^\top} = \dim \Omega_{M^\top}.$$

Applying Lemma 2.5.24 shows that $\text{rk}_r(M) = \text{rk}_r(M^\top)$. □

Proposition 2.5.26, together with Proposition 2.5.21 which characterizes the relevant concepts in term of rank, shows that the following are well-defined:

Definition 2.5.27. *If E is a subset of $\mathcal{E}(m \times n, r)$, we will denote by $G(E)$ the graph with edge set E , where we assume that m, n is minimal and no superfluous, isolated vertices are present.*

Let G be a bipartite graph with edge set $E \subseteq \mathcal{E}(m \times n, r)$. We will say that

- (i) *G is an independent graph (in rank r) if E is an independent set in $\mathcal{D}(m \times n, r)$.*
- (ii) *G is a circuit graph (in rank r) if E is an circuit in $\mathcal{D}(m \times n, r)$.*
- (iii) *H is a basis graph (in rank r) of G if $E(H)$ is a basis of E in $\mathcal{D}(m \times n, r)$.*
- (iv) *the r -rank $\text{rk}_r(G)$ of G is the rank $\text{rk}(E)$ of E w.r.t. $\mathcal{D}(m \times n, r)$.*

Proposition 2.5.26 also relates the completable closure to the concept of graph closure:

Proposition 2.5.28. *Let M be a $(m \times n)$ -mask with labeled bipartite graph G , let N be the completable closure of M in rank r , with labeled bipartite graph H . Let $\gamma : G \hookrightarrow H$ be the corresponding adjacency map. Then, γ is equal to the closing*

$$\text{cl}[\phi_1, \dots, \phi_k] : G \rightarrow \text{cl}[\phi_1, \dots, \phi_k](G) = H,$$

where the set of ϕ_i is the set of all injections of the form $C^- \hookrightarrow C$ such that C is a circuit graph, C^- is the graph C with one edge removed.

Proof. Let M, N be masks such that $N - M$ is a mask. The mask (N, M) is finitely completable if and only if the map

$$\Omega_{N/M} : \Omega_N(\mathcal{M}(m \times n, r)) \rightarrow \Omega_M(\mathcal{M}(m \times n, r))$$

is generically finite. The latter is equivalent to the field extension

$$(\mathbb{C}(\Omega_N(\mathcal{M}(m \times n, r))) / \mathbb{C}(\Omega_M(\mathcal{M}(m \times n, r))))$$

being finite. By Proposition 2.5.22, and the definition of rank, this is equivalent to

$$\text{rk}_r(N) = \text{rk}_r(M).$$

Thus, by uniqueness of the completable closure in Proposition 2.5.15, the mask N is the completable closure of M if and only if $E(N)$ is the biggest superset of $E(M)$ with the same rank. But that is the matroid closure of $E(M)$ with respect to $\mathcal{D}(m \times n, r)$, which can be characterized by closing circuits one-by-one, for a definition see chapter 1.4 of Oxley [35]. The latter is equivalent to the graph closure described above due to Proposition 2.5.26. □

Propositions 2.5.28 and 2.5.26 imply that the following definition on unlabeled bipartite graphs captures all the information about the generic completability of masks associated with the type.

Definition 2.5.29. *Let G be a bipartite graph. The closing of $H = \text{cl}[\phi_1, \dots, \phi_k](G)$ from Proposition 2.5.28 is called the completable closure of G in rank r .*

Proposition 2.5.26 also allows us to state a Corollary of Proposition 2.5.22 for matrix completion and partial matrix completion:

Corollary 2.5.30. *The following statements hold:*

(i) *Let Ω be some masking in rank r , let $G = G(\Omega)$. Then*

$$\dim \Omega = d_r(G) - \text{rk}_r(G),$$

where rank has to be taken in $\mathcal{D}(m \times n, r)$. In particular, Ω is finitely completable if and only if $\text{rk}_r(G) = d_r(G)$.

(ii) *Let Ω be some masking in rank r , let $G = G(\Omega)$. Then Ω is finitely completable if and only if $G(\Omega)$ contains a subgraph G' which is independent in rank r and has $e(G') = d_r(G)$.*

(iii) *Let Ω be some partial masking in rank r , with $(G \hookrightarrow H) = G(\Omega)$. Then*

$$\dim \Omega = \text{rk}_r(H) - \text{rk}_r(G).$$

In particular, Ω is finitely completable if and only if $\text{rk}_r(G) = \text{rk}_r(H)$.

Proof. (i) is a reformulation of Lemma 2.5.24, together with the fact that finite completability of Ω is equivalent to $\dim \Omega = 0$.

For (ii), note that (i) implies for finitely completable Ω that $\text{rk}_r(G) = d_r(G)$. By Proposition 2.5.21 (iv) and Proposition 2.5.26, this is equivalent to the fact that there is a subgraph G' of G with $\text{rk}_r(G') = e(G') = d_r(G)$. Note that the latter condition implies that G' is an independent graph.

(iii) follows from Lemma 2.5.24, by applying it to a partial masking (N, M) and the map $\Omega_{N/M}$, while using $\Omega_M = \Omega_{N/M} \circ \Omega_N$ and the dimension formula

$$\dim \Omega_M = \dim \Omega_{N/M} + \dim \Omega_N.$$

□

Corollary 2.5.30 is more than a mere reformulation of the previous results on fiber dimension, since it implicitly references the matroidal structure induced by dependence of the matrix entries. In particular, the matroidal structure can be used to prove:

Theorem 2.5.31. *If a masking Ω is generically finite in rank r , then $G(\Omega)$ is spanned in rank r . In particular, all conditions (i), (ii), (iii) from Proposition 2.3.34 are all necessary for Ω to be generically finite or generically injective.*

Proof. The matroidal property of completability implies that it is no loss of generality to reduce to the case in which Ω is *minimally* finitely completable, i.e., when it is finitely completable but ceases to be so when *any* of its known entries are removed. Said differently, this is when Ω is a basis of $\mathcal{E}(m \times n, r)$.

Corollary 2.5.30 (i) implies that for any basis Ω in the determinantal matroid, $e(G(\Omega)) = d_r(m, n)$. A second application of Corollary 2.5.30 (i) implies that *any* masking Ω' with $e(G(\Omega')) \geq d_r(G(\Omega')) + 1$ contains a circuit in the determinantal matroid. Since bases are all independent, it follows that every subgraph G' of $G(\Omega)$ has $e(G') \leq d_r(G')$. This shows that $G(\Omega)$ is r -sparse with $d_r(m, n)$ edges; i.e., it is r -tight. \square

Given that finitely completable masks must be spanned in rank r , we might wonder if this is a sufficient condition. In fact, it is not. The intuition is that it is possible to glue rank- r -tight graphs along r vertices to get another rank- r -tight graph, but that finitely completable masks glued along less than an $r \times r$ block will necessarily have left over degrees of freedom.

Theorem 2.5.32. *For fixed $r > 1$, there exist infinitely many rank r -tight graphs that are not finitely completable.*

Proof. Let m and n be such $m \geq 2r$ and $n \geq r + 1$. Define the graph $G^{m,n}$ to have m blue vertices v_1, \dots, v_m and n red ones w_1, \dots, w_n . An edge (v_i, w_j) is present in $G^{m,n}$ if and only if $i \leq r$ or $j \leq r$. The graph $G^{m,n}$ is r -closable and has $d_r(m, n)$ edges. Thus, by Proposition 2.3.26, $G^{m,n}$ is finitely completable, and so Theorem 2.5.31 implies that it is rank- r -tight.

Now observe that the set of vertices $I := \{v_{r+1}, v_{r+2}, \dots, v_{2r-1}, w_{r+1}\}$ have no edges between them. Define the graph $H^{m,n}$ to be the graph obtained by gluing two copies G_1 and G_2 of $G^{m,n}$ along I such that red vertices are identified to red ones and similarly blue to blue. Because we didn't identify any edges when creating $H^{m,n}$, it has $2d_r(m, n) = e(G_1) + e(G_2)$ edges. The number of blue vertices in H is $2m - r + 1$, and the number of red vertices is $2n - 1$. Computing, we see that:

$$d_r(H^{m,n}) = d_r(2m - r + 1, 2n - 1) = r(2m + 2n - 2r) = 2d_r(m, n) = e(H^{m,n})$$

So if $H^{m,n}$ is rank- r -sparse, it is, in addition, rank- r -tight. We will now show that: (i) $H^{m,n}$ is rank- r -sparse; (ii) $H^{m,n}$ is not finitely completable.

(i) Let H be any subgraph of $H^{m,n}$. We will show that $e(H) \leq d_r(H)$. Denote by H_i the subgraph $H \cap G_i$, and let m_i and n_i be the number of blue and red vertices in H_i . Because $\#V(G_1 \cap G_2) = r$, we have $\#V(H) \geq \#V(H_1) + \#V(H_2) - r$. A calculation shows that

$$e(H) = e(H_1) + e(H_2) \leq d_r(m_1, n_1) + d_r(m_2, n_2) \leq d_r(H)$$

The leftmost equality follows from the fact that G_1 and G_2 share no edges; the first inequality follows from the rank- r -sparsity of the G_i ; the last follows from the lower bound on $\#V(H)$ given above.

(ii) Consider the edge $e = (v_{r+1}, w_{r+1})$ that has both of its endpoints in the set I defined above. We add e to G_1 , Proposition 2.5.21 (v) implies a unique circuit exists in $G_1 \cup \{e\}$ and similarly in $G_2 \cup \{e\}$. It then follows that $H^{m,n} \cup \{e\}$ contains both of these circuits and they are distinct. Thus, by Proposition 2.5.21 (v), $H^{m,n}$ is not a basis in the completion matroid. Since it has $d_r(H^{m,n})$ edges, Corollary 2.5.30 implies that $H^{m,n}$ is not finitely completable. \square

In particular, this implies:

Corollary 2.5.33. *Let $r > 1$. Then there are rank- r -sparse circuit graphs.*

Examining the proof of Theorem 2.5.32, we see that it did not depend too much on the specific structure of the graph $G^{m,n}$. The important properties were that: (i) $G^{m,n}$ is minimally finitely completable; (ii) the existence of the set of vertices I with no edges between them and at least one red and one blue vertex.

Proposition 2.5.34. *Let $r > 1$, and let G_1 and G_2 be sufficiently large (depending on r) basis graphs. Then, there is a set I_1 of r vertices in G_1 and a set I_2 of r vertices in G_2 such that, if we glue G_1 and G_2 by identifying I_1 with I_2 , the resulting graph is rank r -tight but not finitely completable.*

Proof. The proof of Theorem 2.5.32 applies nearly verbatim, once we have shown that, for sufficiently large basis graphs G , there is a set I of $r - 1$ blue vertices and one red vertex with no edges between them (or symmetrically, with the role of the colors exchanged). Clearly, there are no edges with two blue endpoints.

For the moment, fix $n > r$. We will show that there is an m_0 such that $m \geq m_0$ implies that some red vertex w_j has at most $m - r + 1$ neighbors. In this case, w_j and any $r - 1$ non-neighbors among the blue vertices give the desired set.

Suppose now that each red vertex has at least $m - r + 2$ neighbors. Since G has $d_r(m, n)$ edges, this implies that $d_r(m, n) \geq n(m - r + 2)$. A computation shows that

$$m \leq \frac{2(r-1)n - r^2}{n-r}$$

We take m_0 to be the least integer larger than the r.h.s. above. It can be shown that the r.h.s. achieves its maximum at $n = r + 1$, and thus, by symmetry, if both m and n are at least $r^2 - 1$, the statement holds. \square

We now develop some further properties of circuits in the determinantal matroid. Theorem 2.5.31 and the matroidal property imply that while there are circuits with fewer than $d_r(G) + 1$ edges, they cannot have more. Combined with a degree lower bound for circuit, we can show that the number of red and blue vertices cannot be too unbalanced in a circuit, which implies a bound on the number of circuit graphs with m red vertices.

