Journal of Machine Learning Research 26 (2025) 1-44 Submitted 7/23; Revised 4/25; Published 4/25

Affine Rank Minimization via Asymptotic Log-Det Iteratively Reweighted Least Squares

Sebastian Krämer KRAEMER@IGPM.RWTH-AACHEN.DE
Institut für Geometrie und Praktische Mathematik
RWTH Aachen University
Aachen, 52062, Germany

Editor: Animashree Anandkumar

Abstract

The affine rank minimization problem is a well-known approach to matrix recovery. While there are various surrogates to this NP-hard problem, we prove that the *asymptotic minimization* of log-det objective functions indeed always reveals the desired, lowest-rank matrices—whereas such may or may not recover a sought-after ground truth. Concerning commonly applied methods such as iteratively reweighted least squares, one thus remains with two difficult to distinguish concerns: how problematic are local minima inherent to the approach truly; and opposingly, how influential instead is the numerical realization. We first show that comparable solution statements do not hold true for Schatten-p functions, including the nuclear norm, and discuss the role of divergent minimizers. Subsequently, we outline corresponding implications for general optimization approaches as well as the more specific IRLS-0 algorithm, emphasizing through examples that the transition of the involved smoothing parameter to zero is frequently a more substantial issue than non-convexity. Lastly, we analyze several presented aspects empirically in a series of numerical experiments. In particular, allowing for instance sufficiently many iterations, one may even observe a phase transition for generic recoverability at the absolute theoretical minimum.

Keywords: affine rank minimization, iteratively reweighted least squares, matrix recovery, non-convex optimization, log-det function

1. Introduction

Affine rank minimization (ARM) is, as the name suggests, the minimization of the rank within an affine set of matrices. Given a surjective linear operator $\mathcal{L} : \mathbb{R}^{n\times m} \to \mathbb{R}^\ell$, for $n \leq m \in \mathbb{N}$, $\ell \in \mathbb{N}$, and a vector $y \in \mathbb{R}^\ell = \operatorname{image}(\mathcal{L})$ referred to as measurements, we accordingly consider the problem to find

$$X^* \in \operatorname*{argmin}_{X\in\mathbb{R}^{n\times m}} \operatorname{rank}(X) \quad \text{subject to} \quad \mathcal{L}(X) = y. \tag{1}$$

We are in particular concerned with observations on the *asymptotic minimization* of the so called *log-det* objective functions,

$$f_\gamma(X) := \log \prod_{j=1}^{n} (\sigma_j^2(X) + \gamma) = \log\det(XX^T + \gamma I), \quad \gamma > 0, \tag{2}$$

where $\{\sigma_j(X)\}_{j=1}^n$ are the singular values of $X \in \mathbb{R}^{n\times m}$.

1.1 Asymptotic Minimizers of Log-Det Functions

The motivation as heuristic and surrogate of the log-det family to the rank function and thereby arising approaches have been analyzed in various works, among others by Mohan and Fazel (2012), whereas further closely related versions have early on been considered for instance by Fazel et al. (2003); Fazel (2002); Mohan and Fazel (2010). While the role and rate of decrease of γ is commonly remarked on, we here emphasize that one truly seeks the limits of (global) minimizers of f_γ for $\gamma \searrow 0$ within the preset $\mathcal{L}^{-1}(y) := \{X \in \mathbb{R}^{n\times m} \mid \mathcal{L}(X) = y\}$. To this effect, Theorem 1 provides that indeed the set of accumulation points

$$\mathcal{X}^* := \left\{X^* \mid \exists (X_\gamma)_{\gamma>0} \subset \mathcal{L}^{-1}(y),\ X^* = \lim_{\gamma\searrow 0} X_\gamma,\ f_\gamma(X_\gamma) \equiv \min_{X\in\mathcal{L}^{-1}(y)} f_\gamma(X)\right\} \tag{3}$$

always yields solutions to the ARM problem defined in (1). Notably, neither the minimizers of $\exp\circ f_0 := X \mapsto \det(XX^T)$ nor of f_γ for some specific $\gamma > 0$, nor such of Schatten-p functions do so, if not subject to further assumptions (Figure 1 and Lemma 4). Theorem 1 further captures a rare case which we refer to as *rank minimizers at infinity* (Section 2.4), where the situation can be more involved than for the simplified set defined in (3).

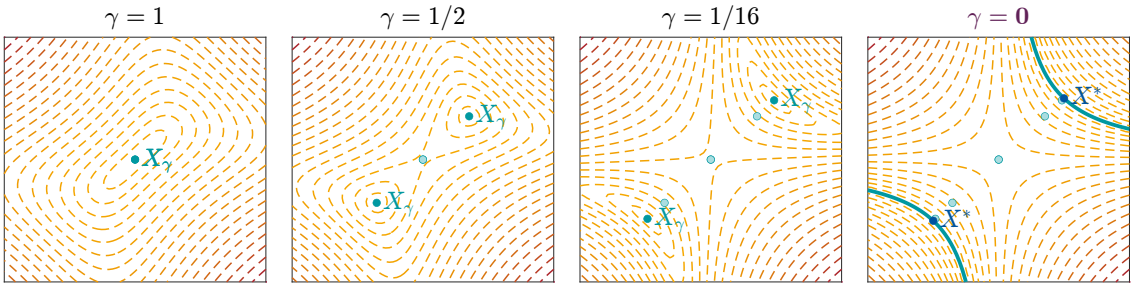

Figure 1: Contour lines and minimizers of $f_\gamma|_{\mathcal{L}^{-1}(y)}$ (as in Example 15) for varying $\gamma > 0$, whereas the fourth picture for $\gamma = 0$ shows the two $X^* \in \mathcal{X}^*$ defined in (3) and the curved lines of minimizers of $\exp\circ f_0|_{\mathcal{L}^{-1}(y)}$ (that is, all singular matrices in $\mathcal{L}^{-1}(y)$).