Proposition 2.5.35. *A circuit graph in rank r has vertex degrees at least $r + 1$. In particular, a circuit graph with m and n red and blue vertices always has $m > r$ and $n > r$.*

Proof. Let $G = (V, E)$ be a circuit graph. Theorem 2.5.14 and Remark 2.5.17 imply that the rank resp. dimension of the \mathbb{C} -vector space, generated by the differentials

$$da_{ij} = du_i \cdot v_j^\top + u_i \cdot dv_j^\top, (ij) \in E$$

must be $\#E - 1$, where $u_i, 1 \leq i \leq m$ and $v_j, 1 \leq j \leq n$ are generic vectors in \mathbb{C}^r , and the components of the r -vectors du_i resp. dv_j are formal basis elements. Equivalently reformulated, this means that there are $\lambda_{ij} \in \mathbb{C}, (ij) \in E$, not all zero, such that

$$\sum_{(ij) \in E} \lambda_{ij} da_{ij} = 0,$$

and that none of the λ_{ij} can be chosen zero if at least one is non-zero. Using the above representation in the basis given by du_i and dv_j , the condition becomes

$$\begin{aligned} 0 &= \sum_{(ij) \in E} \lambda_{ij} \left(du_i \cdot v_j^\top + u_i \cdot dv_j^\top \right) \\ &= \sum_{i=1}^m du_i \sum_{(ij) \in E} \lambda_{ij} v_j^\top + \sum_{j=1}^n dv_j \sum_{(ij) \in E} \lambda_{ij} u_i^\top. \end{aligned}$$

Since the components of du_i and dv_j form a basis of the vector space of differentials, this implies that there are non-zero λ_{ij} such that

$$0 = \sum_{(ij) \in E} \lambda_{ij} v_j^\top \text{ for any (arbitrary but fixed) } i \quad \text{and} \quad 0 = \sum_{(ij) \in E} \lambda_{ij} u_i^\top \text{ for any (a.b.f.) } j.$$

Since the u_i, v_j are generic, and $u_i, v_j \in \mathbb{C}^r$, this can hold only if $N_i \geq r + 1$ and $N_j \geq r + 1$, where

$$N_i = \#\{(ij); (ij) \in E\} \quad \text{and} \quad N_j = \#\{(ij); (ij) \in E\}$$

(note that the definitions of N_i and N_j implies that i resp. j are arbitrary but fixed). Since N_i and N_j are the vertex degrees of the vertex i resp. the vertex j , this implies that each vertex in G has degree at least $r + 1$, which was the statement to prove. \square

Proposition 2.5.36. *The number of blue vertices in a circuit graph in rank r with m red vertices is at most $r(m - r) + 1$.*

Proof. Let G be a circuit graph in rank r with m red vertices and n blue vertices. As noted above, Theorem 2.5.31 implies that $e(G) \leq d_r(G) + 1$, and the degree lower bound Proposition 2.5.35 gives $d_r(G) = r(m + n - r)$. Estimating the number of edges in G from below using, again, Proposition 2.5.35, we get $n(r + 1) \leq r(m + n - r) + 1$. \square

Corollary 2.5.37. *The number of circuits in rank r with m red vertices is at most $2^{mr(m-r)+m}$.*

Bernd Sturmfels and Zvi Rosen have told us they obtained, independently, a similar result with a weaker conclusion. In rank one and $m - 1$ we can give an exact characterization of the circuit graphs.

Proposition 2.5.38. *The following hold:*

- (i) *The circuit graphs in rank $r = 1$ are exactly the cycles.*
- (ii) *The unique circuit graph in rank $r = m - 1$ is exactly $K_{m,m}$.*

Proof. By Lemma 2.3.25 and Proposition 2.3.26, the determinantal matroid is isomorphic to the graphic matroid, which has as its circuits the cycles [35, Proposition 1.1.7]. This proves (i).

For (ii), Proposition 2.3.26 implies that the almost biclique is finitely completable, and, since it has $d_{m-1}(m, m)$ edges, independent. Thus $K_{m,m}$ is a circuit. By Proposition 2.5.35, in any other circuit graph G , every blue vertex must be connected to all of the red vertices, forcing G to contain a copy of $K_{m,m}$, contradicting minimality of circuits. \square

2.6 Completability of random masks

Up to this point we have considered the generic completability of a fixed mask, which we have shown to be equivalent to questions about the associated bipartite graph. We now turn to the case where the masking is sampled at random, which, by Corollary 2.5.30, implies that, generically, this is a question about *random bipartite graphs*.

2.6.1 Random graph models

A *random graph* is a graph valued random variable. We are specifically interested in two such models for bipartite random graphs:

Definition 2.6.1. *The Erdős-Rényi random bipartite graph $G(m, n, p)$ is a bipartite graph on m red and n blue vertices with each edge present with probability p , independently.*

Definition 2.6.2. *The (d, d') -biregular random bipartite graph $G(m, n, d, d')$ is the uniform distribution on graphs with m red vertices, n blue ones, and each red vertex with degree d and each blue vertex with degree d' .*

Clearly, we need $md = nd'$, and if $m = n$, the (d, d') -regular random bipartite graph is, in fact d -regular.

We will call a mask corresponding to a random graph a *random mask*. We now quote some standard properties of random graphs we need.

Proposition 2.6.3. (Connectivity threshold) *The threshold for $G(m, n, p)$ to become connected, w.h.p., is $p = \Theta((m + n)^{-1} \log n)$ [3, Theorem 7.1].*

(Minimum degree threshold) *The threshold for the minimum degree in $G(n, n, p)$ to reach d is $p = \Theta((m + n)^{-1} (\log n + d \log \log n + \omega(1)))$. When $p = cn$, w.h.p., there are isolated vertices [3, Exercise 3.2].*

(d -regular connectivity) *With high probability, $G(m, n, d, d')$ is d -connected [3, Theorem 7.3.2]. (Recall that we assume $m \leq n$).*

(Density principle) *Suppose that the expected number of edges in either of our random graph models is at most Cn , for constant C . Then for every $\epsilon > 0$, there is a constant c , depending on only C and ϵ such that, w.h.p., every subgraph of n' vertices spanning at least $(1 + \epsilon)n'$ edges has $n' \geq cn$ [21, Lemma 5.1].*

(Emergence of the k -core) *Define the k -core of a graph to be the maximal induced subgraph with minimum k . For each k , there is a constant c_k such that $p = c_k/n$ is the first-order threshold for the k -core to emerge. When the k -core emerges, it is giant and afterwards its size and number of edges spanned grows smoothly with p [36].*

2.6.2 Completability of incoherent matrices

A fundamental result in the area of matrix completion, proven independently in the papers Candès and Tao [6], Keshavan et al. [24] is

Theorem 2.6.4. *Let A be an incoherent rank r matrix, with $r = O(1)$. Then, with high probability, an Erdős-Rényi mask with $p = \Theta(rn(\log n)^2)$ is sufficient to complete A uniquely.*

We note that the conclusion is not that the mask is *generically uniquely completable*, since the (crucial) incoherence assumption is about the underlying matrix A . In the next section, we will give a generic version of Theorem 2.6.4.

In a sense, 2.6.4 is the best possible. There are incoherent matrices with a block diagonal structure such that *no* sparser sampling can guarantee even finite completable with high probability Candès and Tao [6].

In the generic case, Theorem 2.5.31 implies that a finitely completable mask requires minimum degree r . The minimum degree threshold in the Erdős-Rényi model gives a similar lower bound. Combined with the methods of Section 2.7 below, we may conclude.

Proposition 2.6.5. *Let r be a fixed constant. There are constants c and C (depending on r) such that, if $p = c(n + m)^{-1} \log n$ then, w.h.p., $G(m, n, p)$ is not finitely completable and if $p = C(n + m)^{-1} (\log n)^2$ then, w.h.p., $G(m, n, p)$ is finitely completable.*

2.6.3 Sparser sampling and partial completable

The lower bounds on sample size for completion of rank r *incoherent* matrices do *not* carry over verbatim to the generic setting of this paper. This is because genericity and incoherence are related, but incomparable concepts: there are generic matrices that are not incoherent (consider a very small perturbation of the identity matrix); and, importantly, the block diagonal examples showing the lower bound for incoherent completable are not generic, since many of the entries are zero.

Thus, in the generic setting, we expect sparse sampling to be more powerful. This is demonstrated experimentally in Section 4.2. In the rest of this section, we derive some heuristics for the expected generic completable behavior of sparse random masks. We are particularly interested in the question of: *when are $\Omega(mn)$ of the entries completable from a sparse random mask?* We call this the *completable transition*. We will conjecture that there is a sharp threshold for the completable transition, and that the threshold occurs well below the threshold for $G(n, m, p)$ to be completable.

Let c be a constant. We first consider the emergence of a circuit in $G(n, n, c/n)$. Proposition 2.5.35 implies that any circuit is a subgraph of the $(r + 1)$ -core. By Theorem 2.5.31, having a circuit is a monotone property, which occurs with probability one for graphs with more than $2rn$ edges, and thus the value

$$t_r := \sup\{t : G(n, n, t/n) \text{ is } r\text{-independent, w.h.p.}\}$$

is a constant. If we define C_r as

$$C_r := \sup\{c : \text{the } (r + 1)\text{-core of } G(n, n, c/n) \text{ has average degree at most } 2r, \text{ w.h.p.}\}$$

smoothness of the growth of the $(r + 1)$ -core implies that we have

$$c_{r+1} \leq t_r \leq C_{r+1}$$

where we recall that c_{r+1} is the threshold degree for the $(r + 1)$ -core to emerge. Putting things together we get:

Proposition 2.6.6. *There is a constant C_r such that, if $c < t_r$ then w.h.p., $G(n, n, c/n)$ is r -independent, and, if $c > t_r$ then w.h.p. $G(n, n, c/n)$ contains a giant r -circuit inside the $(r + 1)$ -core. Moreover, t_r is at most the threshold for the $(r + 1)$ -core to reach average degree $2r$.*

Proposition 2.6.6 gives us some structural information about where to look for rank r circuits in $G(n, n, c/n)$: they emerge suddenly inside of the $(r + 1)$ -core and are all giant when they do. If rank r circuits were themselves finitely completable, this would then yield a threshold for the completability transition. Unfortunately, Theorem 2.5.32 tells us that this is not, in general, the case. Nonetheless, we conjecture:

Conjecture 2.6.7. *The constant t_r is the threshold for the completability transition in $G(n, n, c/n)$. Moreover, we conjecture that almost all of the $(r + 1)$ -core is completable above the threshold.*

We want to stress that the conjecture includes a conjecture about the *existence* of the threshold for the completability transition, which hasn't been established here, unlike the existence for the emergence of a circuit. The subtlety is that we haven't ruled out examples of r -independent graphs with no rank- r -spanning subgraph for which, nonetheless, the closure in the rank r completion matroid is giant. Conjecture 2.6.7 is explored experimentally in Sections 4.1 and 4.2. The conjectured behavior is analogous to what has been proved for distance matrices (also known as *bar-joint frameworks*) in dimension 2 in [23].

Our second conjecture is about $2r$ -regular masks.

Conjecture 2.6.8. *With high probability $G(n, n, 2r, 2r)$ is finitely completable. Moreover, we conjecture that it remains so, w.h.p., after removing r^2 edges uniformly at random.*

We provide evidence in Section 4.2. This behavior is strikingly different than the incoherent case, and consistent with proven results about 2-dimensional distance matrices (frameworks) [20, Theorem 4.1].

2.6.4 Denser sampling and solving minor by minor

The conjectures above, even if true, provide only information about matrix *completeness* and not matrix *completion*. In fact, the convex relaxation of Candès and Recht [5] does not seem to do very well on $2r$ -regular masks in our experiments, and the density principle for sparse random graphs implies that, w.h.p., a $2r$ -regular mask has no dense enough subgraphs for our closability algorithm in section 3.2 to even get started. Thus it seems possible that these instances are quite “hard” to complete even if they are known to be completable.

If we consider denser random masks, then the closability algorithm becomes more practical. A particularly favorable case for it is when every missing entry is part of some $K_{r+1, r+1}^-$. In this case, the error propagation will be minimal and, heuristically, finding a $K_{r+1, r+1}^-$ is not too hard, even though the problem is NP-complete in general.

Define the *1-step r -closure* of a bipartite graph G as the graph G' obtained by adding the missing edge to each $K_{r+1, r+1}^-$ in G . If the 1-step closure of G is $K_{n, n}$, we define G to be *1-step r -closable*. We can give an upper bound on the threshold for 1-step r -closability.

Theorem 2.6.9. *There is a constant $C > 0$ such that, if $p = Cn^{-2/(r+2)} \log n$ then, w.h.p., $G(n, n, p)$ is 1-step r -closable.*

Proof. Fix r and set p as in the statement. The probability of a specific copy of $K_{r+1, r+1}^-$ appearing is $p^{(r+1)^2-1}$ and there are $\Theta(n^{2r+2})$ potential copies. Since $K_{r+1, r+1}^-$ is its own least probable subgraph, we see that if $p = Cn^{-2/(r+2)}$ the expected number X of edge disjoint copies of $K_{r+1, r+1}^-$ in $G(n, n, p)$ is at least $C'n^2 \log n$ for some absolute constant C' depending on C .

A fundamental result about the number of copies of a small subgraphs [22, Theorem 3.29] implies that X is sharply concentrated around its expectation, so, w.h.p, $C''^{-1}n^2 \log n \leq X \leq C''n^2 \log n$ for a constant C'' depending only on r and C .