The family $\{f_\gamma\}_{\gamma>0}$ of log-det functions considered here is not the only reasonable choice, but it realizes an elegant algebraic property that yields a further motivation for asymptotic minimization. The varieties of low-rank matrices

$$V_{\leq r} = \{X \in \mathbb{R}^{n\times m} \mid \operatorname{rank}(X) \leq r\}, \quad \dim(V_{\leq r}) = r(n+m-r),$$

are nested and each characterized by the common roots of the polynomials $p_{r,M}(X) = \det(X|_M)$ for all $(r+1)\times(r+1)$ submatrices $X|_M$ (*minors*), $r \leq n$. These polynomials then appear in the Taylor series

$$\exp\circ f_\gamma(X) = \gamma^n + \sum_{r=0}^{n-1} p_r(X)^2\gamma^{n-1-r}, \quad p_r(X)^2 := \sum_M p_{r,M}(X)^2 = 0 \ \Leftrightarrow\ X \in V_{\leq r}. \tag{4}$$

Comparison through L'Hôpital's rule hence already suggests that the minimal limit value for $\gamma \searrow 0$ must be given through a matrix X with the lowest possible rank r, for which thus the factor $p_r(X)$ still vanishes.

1.2 Algorithmic Approach via Iteratively Reweighted Least Squares (IRLS-0)

Whereas finding the limits of (global) minimizers of f_γ, $\gamma \searrow 0$, is naturally not guaranteed to lie within the reach of numerically feasible methods, limits of stationary points can quite efficiently be determined via well-known *iteratively reweighted least squares*. Given a starting value $X^{(0)} \in \mathcal{L}^{-1}(y)$, one basic version of that method yields the simple updates

$$X^{(i)} := \operatorname*{argmin}_{X \in \mathcal{L}^{-1}(y)} \|W^{1/2}_{\gamma^{(i-1)}, X^{(i-1)}} X\|_F, \quad W_{\gamma,X} := (XX^T + \gamma I)^{-1}, \tag{5}$$

for $0 < \gamma^{(i)} \leq \gamma^{(i-1)}$, $i \in \mathbb{N}$, as for instance described by Mohan and Fazel (2012). The profound relation to the log-det functions is not directly apparent (cf. Section 4.1), but the suitability of this iteration can be motivated (though certainly not proven) by the fact that

$$\lim_{\gamma \searrow 0} \|W^{1/2}_{\gamma,X} X\|_F^2 = \lim_{\gamma \searrow 0} \sum_{j=1}^{n} \frac{\sigma_j(X)^2}{\sigma_j(X)^2 + \gamma} = \operatorname{rank}(X).$$

A suitable discretization of the transition of the smoothing parameter $\gamma \searrow 0$ is of significant importance in practice, as the iterates may otherwise converge or seem to converge to arbitrary points that only appear to be local minima (Examples 15 and 20 and Lemma 11). In turn, slower decrease of the sequence $(\gamma^{(i)})_{i \in \mathbb{N}_0}$ or, if in turn applicable, faster progress trough an adapted method (as by Kümmerle and Verdun, 2021) may overcome the problem (Proposition 9 and Experiment 17). If such care is taken, ultimately only the behavior of the log-det functions is decisive, where then the approach even allows to observe a phase transition near the exact theoretical minimum for generic recoverability (Section 5.2.2).

1.3 Affine Rank Minimization Opposed to Matrix Recovery

Due to its numerous potential applications (for surveys, see Davenport and Romberg, 2016; Nguyen et al., 2019), affine rank minimization is frequently analyzed through the perspective of matrix recovery, or predominantly an approach to the latter, where one desires to reconstruct a low-rank ground truth $X^{(\text{true})}$ from measurements $y = \mathcal{L}(X^{(\text{true})})$. Evaluation of a rank minimization algorithm by means of its recovery properties can however make its suitability indistinguishable from the inherent benevolence of the problem setting; that is, the supposed ground truth may simply not be the solution to (1) depending on the number and type of measurements. Here, we are almost exclusively interested in the affine rank minimization problem itself, regardless of whether minimizers $X^* \in \mathcal{X}^*$ pose any sort of reconstruction. While we additionally reflect on the latter in the numerical experiments, we instead refer to works by Breiding and Vannieuwenhoven (2022) and Breiding et al. (2023) for a theoretical analysis of that matter. In particular, as Theorem 1 ensures that matrices in $\mathcal{X}^*$ solve the ARM problem, there are basically two remaining questions:

1. In which cases are there obstructive local minima?
2. How well is the continuous asymptotic minimization numerically realizable?