We now define the \mathcal{E} to be the event that $G(n, n, p)$ is 1-step r -closable. Also define the event \mathcal{B} to be the event that, in $G(n, n, p)$, no pair of vertices (i, j) is the “missing” edge in more than $D \log n$ copies of $K_{r+1, r+1}^-$, for some sufficiently large constant D . Since both \mathcal{E} and $\neg \mathcal{B}$ are both increasing events, the FKG inequality (e.g., [22, Theorem 2.12]) implies that $\Pr(\mathcal{E}) \leq \Pr(\mathcal{E}|\neg \mathcal{B})$. Using this estimate we get

$$\begin{aligned} \Pr(\mathcal{E}) &= \Pr(\mathcal{E}|\mathcal{B})\Pr(\mathcal{B}) + \Pr(\mathcal{E}|\neg \mathcal{B})\Pr(\neg \mathcal{B}) \\ &\geq \Pr(\mathcal{E}|\mathcal{B})\Pr(\mathcal{B}) + \Pr(\mathcal{E})\Pr(\neg \mathcal{B}) \end{aligned}$$

Rearranging, we conclude that $\Pr(\mathcal{E}) \geq \Pr(\mathcal{E}|\mathcal{B})$.

Conditioning on \mathcal{B} and the fact that there are at least $C''^{-1}n^2 \log n$ edge disjoint copies of $K_{r+1, r+1}^-$, we consider the process that reveals each of these copies one at a time. Since, for each pair of vertices (i, j) , the probability that the next revealed copy has (i, j) as its missing edge is $\Omega(1/n^2)$, and C'' is arbitrary, we can increase the probability that any fixed pair is covered by at least one copy to $1 - 1/n^3$ and then apply a union bound. \square

An interesting question is determining the threshold for $G(n, n, p)$ to be r -closable. Experimentally, it appears that: $p = n^{-2/(r+2)}$ is, in fact the order of the true threshold; when $G(n, n, p)$ is closable it is $O(1)$ -step closable.

2.7 Sufficient Sampling Densities for Algebraic Compressed Sensing

For different conditioned sampling methods of both matrices and masks, the asymptotic behavior of completion has been well-analyzed, most notably in the case of uniform sampling of masks, see e.g. Candès and Recht [5], Candès and Tao [6], Gross [16]. While much is implied between the lines, none of the available literature addresses the question of conditioning only on the mask while removing the conditioning on the matrices directly.

In fact, it turns out that in a novel more general algebraic framework, the analytic arguments found in the previous work can be modified to provide identifiability results in the setting where a point on an algebraic variety is to be reconstructed from a set of random projections. Particularly, with Theorem 2.7.8, we will obtain a result which relates the necessary number of observations directly to intrinsic properties of the variety (namely, its incoherence, which we will define), notably without further conditioning how the point of the variety was sampled. We believe that this result is the canonical expression of a general principle in compressed sensing that relates the necessary sampling density to properties of the signal space without further assumptions.

2.7.1 Finiteness of Random Maskings

In the following, we will examine compressed sensing under algebraic constraints. That is, given a signal $x \in \mathbb{C}^n$, where the inclusion into \mathbb{C}^n is to be considered as a parametric representation of the signal, and given an algebraic variety $X \subseteq \mathbb{C}^n$ which describes the compression constraints, such that $x \in X$, we attempt to reconstruct the signal x from random coordinate projections of x , under consideration of the compression constraints X . The main result of this section will

characterize the sampling density, i.e., the number of random coordinate projections of x needed to reconstruct a generic x , in terms of X , without further sampling assumptions on x .

As a corollary, we will obtain upper reconstruction bounds for matrix completion, where x is a low-rank matrix, and X is the variety of low-rank matrices $\mathcal{M}(m \times n, r)$.

First we introduce some formal concepts which describe the setting of compressed sensing under algebraic constraints, in particular the sampling process which we will assume to randomly, independently and uniformly sample coordinate projections of the signal without repetition.

Definition 2.7.1. *Let $X \subseteq \mathbb{C}^n$ be an algebraic varieties. Fix coordinates (X_1, \dots, X_n) for \mathbb{C}^n . Let $S(p)$ be a the Bernoulli random experiment yielding a random subset of $\{X_1, \dots, X_n\}$ where each X_i is contained in $S(p)$ independently with probability p . We will call the projection map*

$$\begin{aligned} \Omega : X &\rightarrow Y \\ (x_1, \dots, x_n) &\mapsto (\dots, x_i, \dots : X_i \in S(p)) \end{aligned}$$

of X onto the coordinates in $S(p)$, which is an algebraic-map-valued random variable, an algebraic random masking of X with selection probability p .

Intuitively, Ω takes a signal x from the signal space \mathbb{C}^n , fulfilling the constraints in X , and independently samples a Bernoulli set $\Omega(x)$ of random coordinate projections x_i with sampling density p .

The constraints in X will play a crucial role in determining the necessary sampling density which allows reconstruction of the signal. Namely, the central property of X which will determine the necessary density is the so-called coherence, which describes the degree of randomness of a generic tangent plane to X ; intuitively, it can be interpreted as the infinitesimal randomness of a signal.

Definition 2.7.2. *Let H be a k -flat²³ in \mathbb{C}^n . Let $\mathcal{P} : \mathbb{C}^n \rightarrow H \subseteq \mathbb{C}^n$ the orthogonal projection operator onto H , let e_1, \dots, e_n a fixed orthonormal basis of \mathbb{C}^n . Then the coherence of H with respect to the basis e_1, \dots, e_n is defined as*

$$\text{coh}(H) = \max_{1 \leq i \leq n} \|\mathcal{P}(e_i) - \mathcal{P}(0)\|^2$$

The coherence of a k -flat is bounded in both directions:

Proposition 2.7.3. *Let H be a k -flat in \mathbb{C}^n . Then,*

$$\frac{k}{n} \leq \text{coh}(H) \leq 1,$$

and both bounds are achieved.

Proof. Without loss of generality, we can assume that $0 \in H$ and therefore that \mathcal{P} is linear, since coherence, as defined in Definition 2.7.2, is invariant under translation of H .

First we show the upper bound. For that, note that for an orthogonal projection operator $\mathcal{P} : \mathbb{C}^n \rightarrow \mathbb{C}^n$ and any $x \in \mathbb{C}^n$, one has $\|\mathcal{P}(x)\| \leq \|x\|$. Thus, by definition,

$$\text{coh}(H) = \max_{1 \leq i \leq n} \|\mathcal{P}(e_i)\|^2 \leq \max_{1 \leq i \leq n} \|e_i\|^2 = 1.$$

²³A k -flat is a linear subspace of dimension k which does not necessarily contain 0. Other names are affine subspace or affine linear variety.

For strictness, take H as the span of e_1, \dots, e_k .

Let us now show the lower bound. We proceed by contradiction. Assume $\|P(e_i)\|^2 < \frac{k}{n}$ for all i . This would imply

$$k = n \cdot \frac{k}{n} > \sum_{i=1}^n \|\mathcal{P}(e_i)\|^2 = \|\mathcal{P}\|_2^2 = k$$

which is a contradiction, where in the last equality we used the fact that orthonormal projections onto a k -dimensional space have Frobenius norm k .

The tightness of the lower bound follows²⁴ from the existence of tight frames [7], by taking the rows of any tight frame of n vectors in \mathbb{R}^k as coordinate vectors, and H the span of e_1, \dots, e_k , where e_1, \dots, e_n is the standard orthonormal basis. \square

A similar definition of coherence, as in Definition 2.7.2, was used by Candès and Recht [5]; we decided to remove dimensional normalization in order to make the definition more intrinsic, i.e., not to depend on the dimension of the embedding. For completeness, we also state the original concept:

Definition 2.7.4. *Let H be a k -flat in \mathbb{C}^n . The normalized coherence, or coherence in the sense of Candès and Recht [5], is the quantity $\frac{n}{k} \text{coh}(H)$.*

In our definition, one always has that $\text{coh}(H) \leq 1$, possibly attaining the upper bound, while the normalized coherence has 1 as a possibly attainable lower bound. A different version of Proposition 2.7.3 was implicitly stated, but not proved in Candès and Recht [5]. While it is straightforward and probably folklore, we decided to state it nevertheless since it will play an important role in proving Proposition 2.7.10 which allows to apply the main results of this section to matrix completion.

Definition 2.7.5. *Let $X \subseteq \mathbb{C}^n$ be an (real or complex) irreducible analytic variety of dimension d (affine or projective). Let $x \in X$ a smooth point, and let $T_{X,x}$ be the tangent d -flat of X at x . We define*

$$\text{coh}(x \in X) := \text{coh}(T_{X,x}).$$

If it is clear from the context in which variety we consider x to be contained, we also write $\text{coh}(x) = \text{coh}(x \in X)$. Furthermore, we define the coherence of X to be

$$X = \inf_{x \in \text{Sm}(X)} \text{coh}(x),$$

where $\text{Sm}(X)$ denotes the smooth locus of X .

Note that if X is a k -flat, then the definitions of $\text{coh}(X)$, given by Definitions 2.7.2 and 2.7.5 agree. Also, if X is projective, then X^* is compact, so the infimum is in fact a minimum. These observations, together with Proposition 2.7.3, imply

Proposition 2.7.6. *Let X be a complex algebraic variety in \mathbb{C}^n . Then,*

$$\frac{1}{n} \dim X \leq \text{coh}(X) \leq 1,$$

and both bounds are achieved.

²⁴We cordially thank Andriy Bondarenko for pointing this out.

Definition 2.7.7. A complex algebraic variety X is called *maximally incoherent* if

$$\text{coh}(X) = \frac{1}{n} \dim X.$$

The following theorem relates the coherence of a variety X to the sampling density of a generic (constrained) signal $x \in X$, which is needed to achieve reconstruction of x , up to finite choice. The proof integrates some ideas of Candès and Recht [5] into our general algebraic setting. Also, the proof uses two lemmata, namely Lemmata 2.7.12 and 2.7.13, which can be found at the end of the section.

Theorem 2.7.8. *Let $X \subseteq \mathbb{C}^n$ be an irreducible algebraic variety, let Ω be an algebraic random masking with selection probability p , let $x \in X$ be a smooth point. There is an absolute constant C such that if*

$$p \geq C \cdot \lambda \cdot \text{coh}(X) \cdot \log n, \quad \text{with } \lambda \geq 1,$$

then Ω is generically finite with probability at least

$$1 - 3n^{-\lambda}.$$

Proof. By the definition of coherence, for every $\delta > 0$, there exists an x such that X is smooth at x , and $\text{coh}(x) \leq (1 + \delta) \text{coh}(X)$. Now let $y = \Omega(x)$, we can assume by possible changing x that $\Omega(X)$ is also smooth at y . Let T_y, T_x be the respective tangent spaces at y and x . Note that y is a point-valued discrete random variable, and T_y is a flat-valued random variable. By the equivalence of the statements (iv) and (v) in Lemma 2.7.12, it suffices to show that the operator

$$P = p^{-1} \theta \circ d\Omega - \text{id}$$

is contractive, where θ is projection, from T_y onto T_x , with probability at least $1 - 3n^{-\lambda}$ under the assumptions on p . Let $Z = \|P\|$, and let e_1, \dots, e_n be the orthonormal coordinate system for \mathbb{C}^n , and \mathcal{P} the projection onto T_x . Then the projection $\theta \circ d\Omega$ has, when we consider T_x to be embedded into \mathbb{C}^n , the matrix representation

$$\sum_{i=1}^n \varepsilon_i \cdot \mathcal{P}(e_i) \otimes \mathcal{P}(e_i),$$

where ε_i are independent Bernoulli random variables with probability p for 1 and $(1 - p)$ for 0. Thus, in matrix representation,

$$P = \sum_{i=1}^n \left(\frac{\varepsilon_i}{p} - 1 \right) \cdot \mathcal{P}(e_i) \otimes \mathcal{P}(e_i).$$

By Rudelson's Lemma 2.7.13, it follows that

$$\mathbb{E}(Z) \leq C \sqrt{\frac{\log n}{p}} \max_i \|\mathcal{P}(e_i)\|$$

for an absolute constant C provided the right hand side is smaller than 1. The latter is true if and only if

$$p \geq C^{-2} \log n \max_i \|\mathcal{P}(e_i)\|^2.$$

Now let U be an open neighborhood of x such that $\text{coh}(z) < (1 + \delta) \text{coh}(X)$ for all $z \in U$. Then, one can write

$$Z = \sup_{y_1, y_2 \in U'} \left\| \sum_{i=1}^n \left(\frac{\varepsilon_i}{p} - 1 \right) \cdot \langle y_1, \mathcal{P}(e_i) \rangle \langle y_2, \mathcal{P}(e_i) \rangle \right\|$$

with a countable subset $U' \subsetneq U$. By construction of U' , one has

$$\left\| \left(\frac{\varepsilon_i}{p} - 1 \right) \cdot \langle y_1, \mathcal{P}(e_i) \rangle \langle y_2, \mathcal{P}(e_i) \rangle \right\| \leq p^{-1}(1 + \delta) \text{coh}(X).$$