The first question may typically be of more concern given the non-convexity of the approach. However, we outline that this problem may be overestimated, because the critical nature of necessarily discretized asymptotic minimization may be underestimated—in particular since it turns out that it is quite difficult to distinguish these two issues (cf. Section 1.2).

1.4 Related Approaches and Motivations

1.4.1 Affine Cardinality Minimization

A theoretical predecessor to the ARM problem is affine cardinality minimization (ACM), where one seeks to minimize the cardinality $\text{card}(x)$, that is, the number of non-zero coefficients of x within an affine set of vectors. The analogy between methods for both cases is provided through the similar roles of the sorted, absolute values $|x_{i_1}| \geq \ldots > |x_{i_{s+1}}| = 0$ for an s-sparse vector and singular values $\sigma_1(X) \geq \ldots > \sigma_{r+1}(X) = 0$ of a rank-r matrix X,

$$\text{card}((\sigma_i(X))_{i=1}^n) = \text{rank}(X).$$

Like ACM, also the ARM problem is in general NP-hard (Natarajan, 1995).

1.4.2 Log-Det as Limit Case of Schatten-p Functions and IRLS-p

Whereas the log-det objective functions (2) can be motivated as generalization of the vector case as by Candès et al., 2008 (cf. Section 1.4), they can also be derived as limit cases of smoothed Schatten-p functions $S_{\gamma,p}(X)$ as by Mohan and Fazel (2012). This follows from the observation that

$$\lim_{p \searrow 0} \frac{S_{\gamma,p}(X) - n}{p} = \frac{1}{2} f_\gamma(X), \quad S_{\gamma,p}(X) := \sum_{j=1}^n (\sigma_j^2(X) + \gamma)^{\frac{p}{2}}, \quad 0 < p \leq 1. \tag{6}$$

Naturally, a similar derivation has been made for vectors by Daubechies et al. (2010). The case $p = 1$ yields a smoothed version of the nuclear norm $\|\cdot\|_*$ and is comprehensively introduced in a work by Fornasier et al. (2011). The minimization approach through IRLS-p is likewise closely related. In fact, the only principal modification required to generalize (5) is to replace (see Algorithm 2 for more details)

$$W_{\gamma,X} := (XX^T + \gamma I)^{p/2-1}, \quad p \in [0, 1].$$

However, the smoothed Schatten-p functions, $p > 0$, may in contrast to the log-det case $p = 0$ not always yield the desired minimizers, independent of the algorithmic approach to minimize such. We discuss this in more detail in Section 2.3.

1.4.3 Reconstruction Bounds and Convexification

The convex relaxation of the cardinality to the ℓ_1-norm has a long and well known history (Candès and Recht, 2009; Candes and Tao, 2010; Gross, 2011; Recht, 2011) associated to compressed sensing; one of the earlier works being by Fazel (2002). Major tools arising from this area to derive lower bounds on the number of measurements to almost always or with high probability guarantee successful recoveries include the partially related null space property (Oymak et al., 2011; Recht et al., 2010b), the restricted isometry property (Recht et al., 2010a) as well as the incoherence property for matrix completion (Candes and Tao, 2010). Optimization algorithms for the analogous convex relaxation of the rank, which yields the nuclear norm, include methods known as soft thresholding (Cai et al., 2010), fixed-point continuation (Goldfarb and Ma, 2011) and more, among others by Hastie et al. (2015); Keshavan et al. (2010a,b); Lee and Bresler (2010). Also iteratively reweighted least squares

(IRLS-1) has been applied as by Fornasier et al., 2011 (cf. Section 1.4.2). Although some of the methods involve asymptotically *optimal* bounds (up to polylogs), these can be quite far from truly optimal bounds or require too demanding assumptions. Based on empirical results, ℓ_1- or nuclear norm minimization seems to be outperformed by reweighted versions (Chartrand and Staneva, 2008; Chartrand and Yin, 2008; Daubechies et al., 2010; Mohan and Fazel, 2012) which can also be observed in our numerical experiments (Section 5). Still, the strength of (likely) reconstruction guarantees has so far mostly been on the side of convexification.

1.4.4 Data-Sparse Approaches and Machine Learning Context

Other methods optimize data-sparse models, commonly low-rank factorizations, for which the article by Chi et al. (2019) provides a survey. In machine learning terminology, for instance in recommender systems, one may refer to such as matrix factorization based collaborative filtering algorithms, where the herein applied term rank relates to the number of so called latent factors. The log-det objective however serves as template for a practically viable automated tuning of that hyperparameter even in case of noise.

Naturally, these models scale better with growing dimensions, though one can also implicitly make use of such (Fornasier et al., 2011; Kümmerle and Verdun, 2021). Through additional regularization, explicit rank estimates can likewise be avoided (Sun and Luo, 2016). For instance, the method by Giampouras et al. (2020) aims at the minimization of the Schatten-p functions (6). Those that work on a fixed rank instead include Riemannian optimization (Chi et al., 2019; Vandereycken, 2013), nonlinear successive overrelaxation (Xu et al., 2012), hard thresholding (Tanner and Wei, 2013) and others (Balzano et al., 2010; Dai et al., 2011). However, at the latest in the setting of tensor recovery, a non-invasive adaption of the multivariate rank may pose an obstructive problem (Grasedyck and Krämer, 2019), making the herein considered approaches particularly interesting.