Applying Talagrand's Inequality in the form [5, Theorem 9.1], one obtains

$$P(\|Z - \mathbb{E}(Z)\| > t) \leq 3 \exp\left(-\frac{t}{KB} \log\left(1 + \frac{t}{2}\right)\right)$$

with an absolute constant K and $B = p^{-1}(1 + \delta) \text{coh}(X)$. Since δ was arbitrary, it follows that

$$P(\|Z - \mathbb{E}(Z)\| > t) < 3 \exp\left(-\frac{p \cdot t}{K \text{coh}(X)} \log\left(1 + \frac{t}{2}\right)\right).$$

Substituting $p = C \cdot \lambda' \cdot \text{coh}(X) \cdot \log n$, and proceeding as in the proof of Theorem 4.2 in [5] (while changing absolute constants), one arrives at the statement. \square

Proof. Without loss of generality we can assume that X is projective and thus compact (e.g., by using Chow's lemma). Thus, there exists $x \in X$ such that for the tangent space $T_x X$ at x it holds that $\text{coh}(T_x X) = \text{coh}(X)$. Now let $y = \Omega(x)$, note that y is a point-valued discrete random variable. By the equivalence of the statements (iv) and (v) in Lemma 2.7.12, it suffices to show that the operator

$$Z = \left\| p^{-1} \theta \circ d\Omega - \text{id} \right\|$$

is contractive, where θ is projection, from T_y onto T_x , with probability at least $1 - 3n^{-\lambda}$ under the assumptions on p . Let e_1, \dots, e_n be the orthonormal coordinate system we choose for \mathbb{C}^n , and \mathcal{P} the projection onto T_x . Then the projection $\theta \circ d\Omega$ has, when we consider T_x to be embedded into \mathbb{C}^n , the matrix representation

$$\sum_{i=1}^n \varepsilon_i \cdot \mathcal{P}(e_i) \otimes \mathcal{P}(e_i),$$

where ε_i are independent Bernoulli random variables with probability p for 1 and $(1 - p)$ for 0. Thus, in matrix representation,

$$Z = \left\| \sum_{i=1}^n \left(\frac{\varepsilon_i}{p} - 1 \right) \cdot \mathcal{P}(e_i) \otimes \mathcal{P}(e_i) \right\|.$$

By Rudelson's lemma 2.7.13, it follows that

$$\mathbb{E}(Z) \leq C \sqrt{\frac{\log n}{p}} \max_i \|\mathcal{P}(e_i)\|$$

for an absolute constant C provided the right hand side is smaller than 1. The latter is true if and only if

$$p \geq C^{-2} \log n \max_i \|\mathcal{P}(e_i)\|^2.$$

Now let $\delta > 0$, and let U be an open neighborhood of x such that $\text{coh}(T_y X) < (1 + \delta) \text{coh}(X)$. Then, one can write

$$Z = \sup_{y_1, y_2 \in U'} \left\| \sum_{i=1}^n \left(\frac{\varepsilon_i}{p} - 1 \right) \cdot \langle y_1, \mathcal{P}(e_i) \rangle \langle y_2, \mathcal{P}(e_i) \rangle \right\|$$

with a countable subset $U' \subsetneq U$. By construction of U' , one has

$$\left\| \left(\frac{\varepsilon_i}{p} - 1 \right) \cdot \langle y_1, \mathcal{P}(e_i) \rangle \langle y_2, \mathcal{P}(e_i) \rangle \right\| \leq p^{-1} (1 + \delta) \text{coh}(X).$$

Applying Talagrand's Theorem 9.1 from Candès and Recht [5], one obtains

$$P(\|Z - \mathbb{E}(Z)\| > t) \leq 3 \exp\left(-\frac{t}{KB} \log\left(1 + \frac{t}{2}\right)\right)$$

with an absolute constant K and $B = p^{-1}(1 + \delta) \text{coh}(X)$. Since δ was arbitrary, it follows that

$$P(\|Z - \mathbb{E}(Z)\| > t) < 3 \exp\left(-\frac{p \cdot t}{K \text{coh}(X)} \log\left(1 + \frac{t}{2}\right)\right).$$

Substituting $p = C \cdot \lambda' \cdot \text{coh}(X) \cdot \log n$, and proceeding as in the proof of Theorem 4.2 in Candès and Recht [5] (while changing absolute constants), one arrives at the statement. \square

Corollary 2.7.9. *Keep the notations of Theorem 2.7.8. If X is moreover maximally incoherent, and*

$$pn \geq C \cdot \lambda \cdot \dim(X) \cdot \log n, \quad \text{with } \lambda \geq 1,$$

then Ω is generically finite with probability at least

$$1 - 3n^{-\lambda}.$$

Proposition 2.7.10. $\mathcal{M}(m \times n, r)$ *is maximally incoherent.*

Proof. Let $A \in \mathbb{C}^{m \times n}$ be any matrix of rank r or less, with $A = UV^\top$ and $U \in \mathbb{C}^{m \times r}$, $V \in \mathbb{C}^{n \times r}$. Let H be the tangent space to $\mathcal{M}(m \times n, r)$ at A , and H_U resp. H_V the row-spans of U resp. V . The calculation leading to [5, equation 4.9] shows that

$$\text{coh}(H) = \text{coh}(H_U) + \text{coh}(H_V) - \text{coh}(H_U) \text{coh}(H_V).$$

Now for any pair of r -flats H_U and H_V in m -resp. n -space, there exists an A as above; on the other hand, Proposition 2.7.3 shows that there exist H_U, H_V such that

$$\text{coh}(H_U) = \frac{r}{m} \quad \text{and} \quad \text{coh}(H_V) = \frac{r}{n}.$$

Thus, substituting, this implies that there exists H with

$$\text{coh}(H) = \text{coh}(H_U) + \text{coh}(H_V) - \text{coh}(H_U)\text{coh}(H_V) = \frac{r \cdot (m + n - r)}{mn}.$$

Since $\text{coh}(\mathcal{M}(m \times n, r)) \leq \text{coh}(H)$ for any such H , this implies together with the lower bound from Proposition 2.7.6 that

$$\text{coh}(\mathcal{M}(m \times n, r)) = \frac{r \cdot (m + n - r)}{mn}.$$

□

Corollary 2.7.11. *Let M be an Erdős-Rényi random mask of size $(m \times n)$ and sampling probability p . There is an absolute constant C such that if*

$$p \geq C \cdot \lambda \cdot \frac{r \cdot (m + n - r)}{mn} \cdot (\log m + \log n), \quad \text{with } \lambda \geq 1,$$

then Ω is generically finite with probability at least

$$1 - 3(mn)^{-\lambda}.$$

Finally, we state the lemmata which were used in the proof of Theorem 2.7.8. The first lemma relates local injectivity to generic finiteness and contractivity of a linear map. It is related to Corollary 4.3 in Candès and Recht [5].

Lemma 2.7.12. *Let $\varphi : X \rightarrow Y$ be a surjective map of complex algebraic varieties, let $x \in X$, and $y = \varphi(x)$ be smooth points of X resp. Y . Let*

$$d\varphi : T_x X \rightarrow T_y Y$$

be the induced map of tangent spaces²⁵. Then, the following are equivalent:

- (i) *There is a complex open neighborhood $U \ni x$ such that the restriction $\varphi : U \rightarrow \varphi(U)$ is bijective.*
- (ii) *$d\varphi$ is bijective.*
- (iii) *There exists an invertible linear map $\theta : T_y Y \rightarrow T_x X$.*
- (iv) *There exists a linear map $\theta : T_y Y \rightarrow T_x X$ such that the linear map*

$$\theta \circ d\varphi - \text{id},$$

where id is the identity operator, is contractive²⁶.

²⁵ $T_x X$ is the tangent plane of X at x , which is identified with a vector space of formal differentials where x is interpreted at 0. Similarly, $T_y Y$ is identified with the formal differentials around y . The linear map $d\varphi$ is induced by considering $\varphi(x + dv) = y + dv'$ and setting $d\varphi(dv) = dv'$; one checks that this is a linear map since x, y are smooth. Furthermore, $T_x X$ and $T_y Y$ can be endowed with the Euclidean norm and scalar product it inherits from the tangent planes. Thus, $d\varphi$ is also a linear map of normed vector spaces which is always bounded and continuous, but not necessarily proper.

²⁶A linear operator \mathcal{A} is contractive if $\|\mathcal{A}(x)\| < 1$ for all x with $\|x\| < 1$.

If moreover X is irreducible, then the following is also equivalent:

(v) φ is generically finite.

Proof. (ii) is equivalent to the fact that the matrix representing $d\varphi$ is an invertible matrix. Thus, by the properties of the matrix inverse, (ii) is equivalent to (iii), and (ii) is equivalent to (i) by the constant rank theorem (e.g., 9.6 in Rudin [39]).

By the upper semicontinuity theorem (I.8, Corollary 3 in Mumford [31]), (i) is equivalent to (v) in the special case that X is irreducible, the reasoning is completely analogous to the proof of Theorem 2.3.5.

(ii) \Rightarrow (iv): Since $d\varphi$ is bijective, there exists a linear inverse $\theta : T_y Y \rightarrow T_x X$ such that $\theta \circ d\varphi = \text{id}$. Thus

$$\theta \circ d\varphi - \text{id} = 0$$

which is by definition a contractive linear map.

(iv) \Rightarrow (iii): We proceed by contradiction. Assume that no linear map $\theta : T_y Y \rightarrow T_x X$ is invertible. Since φ is surjective, $d\varphi$ also is, which implies that for each θ , the linear map $\theta \circ d\varphi$ is rank deficient. Thus, for every θ , there exists a non-zero $\alpha \in \text{Ker } \theta$. By linearity and surjectivity of $d\Omega$, there exists a non-zero $\beta \in T_x X$ with $d\Omega(\beta) = \alpha$. Without loss of generality we can assume that $\|\beta\| = 1$, else we multiply α and β by the same constant factor. By construction,

$$\|[\theta \circ d\varphi - \text{id}](\beta)\| = \|\theta(\alpha) - \beta\| = \|\beta\| = 1,$$

so θ cannot be contractive. Since θ was arbitrary, this proves that (iv) cannot hold if (iii) does not hold, which is equivalent to the claim. \square

The second lemma is a consequence of Rudelson's Lemma, see Rudelson [38], for Bernoulli samples.

Lemma 2.7.13. *Let y_1, \dots, y_M be vectors in \mathbb{R}^n , let $\varepsilon_1, \dots, \varepsilon_M$ be i.i.d. Bernoulli variables, taking value 1 with probability p and 0 with probability $(1 - p)$. Then,*

$$\mathbb{E} \left(\left\| 1 - \sum_{i=1}^M \left(\frac{\varepsilon_i}{p} \right) y_i \otimes y_i \right\| \right) \leq C \sqrt{\frac{\log n}{p}} \max_{1 \leq i \leq M} \|y_i\|$$

with an absolute constant C , provided the right hand side is 1 or smaller.

Proof. The statement is exactly Theorem 3.1 in Candès and Romberg [4], up to a renaming of variables, the proof can also be found there. It can also be directly obtained from Rudelson's original formulation in Rudelson [38] by setting substituting $\frac{\varepsilon_i}{\sqrt{p}} y_i$ in the above formulation for y_i in Rudelson's formulation and upper bounding the right hand side in Rudelson's estimate. \square

2.7.2 Finiteness of Random Projections

The results of section 2.7.1, in particular Theorem 2.7.8 might lead to the belief that the log-factor in the number of samples always, or almost always necessary for identifiability, in terms of the chosen projections. That, however is not true. While Theorem 2.7.8 gives a bound which is valid for *any* coordinate system and the coherence definition associated to it, the following theorem states that for a *general* system of coordinates, a much lower bound and a stricter statement is true:

Theorem 2.7.14. *Let $X \subseteq \mathbb{K}^n$ be an algebraic variety or a compact analytic variety, let $\Omega : \mathbb{K}^n \rightarrow \mathbb{K}^m$ a generic linear map. Let $x \in X$ be a smooth point. Then, $X \cap \Omega^{-1}(\Omega(x))$ is finite if and only if $k \geq \dim(X)$, and $X \cap \Omega^{-1}(\Omega(x)) = \{x\}$ if $m > \dim(X)$.*

Proof. The theorem follows from the the more general height-theorem-like statement that

$$\text{codim}(X \cap H) = \text{codim}(X) + \text{codim}(H) = \text{codim}(X) + n - k,$$

where H is a generic k -flat with $k \leq \text{codim}(X) + n + 1$, a proof of which can be found for example in the Appendix of [27]. Then, the first statement about generic finiteness follows by taking a generic $y \in \Omega(X)$ and observing that $\Omega^{-1}(y) = H \cap X$ where H is generic if $k \leq \dim(X)$. That implies in particular that if $k = \dim(X)$, then the fiber $\Omega^{-1}(\Omega(x))$ for a generic $x \in X$ consists of finitely many points, which can be separated by an additional generic projection, thus the statement follows. \square

Theorem 2.7.14 can be interpreted in two ways. On one hand, it means that any point on X can be reconstructed from exactly $\dim(X)$ random linear projections. On the other hand, it means that if the chosen coordinate system in which X lives is random, then $\dim(X)$ measurements suffice for (finite) identifiability of the map - no more structural information is needed. In view of Theorem 2.7.8, this implies that the log-factor and the probabilistic phenomena in identifiability occur when the chosen coordinate system is degenerate with respect to the variety X in the sense that it is intrinsically aligned.

3. Algorithms

3.1 Randomized Algorithms for Completability

In the following, we describe some algorithms which can be derived from the theory of differentials and matroids in section 2.5. The algorithms in this section answer the question which entries of a rank r matrix can be in principle computed from the given ones. As Theorems 2.3.5 and 2.4.4 show, for a generic matrix this depends only on the position of the known resp. measured entries, encoded in the so-called mask, and not on the values of the entries. Calculations in the vector space of evaluated differentials then allow to simply determine the entries which can be reconstructed up to finite choice.

First, with Algorithm 1, we present a randomized algorithm which checks whether all missing entries can be reconstructed up to finite choice; in Singer and Cucuringu [43], a very similar algorithm was already conjectured to be correct. In step 1, a $(m \times r)$ -matrix U and a $(n \times r)$ -matrix V are sampled from a generic probability distribution. One can show (see for example the genericity section in the appendix of Király et al. [27]) continuous probability distribution, e.g., any

Gaussian distribution on the matrices, will be generic and fulfill the properties of Definition 2.2.1. Thus, $U \cdot V^T$ will be a generic $(m \times n)$ -matrix of rank r . In step 2, the differentials of A for all its known entries are calculated. The differentials da_{ij} are contained in the formal \mathbb{C} -vector generated by all dU_{ij} and dV_{ij} ; thus they can be conveniently represented as a $(rm + rn)$ -vector, or a $(r \times m + n)$ -matrix. In step 3, their span is calculated. By checking their span in step 4 and its dimension, e.g., numerically, one can decide whether M was finitely completable, which follows from the equivalence of Theorem 2.5.14 (i) and (v), and Corollary 2.5.30, proving correctness of Algorithm 1.

Algorithm 1 Finite completability in rank r .

Input: An $(m \times n)$ mask M . *Output:* Whether M is finitely completable in rank r .

- 1: Randomly sample $U \in \mathbb{C}^{m \times r}, V \in \mathbb{C}^{n \times r}$.
- 2: For all $(ij) \in E(M)$, calculate

$$da_{ij} := \sum_{k=1}^r (V_{jk} dU_{ik} + U_{ik} dV_{jk})$$

where the dU_{ij}, dV_{ij} are to be considered as formal basis vectors of a $(rm + rn)$ -dimensional vector space.

- 3: Set $V_M = \text{span}\{da_{ij} ; (ij) \in E(M)\}$.
 - 4: If $\dim(V_M)_{\mathbb{C}} = r \cdot (m + n - r)$, return “finitely completable”, else “not finitely completable”.
-

Algorithm 2 Matroid rank.

Input: An $(m \times n)$ mask M . *Output:* The matroid rank of M in rank r .

- 1: Randomly sample $U \in \mathbb{C}^{m \times r}, V \in \mathbb{C}^{n \times r}$.
- 2: For all $(ij) \in E(M)$, calculate

$$da_{ij} := \sum_{k=1}^r (V_{jk} dU_{ik} + U_{ik} dV_{jk})$$

where the dU_{ij}, dV_{ij} are to be considered as formal basis vectors of a $(rm + rn)$ -dimensional vector space.

- 3: Set $V_M = \text{span}\{da_{ij} ; (ij) \in E(M)\}$.
 - 4: Return $\text{rk}_r(M) = \dim_{\mathbb{C}} V_M$.
-

Similar principles can be used to calculate the number of degrees of freedom contained in a set of given entries of a matrix. Theorem 2.3.5 again shows that, generically, it does only depend on the position of the entries. In Algorithm 2 we perform the same steps as in Algorithm 1, up to step 4, where we give back the (numerical) dimension of the span instead checking whether it equates to $r \cdot (m + n - r)$. Indeed, Algorithm 1 can be obtained as the algorithm which just checks whether $\text{rk}_r(M) = r \cdot (m + n - r)$. The correctness of Algorithm 2 follows from the equivalence of Theorem 2.5.14 (i) and (v), and Proposition 2.5.22. Corollary 2.5.30 exhibits the relation to Algorithm 1. Also note that since many objects in matroid theory like circuits, independence, bases, etc., can be characterized by an evaluation of the rank - compare Proposition 2.5.21 - Algorithm 2 can be in fact used to classify or find such objects and determine their properties.

For example, Algorithm 2 can be used in classical matroid theoretical algorithms to find or count circuits, bases, or for determining the structure of the whole determinantal matroid.

Algorithm 3 Completable closure.

Input: An $(m \times n)$ mask M . *Output:* The completable closure N of M in rank r .

- 1: Randomly sample $U \in \mathbb{C}^{m \times r}, V \in \mathbb{C}^{n \times r}$.
- 2: For all (ij) , calculate

$$da_{ij} := \sum_{k=1}^r (V_{jk} dU_{ik} + U_{ik} dV_{jk})$$

where the dU_{ij}, dV_{ij} are to be considered as formal basis vectors of a $(rm + rn)$ -dimensional vector space.

- 3: Set $V_M = \text{span}\{da_{ij}; (ij) \in E(M)\}$.
 - 4: For each (ij) , calculate whether $da_{ij} \in V_M$.
 - 5: Define N as $E(N) = \{(ij); da_{ij} \in V_M\}$.
-

The randomized strategy, in its most general setting, allows to compute the set of all entries which are in principle reconstructible, up to finite choice, from the known entries; Algorithm 3 can be used to do that. Steps 1 to 3 are analogous as in Algorithms 1 and 2, with the small distinction that in step 2, all differentials are computed, since they correspond to the entries, and one has to check for all entries whether they can be reconstructed or not. In step 2, one numerically then checks for each entry whether its differential is in the span of those in the mask or not; an entry is reconstructible if and only if it is contained in the span, which follows from the equivalence of Theorem 2.5.14 (i) and (v). Algorithm 1 can be seen as a special and simplified case of Algorithm 3 where it is only checked whether the completable entries are all entries, or not. Note also that Algorithm 3 can be applied to determine whether a partial mask (N, M) is finitely completable, by checking whether $G(N)$ lies in the completable closure of $G(M)$ (i.e., whether the completable closure of M minus N is non-negative).

3.2 Algorithms for checking Closability

An algorithm for one step r -closure is shown in Algorithm 4. Roughly speaking we look through each missing edge $(i, j) \in V \times W \setminus E$ and check whether it can be closed by known edges. This can be done by finding neighbors $J = N(i)$ of i in W and neighbors $I = N(j)$ of j in V , and checking if the subgraph $(I, J, I \times J \cap E)$ contains an $r \times r$ bi-clique. This is shown in Algorithm 6. Since an $r \times r$ clique cannot contain any vertex with degree less than r , we prune these vertices beforehand; this is shown in Algorithm 5.

One can decide if a bipartite graph (V, W, E) is r -closable or not by repeatedly applying Algorithm 4 and checking if the graph is a complete bipartite graph when there is no more edge to add; this is shown in Algorithm 7.

3.3 Algebraic Reconstruction of Matrices

The algorithm described in the previous subsection can be used to actually perform matrix completion. Each entry in the associative array $C : (i, j) \rightarrow (I, J)$ provides a valid $(r + 1) \times (r + 1)$

Algorithm 4 CloseOneStep($(V, W, E), r$)

Inputs: bipartite graph (V, W, E) , rank r .

Output: associative array $C: (i, j) \mapsto (I, J)$ where $(i, j) \in (V \times W) \setminus E$ and $I \subseteq V, J \subseteq W$ such that $(i', j') \in E$ for all $(i', j') \in (I \cup \{i\}) \times (J \cup \{j\})$ except $(i', j') = (i, j)$.

for each missing edge (i, j) in $V \times W \setminus E$ **do**

 Let $I \leftarrow N(j), J \leftarrow N(i)$, and $E' \leftarrow I \times J \cap E$.

$(I, J, E') \leftarrow \text{PruneNodesWithDegreeLessThan}((I, J, E'), r)$.

if $|I| < r$ or $|J| < r$ **then**

 Continue.

end if

$(I', J') \leftarrow \text{FindAClique}((I, J, E'), r, r)$.

if $|I'| > 0$ and $|J'| > 0$ **then**

$C(i, j) \leftarrow (I', J')$.

end if

end for

Return C .

Algorithm 5 PruneNodesWithDegreeLessThan($(V, W, E), d_1, d_2$)

Inputs: bipartite graph (V, W, E) , minimum degree d_1 for the row vertices and d_2 for the column vertices.

Output: pruned bipartite graph (V, W, E) .

while true **do**

$V' \leftarrow \{v \in V : |N(v)| < d_1\}$.

$W' \leftarrow \{w \in W : |N(w)| < d_2\}$.

if $V' = \emptyset$ and $W' = \emptyset$ **then**

 Return $(V, W, (V \times W) \cap E)$.

end if

$V \leftarrow V \setminus V'$.

$W \leftarrow W \setminus W'$.

end while

Algorithm 6 FindAClique($(V, W, E), d_1, d_2$)

Inputs: bipartite graph (V, W, E) , size of the bipartite clique to be found $d_1 \times d_2$.

Output: vertex sets of a clique (I, J) .

if $|V| < d_1$ or $|W| < d_2$ **then**

 Return (\emptyset, \emptyset) .

end if

$(V, W, E) \leftarrow \text{PruneNodesWithDegreeLessThan}((V, W, E), d_2, d_1)$.

for each $v \in V$ **do**

if $d_1 = 1$ and $|N(v)| \geq d_2$ **then**

 Return $(\{v\}, N(v))$.

end if

$V \leftarrow V \setminus \{v\}$, $W' \leftarrow N(v)$, $E' \leftarrow (V \times W') \cap E$.

$(I', J') \leftarrow \text{FindAClique}((V, W', E'), d_1 - 1, d_2)$.

if $|I'| > 0$ and $|J'| > 0$ **then**

 Return $(I' \cup \{v\}, J')$.

end if

end for

Return (\emptyset, \emptyset) .

Algorithm 7 IsClosable($(V, W, E), r$)

Inputs: bipartite graph (V, W, E) , rank r .

Output: binary (true means closable, and false means not closable).

repeat

$n_{\text{non-zero}} \leftarrow |E|$.

$C \leftarrow \text{CloseOneStep}((V, W, E), r)$.

$E \leftarrow E \cup \text{keys}(C)$.

until $n_{\text{non-zero}} = |E|$

Return *true* if $E = I \times J$ else *false*.

vanishing minor condition, which we can exploit to fill one missing entry. Our algorithm is implemented in a breadth-first manner to minimize propagation of numerical errors. The details are described in Algorithm 8.

Algorithm 8 CompletionByClosure($A, (V, W, E), r$)

Inputs: partially observed matrix A , bipartite graph (V, W, E) , rank r .

Output: completed matrix A , list of associative arrays C_{save} .

$C_{save} = []$.

repeat

$n_{non-zero} \leftarrow |E|$.

$C \leftarrow \text{CloseOneStep}((V, W, E), r)$.

$C_{save} \leftarrow [C_{save}, C]$

for each $(i, j) \in \text{keys}(C)$ **do**

$E \leftarrow E \cup \{(i, j)\}$.

$A(i, j) \leftarrow A(i, J)A(I, J)^+A(I, j)$ where $(I, J) = C(i, j)$.

end for

until $n_{non-zero} = |E|$

Return (A, C_{save}) .

4. Experiments

4.1 Randomized Algorithms for Completability

In this section we will investigate the set of entries which is completable from a set of given entries. In section 2.4 we have seen that the completable set of entries does not depend on the value of the entries but only on their position, and Algorithm 3 provides a method to do so.

In order to illustrate the input and the output of Algorithm 3, we first give an example pair of input and output.