1.4.5 Derivation and Interpretation of IRLS

The iteratively reweighted least squares algorithm itself as well as related variants have been motivated as majorization-minimization algorithm for the vector case (Candès et al., 2008), or for instance through entropy minimization (Rao and Kreutz-Delgado, 1999). So called log-thresholding has further appeared in the work by Malioutov and Aravkin (2014). A relaxation of the affine constraint into a penalty term to allow for more noise has been analyzed for instance by Lai et al. (2013). In Section 4.1, we outline the derivation through the majorizing linearization of the convex log-det functions (Candès et al., 2013) as it is most convenient to our intentions. Moreover, this idea has more recently been generalized (Kümmerle and Verdun, 2021; Kümmerle, 2019), replacing *one-sided* weights with superior *both-sided* weights (cf. Section 5.1.2). The modified IRLS-0 method derived therein is also among the first to provide quadratic, local convergence rates, given incoherence assumptions in the completion setting.

1.5 Organization of the Remainder of the Paper

In order to distinguish principle properties of log-det minimization from the ones that stem from the discretization of the asymptotic process as well as the specific behavior of

IRLS-p optimization, we keep statements initially more general. Throughout the paper, we thereby motivate and further narrow down observations to more particular algorithms and parameter adaptions. In Section 2, we prove the earlier introduced statement that asymptotic log-det minimization in contrast to Schatten-p minimization always solves ARM, with certain peculiarities for settings where the minimal rank is only attained at infinity. Concerns resulting from the discretization of the transition $\gamma \searrow 0$ for a fairly general optimization template are considered in Section 3, highlighting the need for, but also the sufficiency of a slow enough yet steady decrease of that parameter. A short derivation and clarification of IRLS-p, with focus on $p = 0$, is given in Section 4, for which we exemplarily analyze stagnation based issues in Section 4.3 and Appendix B. All proofs of lemmas are postponed to Appendix A. Section 5 continued by Appendices C and D contain a series of detailed numerical experiments on the inherent properties of log-det and Schatten-p minimization with particular attention to the rate of decrease of the smoothing parameter. Appendices E and F contain and explain further visualizations of the resulting data.

2. Asymptotic Minimization

The following Theorem 1 characterizes, under no assumptions towards the affine preset, more general properties of asymptotic log-det minimizers compared to the simplified statement about limits $\mathcal{X}^*$ in (3). Throughout this section, we denote with

$$r^* = \min_{X \in \mathcal{L}^{-1}(y)} \operatorname{rank}(X) \quad \text{and} \quad \mathcal{G}_\gamma^* := \operatorname*{argmin}_{X \in \mathcal{L}^{-1}(y)} f_\gamma(X)$$

the minimally attainable rank and the sets of (global) minimizers of f_γ within the affine preset. Let further

$$(X_k)_{k\in\mathbb{N}} \subset \mathcal{L}^{-1}(y),\ X_k \in G_{\gamma_k},\ k \in \mathbb{N}, \quad \text{with} \quad (\gamma_k)_{k\in\mathbb{N}} \subset \mathbb{R}_{>0},\ \lim_{k\to\infty} \gamma_k = 0, \tag{7}$$

be a sequence of (global) minimizers for a monotonically decreasing sequence of smoothing parameters that converge to zero. Whereas indeed diverging sequences exist in settings in which a rank minimizer lies at infinity (Section 2.4), already minor assumptions exclude this possibility and guarantee sequences of global minimizers to be bounded, which thereby reveal explicit solutions to the ARM problem (Section 2.2). Contrarily, the subsequent Lemma 4 constitutes a simple counterexample demonstrating that this does not generally hold true for Schatten-p minimization. Further, whereas for $p > 0$, minimizers remain bounded, Section 2.4 discusses that this is not necessarily desired behavior.

2.1 Main Theorem

We first describe the general situation including diverging sequences. Minor assumptions on the operator $\mathcal{L}$ satisfied in generic settings however already suffice to guarantee the boundedness of the sequence (Corollary 2). Further, under mildly stronger assumptions, the sequence even converges to a unique rank minimizer (Corollary 3).

Theorem 1. *For a sequence as in* (7)*, it holds* $\lim_{k\to\infty} \sigma_{r^*+1}(X_k) = 0$ *and*

$$\limsup_{k\to\infty} \prod_{j=1}^{r^*} \sigma_j(X_k) \leq \inf_{X \in \mathcal{L}^{-1}(y),\ \operatorname{rank}(X)=r^*} \prod_{j=1}^{r^*} \sigma_j(X). \tag{8}$$

In particular, every accumulation point A of $(X_k)_{k\in\mathbb{N}}$ is rank r^* and satisfies

$$A \in \operatorname*{argmin}_{X\in\mathcal{L}^{-1}(y),\ \operatorname{rank}(X)=r^*} \prod_{i=1}^{r^*} \sigma_i(X).$$

Further, every diverging subsequence $(X_{k_\ell})_{\ell\in\mathbb{N}}$ necessitates $\lim_{\ell\to\infty}\sigma_{r^*}(X_{k_\ell})=0$.