Example 4.1.1. *The example input for Algorithm 3 consists of the mask M , which has ones at the position of the known entries, the output is the mask N which has ones at the positions of the entries which can be reconstructed up to finite choice. The rank is set to $r = 2$. For the input*

$$M = \begin{pmatrix} 0 & 1 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 1 & 1 & 0 & 1 & 1 \\ 1 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 & 1 \\ 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 1 & 1 & 1 & 0 & 0 & 1 & 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 1 & 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \end{pmatrix}$$

Algorithm 3 computes the output

$$M = \begin{pmatrix} 1 & 1 & 1 & 0 & 0 & 1 & 1 & 1 & 1 & 1 \\ 0 & 1 & 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 & 1 & 1 & 0 & 1 & 1 \\ 1 & 1 & 1 & 0 & 1 & 1 & 1 & 0 & 1 & 1 \\ 1 & 1 & 1 & 0 & 0 & 1 & 1 & 0 & 1 & 1 \\ 1 & 1 & 1 & 0 & 0 & 1 & 1 & 0 & 1 & 1 \\ 1 & 1 & 1 & 0 & 0 & 1 & 1 & 0 & 1 & 1 \\ 1 & 1 & 1 & 0 & 0 & 1 & 1 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 & 1 & 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \end{pmatrix}$$

For a quantitative analysis, we perform experiments to investigate how the expected number of completable entries is influenced by the number of known entries. In particular, section 2.6 suggests that a phase transition between the state where only very few additional entries can be reconstructed and the state where a large set of entries can be reconstructed should take place at some point. Figure 3 shows that this is indeed the case when slowly increasing the number of known entries: first, the set of reconstructible entries is roughly equal to the set of known entries, but then, a sudden phase transition occurs and the set of reconstructible entries quickly reaches the set of all entries.

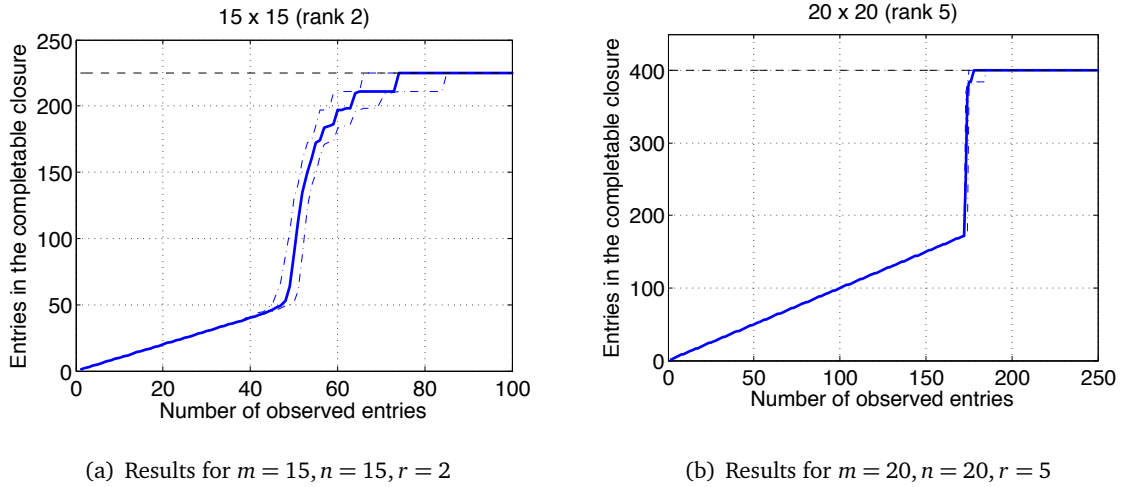


Figure 3: Expected number of completable entries (in rank r) versus the number of known entries where the positions of the known entries are uniformly randomly sampled in an $(m \times n)$ -matrix. The expected number of completable entries was estimated for each data points from repeated calculations of the completable closure (200 for $r = 2$, and 20 for $r = 5$). The blue solid line is the median, the blue dotted lines are the two other quartiles. The black dotted line is the total number of entries, $m \cdot n$.




Figure 4: Phase transition curves of various conditions for 100×100 matrices at rank 3.

4.2 Phase Transitions

Figure 4 shows phase transition curves of various conditions for 100×100 matrices at rank 3. We consider uniform sampling model here. More specifically, we generated random 100×100 masks with various number of observed entries by first randomly sampling the order of edges (using MATLAB `randperm` function) and sequentially adding 100 entries at a time from 100 to 6000. Therefore, we made sure to preserve the monotonicity of the properties considered here. This experiment was repeated 100 times and averaged to obtain estimates of success probabilities. The conditions plotted are (a) minimum degree at least r , (b) r -connected, (c) finitely completable at rank r , (d) r -closable, (e) nuclear norm successful, and (f) one step r -closable. We solved the following minimization problem

$$\begin{aligned} \min_X \quad & \|X\|_*, \\ \text{subject to} \quad & X(i, j) = A(i, j) \quad \forall (i, j) \in E, \end{aligned}$$

where $\|X\|_* = \sum_{j=1}^r \sigma_j(X)$ is the nuclear norm of X . The success of nuclear norm minimization is defined as the relative error $\|\hat{X} - A\|_2 / \|A\|_2$ less than 0.01, where \hat{X} is the minimizer of the above minimization problem.

The success probabilities of the (a) minimum degree, (b) r -connected, and (c) finitely completable are almost on top of each other, and exceeds chance (probability 0.5) around $|E| \simeq 1,000$. The success probability of the (d) r -closable curve passes through 0.5 around $|E| \simeq 1,300$. Therefore the optimality gap of the r -closure method is small. On the other hand, the nuclear norm minimization required about 2,200 entries to succeed with probability larger than 0.5.

Figure 5 shows the same plot as above for 100×100 matrices at rank 6. The success probabilities of the (a) minimum degree, (b) r -connected, (c) finitely completable are again almost the same, and succeeds chance probability 0.5 around $|E| \simeq 1,400$. On the other hand, the number of entries required for r -closability is at least 3,700, whereas that required for the nuclear norm minimization to succeed is only 3,100.

Figure 6 shows the phase transition from a non-completable mask to a finitely completable mask for almost $2r$ -regular random masks. Here we first randomly sampled $n \times n$ $2r$ -regular

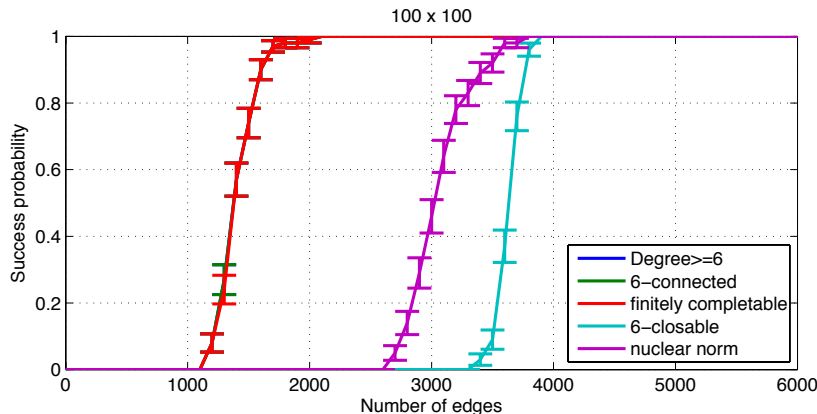


Figure 5: Phase transition curves of various conditions for 100×100 matrices at rank 6.

masks using Steger & Wormald algorithm [46]. Next we randomly permuted the edges included in the mask and the edges not included in the mask independently and concatenated them into a single list of edges. In this way, we obtained a length mn ordered list of edges that become $2r$ -regular exactly at the $2rn$ th edge. For each ordered list sampled this way, we took the first $2rn - i$ edges and checked whether the mask corresponding to these edges was finitely completable or not for $i = -15, -14, \dots, 5$. This procedure was repeated 100 times and averaged to obtain a probability estimate. In order to make sure that the phase transition is indeed caused by the regularity of the mask, we conducted the same experiment with row-wise $2r$ -regular masks, i.e., each row of the mask contained exactly $2r$ entries while the number of non-zero entries varied from a column to another.

In Figure 6, the phase transition curves for different n at rank 2 and 3 are shown. The two plots in the top part show the results for the $2r$ -regular masks, and the two plots in the bottom show the same results for the $2r$ -row-wise regular masks. For the $2r$ -regular masks, the success probability of finite completability sharply rises when the number of edges exceeds $2rn - r^2$ ($i = -4$ for $r = 2$ and $i = -9$ for $r = 3$); the phase transition is already rather sharp for $n = 10$ and for $n \geq 20$ it becomes almost zero or one. On the other hand, the success probabilities for the $2r$ -row-wise regular masks grow rather slowly and approach zero for large n . This is natural, since it is likely for large n that there is some column with non-zero entries less than r , which violates the necessary condition in Proposition 2.3.31.

5. Conclusion

In this paper we have shown that low-rank matrix completion is a task with both algebraic and combinatorial structure. We have shown that this structure can be made use of both in the theoretical analysis and the construction of algorithms for matrix completion. We thus reason that using the inherent algebraic structure of a Machine Learning problem is beneficial and thus preferable to structure agnostic methodology.

For the problem of matrix completion, we have also shown that its behavior depends crucially on the sampling process while only marginally on the generative truth. That is, given some

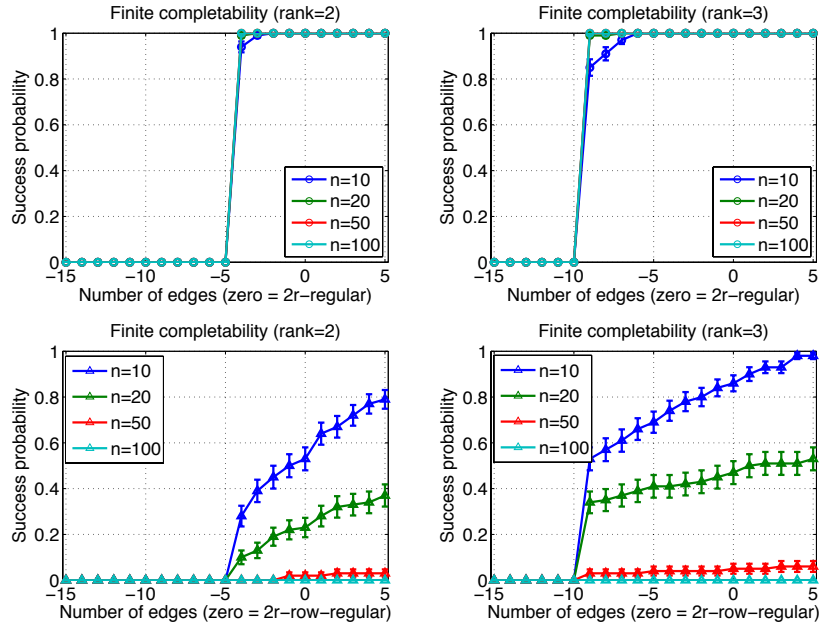


Figure 6: Phase transition in an almost regular mask.

entries of a low-rank matrix, the set of entries which can be reconstructed from the entries and the condition that the matrix has low-rank, depends only on the position of the known entries, and not on their particular values. Similarly, the properties of the reconstruction process can be made independent from the values of the entries. We argue that this is more natural than assuming the converse, i.e., that the intrinsic properties of the task are determined by the generative truth, since the generative truth may change while the problem itself and thus its intrinsic properties should not. Indeed, our results on algebraic compressed sensing, which generalize our findings on matrix completion, seem to imply that this is in fact a general principle in compressed sensing: the properties of the sampling, for example reconstructibility, necessary and sufficient sampling densities, etc., should be independent of particular signal, and only dependent on sampling and compression properties.

We have also presented several combinatorial objects which can be used to study the possible set of completable entries in an incomplete low-rank matrix. Namely, one can associate bipartite graphs to patterns of entries and study the degrees of freedom, or the set of completable entries by analyzing properties of the associated matroid. Moreover, the asymptotics of the necessary number of entries for reconstruction can be now studied via the asymptotics of the bipartite graphs and the matroids. The theory of formal differentials can be used to design algorithms for calculating the combinatorial objects and their implication for reconstructability.

The algorithms presented in this paper do not only allow the theoretical studies of the phase transitions involved in the matrix completion of large matrices, but also give efficient tools to the hand of the practitioner to determine which entries of a matrix can be completed or not, i.e., which reconstructed entries can be trusted or not, and methods, using the algebraic combinatorial structure, to calculate the reconstruction itself.

We conjecture that the methods and principles presented in this paper can also be applied to

a wider class of problems with algebraic-combinatorial structure, in particular

- **Matrix completion with different constraints.** Completing low-rank matrices which are symmetric, antisymmetric, Hermitian, real, or endowed with other combinatorial or algebraic constraints can be studied by using analogous methods. Also, the task of completing matrices which instead of the low-rank constraint fulfill different algebraic boundary conditions, e.g., different types of rigidity, or sparsity, or hybrid properties, can be recast in our framework.
- **Tensor methods.** Similar to matrices, the low-rank tensors can be expressed as being contained in an algebraic manifold. Projections of tensors can be treated in a similar way to matrices, and the algebraic-combinatorial structures generalize, involving multigraphs and their structured matroids.
- **Algebraic compressed sensing.** If the signal is parameterized by a finite dimensional set of parameters, and the compression constraints can be given by polynomial equations, many of our methods which are not specific for low-Rank matrix completion apply, in particular our theory for the study of the sampling properties independently from signal properties.