Proof Let $c\in\mathbb{R}_{\geq 0}$ denote the righthand side infinum value of the to be proven inequality (8). For any rank-r^* matrix $\hat{X}\in\mathcal{L}^{-1}(y)$ and $\overline{r}\geq r^*$, we have

$$q_k(\hat{X}) := \frac{\exp f_{\gamma_k}(\hat{X})}{\exp f_{\gamma_k}(X_k)} = \prod_{j=1}^{n}\frac{\sigma_j(\hat{X})^2+\gamma_k}{\sigma_j(X_k)^2+\gamma_k} \leq \prod_{j=1}^{\overline{r}}\frac{\sigma_j(\hat{X})^2+\gamma_k}{\sigma_j(X_k)^2+\gamma_k} =: q_k^{\overline{r}}(\hat{X}).$$

Assuming to the contrary that $\limsup_{k\to\infty}\sigma_{r^*+1}(X_k)\neq 0$, it follows $\liminf_{k\to\infty}q_k(\hat{X}) = \liminf_{k\to\infty}q_k^{r^*+1}(\hat{X}) = 0 < 1$. If $\limsup_{k\to\infty}\prod_{j=1}^{r^*}\sigma_j(X_k) = c+2\varepsilon > c$ for some $\varepsilon>0$, then by definition of the infinum, there exists a matrix $\hat{X}\in\mathcal{L}^{-1}(y)$ with $\operatorname{rank}(\hat{X})=r^*$ and $\prod_{j=1}^{r^*}\sigma_j(\hat{X}) = \hat{c}\leq c+\varepsilon$. Thus, $\liminf_{k\to\infty}q_k(\hat{X})\leq\liminf_{k\to\infty}q_k^{r^*}(\hat{X})<1$. In both cases, it follows that for some $k\in\mathbb{N}$, it is $f_{\gamma_k}(\hat{X}) < f_{\gamma_k}(X_k)$, which is in contradiction to the minimizing property of $X_k\in\mathcal{G}^*_{\gamma_k}$. Thus, the first two assertions of the theorem must hold true.
Let now $A\in\mathcal{L}^{-1}(y)$ be an accumulation point of the sequence as limit of the subsequence $(X_{k_\ell})_{\ell\in\mathbb{N}}$. Then since also $\lim_{k\to\infty}\sigma_{r^*+1}(X_{k_\ell})=0$, the matrix A must have rank r^* (as this is also the minimal possible rank in the affine set). In particular, we obtain

$$c\leq\prod_{j=1}^{r^*}\sigma_j(A)\leq\limsup_{k\to\infty}\prod_{j=1}^{r^*}\sigma_j(X_k)\leq c.$$

Lastly, assume there is a diverging subsequence $(X_{k_\ell})_{\ell\in\mathbb{N}}$. Since $\limsup_{\ell\to\infty}\prod_{j=1}^{r^*}\sigma_j(X_{k_\ell})$ is bounded by $c<\infty$ but $\sigma_1(X_{k_\ell})\to\infty$, it must $\sigma_{r^*}(X_{k_\ell})\to 0$. ∎

We additionally note here that by the prior proof, the existence of any accumulation point implies that the infinum is attained at least in that accumulation point.

2.2 Boundedness and Convergence of Asymptotic Minimizers

So far no assumptions have been made on the affine preset $\mathcal{L}^{-1}(y)$. Provided that comparatively weak restricted isometry properties hold true, one can already substantially limit the behavior of the global minimizers. For that purpose, we say the operator $\mathcal{L}$ fulfills the r-RIP, if Definition 3.1 by Recht et al. (2010a) holds true for a constant $\delta_r\in[0,1)$. Further, for two matrix subsets $S_1,S_2\subseteq\mathbb{R}^{n\times m}$, we denote $\operatorname{dist}(S_1,S_2) := \inf_{X_1\in S_1, X_2\in S_2}\|X_1-X_2\|$.

Corollary 2. *If $\mathcal{L}$ satisfies the (r^*-1)-RIP, then $\operatorname{dist}(\mathcal{L}^{-1}(y),V_{\leq r^*-1})>0$. If the latter holds true, then $(X_k)_{k\in\mathbb{N}}$ is bounded.*

Proof Assume to the contrary that $\operatorname{dist}(\mathcal{L}^{-1}(y),V_{\leq r^*-1})=0$. Then there must exist a sequence $(X_k)_{k\in\mathbb{N}}\subset\mathcal{L}^{-1}(y)$ together with a corresponding sequence of best rank-(r^*-1) approximations $(X_{k,r^*-1})_{k\in\mathbb{N}}$ with $\lim_{k\to\infty}\|X_{k,r^*-1}-X_k\|=0$. This implies

$$\|\mathcal{L}(X_{k,r^*-1})-y\| = \|\mathcal{L}(X_{k,r^*-1}-X_k)\|\leq\|\mathcal{L}\|\|X_{k,r^*-1}-X_k\|\to 0,\quad k\to\infty.$$

As r^* is the minimal attainable rank, $(X_{k,r^*-1})_{k\in\mathbb{N}}$ can not have an accumulation point. Thus $\lim_{k\to\infty}\|\mathcal{L}(X_{k,r^*-1})\| = \|y\| < \infty$, $\operatorname{rank}(X_{k,r^*-1}) \equiv r^* - 1$, but $\lim_{k\to\infty}\|X_{k,r^*-1}\| = \lim_{k\to\infty}\|X_k\| = \infty$. Hence, the (r^*-1)-RIP can not hold true. For the second part, it only remains to consider that the existence of a sequence with $\lim_{k\to\infty}\sigma_r(X_k) = 0$ would require $\operatorname{dist}(\mathcal{L}^{-1}(y), V_{\leq r-1}) = 0$. The rest then follows by Theorem 1. ■

Under maybe the most common assumption on a unique rank minimizer, convergence of the sequence follows. Alternatively, criteria by Breiding and Vannieuwenhoven (2022) and Breiding et al. (2023) provide uniqueness in generic settings, such that then the prior, weaker RIP condition suffices.