Concluding, we argue that the additional use of algebraic or combinatoric structure in a Machine Learning problem can only be beneficial compared to not using it. Since algebraic geometry, combinatorial algebra and discrete mathematics are the proper tools to analyze and utilize such kind of structure, we claim that machine learning, as well as the mentioned fields, can only profit from a more widespread interdisciplinary collaboration with and between each other.

References

- [1] Evrim Acar, Daniel M. Dunlavy, and Tamara G. Kolda. Link prediction on evolving data using matrix and tensor factorizations. In *Data Mining Workshops, 2009. ICDMW'09. IEEE International Conference on*, pages 262–269. IEEE, 2009.
- [2] Andreas Argyriou, Craig A. Micchelli, Massimiliano Pontil, and Yi Ying. A spectral regularization framework for multi-task structure learning. In J.C. Platt, D. Koller, Y. Singer, and S. Roweis, editors, *Advances in NIPS 20*, pages 25–32. MIT Press, Cambridge, MA, 2008.
- [3] Béla Bollobás. *Random graphs*, volume 73 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, second edition, 2001. ISBN 0-521-80920-7; 0-521-79722-5. doi: 10.1017/CBO9780511814068. URL <http://dx.doi.org/10.1017/CBO9780511814068>.
- [4] Emmanuel Candès and Justin Romberg. Sparsity and incoherence in compressive sampling. *Inverse Problems*, 23(3):969–985, 2007. ISSN 0266-5611. doi: 10.1088/0266-5611/23/3/008. URL <http://dx.doi.org/10.1088/0266-5611/23/3/008>.
- [5] Emmanuel J. Candès and Benjamin Recht. Exact matrix completion via convex optimization. *Found. Comput. Math.*, 9(6):717–772, 2009. ISSN 1615-3375. doi: 10.1007/s10208-009-9045-5. URL <http://dx.doi.org/10.1007/s10208-009-9045-5>.
- [6] Emmanuel J. Candès and Terence Tao. The power of convex relaxation: near-optimal matrix completion. *IEEE Trans. Inform. Theory*, 56(5):2053–2080, 2010. ISSN 0018-9448. doi: 10.1109/TIT.2010.2044061. URL <http://dx.doi.org/10.1109/TIT.2010.2044061>.
- [7] Peter G. Casazza and Manuel T. Leon. Existence and construction of finite tight frames. *J. Concr. Appl. Math.*, 4(3):277–289, 2006.
- [8] Robert Connelly. Generic global rigidity. *Discrete Comput. Geom.*, 33(4):549–563, 2005. ISSN 0179-5376. doi: 10.1007/s00454-004-1124-4. URL <http://dx.doi.org/10.1007/s00454-004-1124-4>.
- [9] David A. Cox, John Little, and Donal O’Shea. *Ideals, Varieties, and Algorithms: An Introduction to Computational Algebraic Geometry and Commutative Algebra, 3/e (Undergraduate Texts in Mathematics)*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2007. ISBN 0387356509.
- [10] Michael Dobbins. Realizability of polytopes as a low rank matrix completion problem. Preprint, arXiv:1012.3905, 2011. URL <http://arxiv.org/abs/1012.3905>.
- [11] David Eisenbud. *Commutative Algebra with a View towards Algebraic Geometry*. Springer-Verlag New York, Inc., 1995.
- [12] Rina Foygel and Nathan Srebro. Concentration-based guarantees for low-rank matrix reconstruction. *Arxiv preprint arXiv:1102.3923*, 2011.
- [13] Andrew Goldberg, Xiaojin Zhu, Benjamin Recht, Jun-Ming Xu, and Robert Nowak. Transduction with matrix completion: Three birds with one stone. In J. Lafferty, C. K. I. Williams,

- J. Shawe-Taylor, R.S. Zemel, and A. Culotta, editors, *Advances in Neural Information Processing Systems 23*, pages 757–765. 2010.
- [14] Steven J. Gortler, Alexander D. Healy, and Dylan P. Thurston. Characterizing generic global rigidity. *Amer. J. Math.*, 132(4):897–939, 2010. ISSN 0002-9327. doi: 10.1353/ajm.0.0132. URL <http://dx.doi.org/10.1353/ajm.0.0132>.
- [15] Jack Graver, Brigitte Servatius, and Herman Servatius. *Combinatorial rigidity*, volume 2 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 1993. ISBN 0-8218-3801-6.
- [16] David Gross. Recovering low-rank matrices from few coefficients in any basis, 2010. Eprint arXiv:0910.1879.
- [17] Kirk Haller, Audrey Lee-St.John, Meera Sitharam, Ileana Streinu, and Neil White. Body-and-cad geometric constraint systems. *Comput. Geom. Theory Appl.*, 45(8):385–405, October 2012. ISSN 0925-7721. doi: 10.1016/j.comgeo.2010.06.003. URL <http://dx.doi.org/10.1016/j.comgeo.2010.06.003>.
- [18] Robin Hartshorne. *Algebraic Geometry*. Springer-Verlag New York, Inc., 1977.
- [19] Bill Jackson and Tibor Jordán. Connected rigidity matroids and unique realizations of graphs. *J. Combin. Theory Ser. B*, 94(1):1–29, 2005. ISSN 0095-8956. doi: 10.1016/j.jctb.2004.11.002. URL <http://dx.doi.org/10.1016/j.jctb.2004.11.002>.
- [20] Bill Jackson, Brigitte Servatius, and Herman Servatius. The 2-dimensional rigidity of certain families of graphs. *J. Graph Theory*, 54(2):154–166, 2007. ISSN 0364-9024. doi: 10.1002/jgt.20196. URL <http://dx.doi.org/10.1002/jgt.20196>.
- [21] Svante Janson and Malwina J. Luczak. A simple solution to the k -core problem. *Random Structures Algorithms*, 30(1-2):50–62, 2007. ISSN 1042-9832. doi: 10.1002/rsa.20147. URL <http://dx.doi.org/10.1002/rsa.20147>.
- [22] Svante Janson, Tomasz Łuczak, and Andrzej Ruciński. *Random graphs*. Wiley-Interscience Series in Discrete Mathematics and Optimization. Wiley-Interscience, New York, 2000. ISBN 0-471-17541-2. doi: 10.1002/9781118032718. URL <http://dx.doi.org/10.1002/9781118032718>.
- [23] Shiva Prasad Kasiviswanathan, Cristopher Moore, and Louis Theran. The rigidity transition in random graphs. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1237–1252, Philadelphia, PA, 2011. SIAM.
- [24] Raghunandan H. Keshavan, Andrea Montanari, and Sewoong Oh. Matrix completion from a few entries. *IEEE Trans. Inform. Theory*, 56(6):2980–2998, 2010. ISSN 0018-9448. doi: 10.1109/TIT.2010.2046205. URL <http://dx.doi.org/10.1109/TIT.2010.2046205>.
- [25] Franz J. Király and Ryota Tomioka. A combinatorial algebraic approach for the identifiability of low-rank matrix completion. ICML 2012.

- [26] Franz J. Király and Ryota Tomioka. A combinatorial algebraic approach for the identifiability of low-rank matrix completion, 2012. URL http://www.ml.tu-berlin.de/fileadmin/fg159/kiraly/icml2012_merged.pdf.
- [27] Franz J. Király, Paul von Büna, Frank Meinecke, Duncan Blythe, and Klaus-Robert Müller. Algebraic geometric comparison of probability distributions. *Journal of Machine Learning Research*, 13(Mar):855–903, 2012.
- [28] Justin Malestein and Louis Theran. Generic combinatorial rigidity of periodic frameworks. *Advances in Mathematics*, 233(1):291 – 331, 2013. ISSN 0001-8708. doi: 10.1016/j.aim.2012.10.007. URL <http://dx.doi.org/10.1016/j.aim.2012.10.007>.
- [29] A. Menon and C. Elkan. Link prediction via matrix factorization. *Machine Learning and Knowledge Discovery in Databases*, pages 437–452, 2011.
- [30] James S. Milne. *Étale cohomology*, volume 33 of *Princeton Mathematical Series*. Princeton University Press, Princeton, N.J., 1980.
- [31] David Mumford. *The Red Book of Varieties and Schemes*. Lecture Notes in Mathematics. Springer-Verlag Berlin Heidelberg, 1999. ISBN 054063293X.
- [32] C. St. J. A. Nash-Williams. Edge-disjoint spanning trees of finite graphs. *J. London Math. Soc.*, 36:445–450, 1961. ISSN 0024-6107.
- [33] Sahand Negahban and Martin J. Wainwright. Restricted strong convexity and weighted matrix completion: Optimal bounds with noise. *J. Mach. Learn. Res.*, 13:1665–1697, 2012.
- [34] Sewoong Oh. *Matrix completion: fundamental limits and efficient algorithms*. PhD thesis, Stanford University, Dept. of Electrical Engineering, 2011. URL <http://purl.stanford.edu/by297nr4588>.
- [35] James Oxley. *Matroid theory*, volume 21 of *Oxford Graduate Texts in Mathematics*. Oxford University Press, Oxford, second edition, 2011. ISBN 978-0-19-960339-8.
- [36] Boris Pittel, Joel Spencer, and Nicholas Wormald. Sudden emergence of a giant k -core in a random graph. *J. Combin. Theory Ser. B*, 67(1):111–151, 1996. ISSN 0095-8956. doi: 10.1006/jctb.1996.0036. URL <http://dx.doi.org/10.1006/jctb.1996.0036>.
- [37] Jürgen Richter-Gebert. *Realization spaces of polytopes*, volume 1643 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 1996. ISBN 3-540-62084-2.
- [38] Mark Rudelson. Random vectors in the isotropic position. *J. Funct. Anal.*, 164(1):60–72, 1999. ISSN 0022-1236. doi: 10.1006/jfan.1998.3384. URL <http://dx.doi.org/10.1006/jfan.1998.3384>.
- [39] Walter Rudin. *Principles of mathematical analysis*. McGraw-Hill Book Co., New York, third edition, 1976. International Series in Pure and Applied Mathematics.
- [40] Ruslan Salakhutdinov and Nathan Srebro. Collaborative filtering in a non-uniform world: Learning with the weighted trace norm. In J. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R.S. Zemel, and A. Culotta, editors, *Advances in Neural Information Processing Systems 23*, pages 2056–2064. 2010.

- [41] Bernd Schulze and Walter Whiteley. The orbit rigidity matrix of a symmetric framework. *Discrete Comput. Geom.*, 46(3):561–598, 2011. ISSN 0179-5376. doi: 10.1007/s00454-010-9317-5. URL <http://dx.doi.org/10.1007/s00454-010-9317-5>.
- [42] Peter W. Shor. Stretchability of pseudolines is NP-hard. In *Applied geometry and discrete mathematics*, volume 4 of *DIMACS Ser. Discrete Math. Theoret. Comput. Sci.*, pages 531–554. Amer. Math. Soc., Providence, RI, 1991.
- [43] Amit Singer and Mihai Cucuringu. Uniqueness of low-rank matrix completion by rigidity theory. *Siam J Matrix Anal Appl*, 31(4):1621–1641, 2010.
- [44] Nathan Srebro and Adi Shraibman. Rank, trace-norm and max-norm. In *Proc. of the 18th Annual Conference on Learning Theory (COLT)*, pages 545–560. Springer, 2005.
- [45] Nathan Srebro, Jason D. M. Rennie, and Tommi S. Jaakkola. Maximum-margin matrix factorization. In Lawrence K. Saul, Yair Weiss, and Léon Bottou, editors, *Advances in NIPS 17*, pages 1329–1336. MIT Press, Cambridge, MA, 2005.
- [46] Angelika Steger and Nicholas Wormald. Generating random regular graphs quickly. *Combinatorics Probability and Computing*, 8(4):377–396, 1999.
- [47] Ileana Streinu and Louis Theran. Slider-pinning rigidity: a Maxwell-Laman-type theorem. *Discrete Comput. Geom.*, 44(4):812–837, 2010. ISSN 0179-5376. doi: 10.1007/s00454-010-9283-y. URL <http://dx.doi.org/10.1007/s00454-010-9283-y>.
- [48] Tiong-Seng Tay. Rigidity of multigraphs. I. Linking rigid bodies in n -space. *J. Combin. Theory Ser. B*, 36(1):95–112, 1984. ISSN 0095-8956. doi: 10.1016/0095-8956(84)90016-9. URL [http://dx.doi.org/10.1016/0095-8956\(84\)90016-9](http://dx.doi.org/10.1016/0095-8956(84)90016-9).
- [49] William T. Tutte. On the problem of decomposing a graph into n connected factors. *J. London Math. Soc.*, 36:221–230, 1961. ISSN 0024-6107.
- [50] Walter Whiteley. The union of matroids and the rigidity of frameworks. *SIAM J. Discrete Math.*, 1(2):237–255, 1988. ISSN 0895-4801. doi: 10.1137/0401025. URL <http://dx.doi.org/10.1137/0401025>.