Corollary 3. *If $\mathcal{L}$ satisfies the $2r^*$-RIP, then there is a unique rank-r^* matrix X^* in $\mathcal{L}^{-1}(y)$ (as is well-known, Recht et al., 2010a) and $(X_k)_{k\in\mathbb{N}}$ is bounded (by Corollary 2). If these latter two conditions hold true, then the sequence $(X_k)_{k\in\mathbb{N}}$ converges to X^*.*

Proof Considering the second part, as the sequence is bounded by assumption, there must exist an accumulation point. Since further, by Theorem 1, any such accumulation point must be the by assumption unique rank minimizer, the entire sequence must converge to that matrix. ■

Even though the k-RIP does not hold true (given a non-trivial k) for entry-wise samples as in matrix completion problems, as opposed to Gaussian measurements, there are related approaches as mentioned in Section 1.4.3.

2.3 Comparison to Schatten-p Minimization

Schatten-p minimization for $p > 0$ (cf. (6)) exhibits a distinct difference to the limit case in which we are interested. For $p = 0$, the parameter transition $\gamma \searrow 0$ is a fundamental part of the approach independent of the applied algorithm. Namely, for $\gamma = 0$, every rank-deficient matrix $X \in \mathcal{L}^{-1}(y)$ is already a minimizer (cf. Figure 1). This case is hence only useful for $r^* = n - 1$. Contrarily, for $p > 0$, one ultimately searches for minimizers of $S_{0,p}$ while the smoothing parameter $\gamma > 0$ is only required for a stable application of iteratively reweighted least squares.

While choosing $p = 0$ can always be beneficial, asymptotic minimization may be necessary within this class of approaches, depending on the problem setting. Namely, for every $p > 0$, there are always instances for which the ARM problem is not solved even by global minimizers due to, roughly speaking, a prediction bias towards zeros (as for instance observable later in (9)). As demonstrated in the following Lemma 4, this tends to become increasingly worse for large values $p \in (0, 1]$. This extreme case, $p = 1$, has the in principle strong advantage that the to be minimized (smoothed) nuclear norm is convex, yet we could not observe an advantage in practice regarding the frequency of recovery (cf. Section 5.2.1).

Lemma 4. *For $n \in \mathbb{N}$, let $\mathcal{L} : \mathbb{R}^{n\times n} \to \mathbb{R}^{n^2-1}$ be a linear operator and let $y_\delta \in \mathbb{R}^n$ constitute measurements depending on some $\delta \in \mathbb{R}^{n-1}$ such that*

$$\mathcal{L}^{-1}(y_\delta) = \{X_\delta(a) \mid a \in \mathbb{R}\}, \quad X_\delta(a) := \begin{pmatrix} \delta\delta^T & \delta \\ \delta^T & a \end{pmatrix} \in \mathbb{R}^{n\times n}, \quad \Delta_\delta := \|\delta\|_F^2 \in (0, 1).$$

Then for all $p > 0$ and $\gamma > 0$, the derivative of the (smoothed) Schatten-p function at the one rank-one matrix $X_\delta(1) \in \mathcal{L}^{-1}(y_\delta)$ is given by

$$(S_{\gamma,p} \circ X_\delta)'(1) = \frac{p}{((1+\Delta_\delta)^2+\gamma)^{1-p/2}} > 0.$$

Thus, the minimizer of $S_{\gamma,p}|_{\mathcal{L}^{-1}(y_\delta)}$ does not minimize the rank for any $\delta, \gamma > 0$. Further, considering $\gamma = 0$, we have the conditional statement that

$$4^{\frac{1}{p}} < \frac{(1+\Delta_\delta)^2}{\Delta_\delta} \quad \Rightarrow \quad S_{0,p}(X_\delta(\Delta_\delta)) < S_{0,p}(X_\delta(1)). \tag{9}$$

Thus, given that bound, neither the minimizer of $S_{0,p}|_{\mathcal{L}^{-1}(y_\delta)}$ nor the limit of minimizers of $S_{\gamma,p}|_{\mathcal{L}^{-1}(y_\delta)}$ for $\gamma \searrow 0$ minimize the rank. In particular, the nuclear norm minimizer $(p = 1)$ is always given through $a = \Delta_\delta < 1$.

As the second derivative at $a = 1$ of $S_{\gamma,p} \circ X_\delta(a)$, $p \in (0, 1)$, tends to positive infinity for $\gamma \searrow 0$, the neighbored local minimizers do converge to the rank minimizer. These local minima around $a = 1$ however only exist for $\gamma > 0$ small enough, the latter depending on p and δ. Contrarily, the log-det minimizers converge to the desired solution. This can be concluded by Corollary 2 from $\operatorname{dist}(\mathcal{L}^{-1}(y_\delta), V_{\leq 0}) = \inf_{a\in\mathbb{R}} \|X_\delta(a)\| > \delta > 0$. In fact, since

$$\exp \circ f_\gamma(X_\delta(a)) = \left(\Delta_\delta^2(a-1)^2 + (\Delta_\delta^2 + 2\Delta_\delta + a^2)\gamma + \gamma^2\right)\gamma^{n-2}, \tag{10}$$

its minimizer $X_\delta(a_\delta^*)$ is also each the only stationary point,[1] given through

$$a_\delta^* = \frac{\Delta_\delta^2}{\Delta_\delta^2 + \gamma} \xrightarrow[\gamma\searrow 0]{} 1, \quad \delta > 0.$$

It is further worth noting that Δ_δ is directly related to (in)coherence values of the rank-minimizing matrix. While there are different version, the classical one (Candès and Recht, 2009, Def. 1.2) sets a coherence value which for $X_\delta(1)$ equals $\mu = \frac{n}{\Delta_\delta+1}$. Thus if $\mu > \frac{n}{2}$, then $\Delta_\delta < 1$ and nuclear norm minimization $(p = 1)$ does not return the desired $X_\delta(1)$.

2.4 Rank Minimizers at Infinity

The convergence or even boundedness discussed in Section 2.2 does not necessarily hold true when rank minimizers lie at infinity. One may dismiss these cases as ill-posed, but it is likewise sensible to ask for ones solution to accordingly diverge. Certainly though, the situation is not as clear anymore. Comparative settings are well-known as for example mentioned by Breiding and Vannieuwenhoven (2022); Tanner et al. (2019). The therein considered instance here translates to

$$\mathcal{L}^{-1}(y) = \{X(a) \mid a \in \mathbb{R}\}, \quad X(a) := \begin{pmatrix} a & 1 \\ 1 & 0 \end{pmatrix}.$$

1. This observation holds true generally considering the description in (4). When $\mathcal{L}^{-1}(y)$ is one-dimensional, the affine space can be parameterized by one parameter $a \in \mathbb{R}$. The latter then appears at most quadratically in $\exp \circ f_\gamma$ and thus linearly in $\frac{\partial}{\partial a} \exp \circ f_\gamma = 0$, yielding exactly one stationary point.

One the one hand, $\lim_{a\to\infty} \sigma_2(X(a)) = 0$ holds true while the rank of $X(a)$ is always full. Thus, the 1-RIP (cf. Section 2.2) is not satisfied. However, the determinant is constant with $\det(X(a)X(a)^T) \equiv 1$, and although not captured by Theorem 1, the set of global minimizers for any $\gamma > 0$ is just given through $a = 0$. This is different concerning the following Proposition 5 based on Example 6.

Proposition 5. *There exists a linear operator $\mathcal{L}$ and $y \in \text{image}(\mathcal{L})$, such that*

$$\inf_{X\in\mathcal{L}^{-1}(y),\ \text{rank}(X)=r^*} \prod_{j=1}^{r^*} \sigma_j(X) = 0, \quad \textit{even though} \quad r^* = \min_{X\in\mathcal{L}^{-1}(y)} \text{rank}(X).$$

In such a case, a sequence $(X_k)_{k\in\mathbb{N}}$ as defined in (7) can not have accumulation points and hence diverges.

Proof The setting itself is given in form of the separate Example 6. Regarding the second statement, as the infinum is zero, Theorem 1 here implies that every accumulation point A were to satisfy $\prod_{j=1}^{r^*} \sigma_j(A) = 0$ as well. This would however require $\text{rank}(A) < r^*$. ■

Example 6. Let $\mathcal{L} : \mathbb{R}^{2\times 3} \to \mathbb{R}^4$ be a linear operator and $y \in \mathbb{R}^4$ such that

$$\mathcal{L}^{-1}(y) = \{X(a,b) \mid a,b \in \mathbb{R}\}, \quad X(a,b) := \begin{pmatrix} a & 1 & a+b \\ b+1 & b & b+1 \end{pmatrix}.$$

The rank of $X(a,b)$ is always $r^* = 2$. For $(a_k, b_k) = (k, 1/k)$, we however obtain

$$X_k := X(a_k,b_k) = \begin{pmatrix} k & 1 & k+1/k \\ 1+1/k & 1/k & 1+1/k \end{pmatrix}, \quad \det(X_k X_k^T) = \frac{3}{k^2} + \frac{2}{k^4} \to 0.$$

In fact, the function $g_0 : \mathbb{R}^2 \to \mathbb{R}$, $g_0(a,b) := f_0(X(a,b))$ (cf. Figure 2), has only one stationary point in form of a saddle point at $(a,b) = (1,0)$ with value $g_0(1,0) = \log(2)$, while there are two disconnected valleys without stationary points (in which f_0 monotonically declines in direction of $a \to \infty$). The sequence $(a_k, b_k) = (k, 1/k)$ does not describe one of the valleys, but approaches the one in the positive quadrant.