A. Sheaves of Matroids

Several results in section 2.5.3 can be seen as an instance of a single principle in Algebraic Geometry, where one considers the dependence structure of sections on a scheme, e.g., a complex algebraic variety. The results will imply that the dependencies between the sections exhibit a generic behavior, which implies constant behavior on Zariski open sets of the variety.

The basic idea of associating a generic matroid to a real or complex variety via the formal differential is well-known (e.g, [47] contains an explicit discussion) in the combinatorial rigidity community and has been derived in a number of concrete cases, for example [17, 28, 41, 47, 48, 50], among others. Thus, the form of Theorem A.10 will not surprise experts. However, we are unaware of the general statement appearing in print. Moreover, as the presentation here makes clear, the *rank* natural matroidal construction on a fixed global section of the matroid sheaf is a generic invariant of smooth points, the *exact dependence structure* requires more genericity, a fact which is used implicitly when defining rigidity matroids as is done in [47] or when discussing combinatorial aspects of generic global rigidity [8, 14, 19].

Definition A.1. *Let R be an integral domain with field of fractions K , let M be an R -module, let $\Lambda \subseteq M$ be a finite multiset (a set where we allow finite repetitions). To Λ , we associate a matroid $\mathcal{M}[\Lambda]$ in the following way: The matroid is defined over the power set $\mathcal{P}(\Lambda)$ of Λ , and the independent sets are exactly the subsets $J \subseteq \Lambda$ such that*

$$\dim_K(K \cdot J) = \#J$$

where $K \cdot J$ denotes the K -submodule of $K \otimes_R M$ generated by J . By convention, $K \cdot \emptyset$ is the trivial K -module, i.e., the zero module.

That $\mathcal{M}[\Lambda]$ is in fact a matroid follows since $K \cdot J$ is a K -vector space.

Notice that the rank function of $\mathcal{M}[\Lambda]$ is exactly

$$\text{rk}(J) = \dim_K(K \cdot J).$$

Definition A.2. *Let $\varphi : R \rightarrow S$ be a morphism of integral domains, let M be an R -module. The morphism φ induces a map of matroids*

$$\varphi : \mathcal{M}[\Lambda] \rightarrow \mathcal{M}[\varphi(\Lambda)],$$

where $\varphi(\Lambda)$ is the canonical image of Λ in the tensor product $S \otimes_R M$, considered as S -module. Accordingly, we will write $\varphi(\mathcal{M}[\Lambda]) := \mathcal{M}[\varphi(\Lambda)]$, or $\mathcal{M}[\Lambda] \otimes_R S := \mathcal{M}[\varphi(\Lambda)]$.

Note that the map of matroids is a well-defined homomorphism, as dependent sets are mapped onto dependent sets.

Notations A.3. *In the following, let X be an affine, integral, Noetherian scheme, let \mathcal{F} be a coherent sheaf on X .*

Definition A.4. *We define a presheaf of matroids $\mathcal{M}_{\mathcal{F}}$, the matroid of sections, in the following way: To each Zariski open subset U of X , we associate the set of all matroids*

$$\mathcal{M}_{\mathcal{F}}(U) = \{\mathcal{M}[\Lambda] ; \Lambda \subseteq \mathcal{F}(U), \#\Lambda < \infty\},$$

where $\mathcal{F}(U)$ is considered as an $\mathcal{O}_X(U)$ -module. To a restriction $V \subseteq U$ of open subsets of X , one associates the map of sets of matroids which is induced by the maps of matroids

$$\mathcal{M}[\Lambda] \rightarrow \text{res}_{V,U}(\mathcal{M}[\Lambda])$$

which is induced by the usual restriction morphism of the structure sheaf

$$\text{res}_{V,U} \mathcal{O}_X(U) \rightarrow \mathcal{O}_X(V).$$

Since the sheaf axioms hold, by assumption, for \mathcal{F} , they directly transfer to the matroid of sections $\mathcal{M}_{\mathcal{F}}$, making it a sheaf.

Proposition A.5. *Let $V \subseteq U \subseteq X$ be open subsets, let $M \in \mathcal{M}_{\mathcal{F}}(U)$. Then, as a matroid, M is isomorphic to $\text{res}_{V,U}(M)$.*

Proof. Since X is irreducible, the quotient fields of $\mathcal{O}_X(U)$ and $\mathcal{O}_X(V)$ agree. Thus, the rank functions are equivalent for all J and the image of J in $\mathcal{O}_X(V)$, proving isomorphy of matroids. \square

Remark A.6. *By going to the direct limit, Proposition A.5 implies that for U open in X , and $x \in U$, any matroid $M \in \mathcal{M}_{\mathcal{F}}(U)$ is isomorphic to its canonical image in the stalk $\mathcal{M}_{\mathcal{F},x}$ at x .*

Notations A.7. *Let $x \in X$, let $k(x)$ be the residue field at x . Denote by*

$$.(x): \mathcal{F}_x \rightarrow \mathcal{F}|_x := \mathcal{F}_x \otimes_{\mathcal{O}_{X,x}} k(x)$$

the canonical evaluation of \mathcal{F} at x . It induces a map on subsets $\Lambda \subseteq \mathcal{F}(U)$ with $U \ni x$, and we will write $\Lambda(x)$ for the (element-wise) canonical image of Λ in $\mathcal{F}|_x$ and call it the evaluation of Λ at x .

Definition A.8. *We will denote by*

$$\mathcal{M}_{\mathcal{F}}|_x := \mathcal{M}_{\mathcal{F},x} \otimes_{\mathcal{O}_{X,x}} k(x).$$

the set of matroids $\mathcal{M}[\Lambda]$ with $\Lambda \subseteq \mathcal{F}|_x$, where $\mathcal{F}|_x$ is canonically considered as $k(x)$ -module. This induces a canonical evaluation morphism

$$.(x): \mathcal{M}_{\mathcal{F},x} \rightarrow \mathcal{M}_{\mathcal{F}}|_x,$$

and for $U \ni x$ and $M \in \mathcal{M}_{\mathcal{F}}(U)$, we will write $M(x)$ for its canonical image in $\mathcal{M}_{\mathcal{F}}|_x$ and call it the evaluation of M at x .

Definition A.9. *Let U be an open set in X , let $J \subseteq \mathcal{F}(U)$, let $x \in X$. We will denote by $\text{rk}_x(J)$ the number*

$$\text{rk}_x(J) = \dim_{k(x)}(k(x) \cdot J(x)).$$

Note that by definition, $\text{rk}_x(J)$ is equivalent to the matroid rank of $J(x)$ in any matroid in $\mathcal{M}_{\mathcal{F}}|_x$ (where $J(x)$ is contained in the ground set).

The upper semi-continuity theorem can now be invoked to relate the evaluated matroids to the non-evaluated ones, and provide a genericity result on the non-degeneratedness of the evaluations.

Theorem A.10. *Let U an open set in X , let $J \subseteq \mathcal{F}(U)$. Then, the function*

$$\begin{aligned} U &\rightarrow \mathbb{N} \\ x &\mapsto \text{rk}_x(J) \end{aligned}$$

is upper semi-continuous in the Zariski topology. Moreover, there is an open dense subset $V \subseteq U$ such that

$$\text{rk}_x(J) = \text{rk}_\eta(J) = \text{rk}(J) \quad \text{for all } x \in V,$$

where η is the generic point of X , and the last rank function can be considered in an arbitrary matroid $M \in \mathcal{M}_{\mathcal{F}}(U)$.

In particular, for each matroid $M \in \mathcal{M}_{\mathcal{F}}(U)$, there is an open dense subset $V \subseteq U$ such that we have isomorphisms of matroids

$$M(x) = M(\eta) = M \quad \text{for all } x \in V,$$

where η is the generic point of X .

Proof. Since $\mathcal{O}_{X,x} \cdot J$ is again a coherent sheaf, with

$$k(x) \cdot J(x) = (\mathcal{O}_{X,x} \cdot J)|_x,$$

Theorem A.11 implies upper semi-continuity of the map $x \mapsto \text{rk}_x(J)$. This implies that $\text{rk}_x(J)$ is constant for $x \in V$ where V is open dense in U . Moreover, $\text{rk}_\eta(J)$ is equivalent to $\text{rk}(J)$ for any matroid $M \in \mathcal{M}_{\mathcal{F}}(U)$ due to Remark A.6.

Now let $M = \mathcal{M}[\Lambda]$ with $\Lambda \subseteq \mathcal{F}(U)$ be a matroid in $\mathcal{M}_{\mathcal{F}}(U)$. Since Λ is finite, the power set $\mathcal{P}(\Lambda)$ also is. By the above, for each $J \in \mathcal{P}(\Lambda)$, there is V_J , open dense in U , such that

$$\text{rk}_x(J) = \text{rk}_\eta(J) = \text{rk}(J) \quad \text{for all } x \in V_J,$$

where the last rank function is the rank in M . Set

$$V := \bigcap_{J \in \mathcal{P}(\Lambda)} V_J,$$

which again open dense in U since $\mathcal{P}(\Lambda)$ is finite. Since a matroid is uniquely characterized by the ranks of all subsets of the ground set, see Proposition 2.5.21, it follows that

$$M(x) = M(\eta) = M \quad \text{for all } x \in V,$$

which was the statement to prove. □

Note that Theorem A.10 does not imply that the stalks $\mathcal{M}_{\mathcal{F},x}$ agree on an open dense subset of X . For sake of completeness, we give the form of the upper semi-continuity theorem which was used in the proof of Theorem A.10.

Theorem A.11. *Let \mathcal{F} be a coherent sheaf on a locally Noetherian scheme X . Then, for $i \in \mathbb{N}$ fixed, the function*

$$\begin{aligned} X &\rightarrow \mathbb{N} \\ x &\mapsto \dim_{k(x)}(\mathcal{F}|_x) \end{aligned}$$

is upper semi-continuous in the Zariski topology on X .

Proof. The proof of this theorem is classical and can be found for example in [18], as Example 12.7.2, or by specializing Theorem 12.8 to the given setting. \square

Finally, we want to stress the relation between the sheaf of matroids over the differentials, and algebraic independence sheaves, which has already surfaced in section 2.5.2, as Theorem 2.5.14, and implicitly in [43].

Proposition A.12. *Let $Y = \text{Spec} k$, where k is a field of characteristic zero contained in $K(X)$, and let $X \rightarrow Y$ be the corresponding morphism of schemes. Denote by $\Omega_{X/Y}$ the relative sheaf of differentials of degree one.*

Then, $\mathcal{M}_{\Omega_{X/Y}}$ is the sheaf of all algebraic independence matroids, i.e., it is isomorphic to the sheaf \mathcal{G} constructed in the following way: For $U \subseteq X$ open, the elements of $\mathcal{G}(U)$ are the algebraic independence matroids of finite subsets of $\mathcal{O}_X(U)$ over k , and the restrictions are induced by the restrictions of the structure sheaf \mathcal{O}_X . The isomorphism is given by the canonical differentiation

$$d: \mathcal{O}_X \rightarrow \Omega_{X/Y},$$

inducing a canonical map $\mathcal{G} \rightarrow \mathcal{M}_{\Omega_{X/Y}}$. Thus, for $U \subseteq X$ open, η the generic point of X , and $x \in U$ generic, and $M \in \mathcal{G}(U)$, one has the isomorphisms of matroids

$$M = dM(\eta) = dM(x),$$

Proof. The last directly follows from Theorem A.10, so it suffices to show isomorphy of \mathcal{G} and $\mathcal{M}_{\Omega_{X/Y}}$.

Since d induces a bijection on the underlying sets of M and dM (note that we have allowed multisets, so while d may identify elements, they are kept as copies), it suffices to check that dependent sets in $\mathcal{O}_X(U)$ are mapped to dependent sets in $\Omega_{X/Y}(U)$, and independent sets in $\mathcal{O}_X(U)$ are mapped to independent sets in $\Omega_{X/Y}(U)$. But that is implied by Theorem 16.14 in [11], since k has characteristic zero. \square

Due to Proposition A.12, it is intrinsic to define the following:

Definition A.13. *Keep the situation of Proposition A.12. Then the sheaf of matroids $\mathcal{M}_{\Omega_{X/Y}}$ is called the algebraic independence sheaf of X over Y .*

Remark A.14. *Proposition A.12 indeed gives not only a guarantee that one can always restrict to an open dense subset such that the generic matroidal structure is preserved, but also a tool on algorithmically calculating the generic matroid on irreducible components: namely, sample a random point on the component and calculate the linear matroid on the respective elements in the module of relative differentials, evaluated at that random point. The results of section 2.5 rephrase that in a way which is more specific for the case of Matrix Completion and more hands-on.*