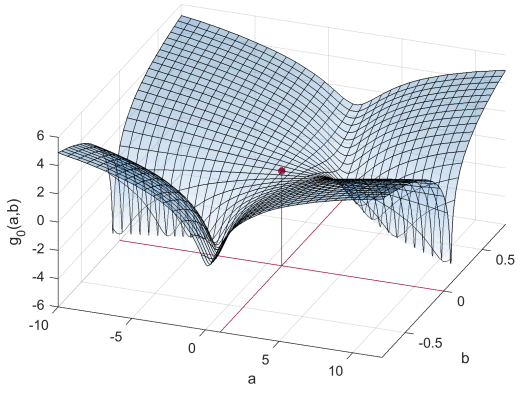

Figure 2: The function g_0 as well as the saddle point $(1,0)$ with $g_0(1,0) = \log(2)$ described in Example 6.

3. Discretized Asymptotic Minimization

The discretization of the transition $\gamma \searrow 0$ is a necessary step towards a numerically feasible method. To distinguish from issues that stem from the eventual optimization algorithm, we define the template Algorithm 1 for a general updating function $\Psi_\gamma : \mathcal{L}^{-1}(y) \to \mathcal{L}^{-1}(y)$ that satisfies the following Assumption 7.

Algorithm 1 Asymptotic Minimization

1: set $X^{(0)} \in \mathcal{L}^{-1}(y)$ and $\gamma^{(0)} > 0$
2: **for** $i = 1, 2, \ldots$ **do**
3: $\quad X^{(i)} := \Psi_{\gamma^{(i-1)}}(X^{(i-1)})$
4: $\quad$ set $\gamma^{(i)} \leq \gamma^{(i-1)}$ according to chosen strategy
5: **end for**

The statements in Assumption 7 are well-established and similarly expected for reasonable Schatten-p minimization (Fornasier et al., 2011; Mohan and Fazel, 2012). Although phrased for a limit $\gamma^* := \lim_{i\to\infty} \gamma^{(i)} > 0$, they mostly determine the behavior of the iterates for the main case of interest, $\gamma^* = 0$, depending on the sequence $(\gamma^{(i)})_{i\in\mathbb{N}_0}$ (Section 3.1). We note that the minimizers remain bounded if the smoothing parameter does, namely $\bigcup_{\gamma\geq\gamma^*>0} G^*_\gamma$ is always bounded (cf. Section 2). Even though empirically neglectable, one may construct affine sets, where for $\gamma^* = 0$, the minimizers are indeed unbounded and cause the iterates to diverge (Section 3.2). While this can be considered desired behavior, certain natural strategies how to adapt the smoothing parameter can in turn, when counted as part of the iteration, introduce artificial fixed points independent of Ψ_γ (Section 3.3), which do have practical relevance. In order to include not just the limit, but explicitly allow $\gamma = 0$ for later purposes, let in the following $\mathcal{S}^*_\gamma$ be each the set of stationary points of the formally rescaled[2] function $\exp \circ f_\gamma|_{\mathcal{L}^{-1}(y)}$, for $\gamma \geq 0$.

Assumption 7. *Let $(X^{(i)})_{i\in\mathbb{N}_0}$ be generated by Algorithm 1 for the weakly decreasing sequence $(\gamma^{(i)})_{i\in\mathbb{N}_0} \subset \mathbb{R}_{\geq 0}$ and let $\gamma^* := \lim_{i\to\infty} \gamma^{(i)}$.*

(i) For each $i \in \mathbb{N}$, it holds $f_{\gamma^{(i)}}(X^{(i)}) \leq f_{\gamma^{(i-1)}}(X^{(i-1)})$.

(ii) If $\gamma^* > 0$, then the sequences $(X^{(i)})_{i\in\mathbb{N}_0}$ as well as $(|f_{\gamma^{(i)}}(X^{(i)})|)_{i\in\mathbb{N}_0}$ remain bounded.

(iii) If the sequences $(X^{(i)})_{i\in\mathbb{N}_0}$ and $(|f_{\gamma^{(i)}}(X^{(i)})|)_{i\in\mathbb{N}_0}$ remain bounded, then each accumulation point of $(X^{(i)})_{i\in\mathbb{N}_0}$ is in $\mathcal{S}^*_{\gamma^*}$ and $\lim_{i\to\infty} \|X^{(i)} - X^{(i-1)}\| = 0$.

3.1 Asymptotics of Stationary Points

To approach global minima is certainly not guaranteed in practice. The local analogous to the asymptotically defined, desired accumulation points in (3) (see Section 2 for more detail) are such of stationary points. That is, we consider the existence of subsequences

2. This function has the same stationary points as $f_\gamma|_{\mathcal{L}^{-1}(y)}$ for $\gamma > 0$, as well as for $\gamma = 0$ if $-\infty$ is formally allowed and treated as minimal value.

$(X^{(i_\ell)})_{\ell\in\mathbb{N}_0}$ such that

$$\lim_{\ell\to\infty} \operatorname{dist}(\mathcal{S}^*_{\gamma^{(i_\ell)}}, X^{(i_\ell)}) = 0 \quad \text{for} \quad \gamma^* = \lim_{\ell\to\infty} \gamma^{(i_\ell)} = 0. \tag{11}$$

We can not provide a constructive, practicable method for this general case. The following Proposition 9 merely ensures that under Assumption 7, there is always the principle possibility to reduce the parameter γ slowly enough such that (11) holds true. Note that while Assumption 7 (*iii*) asks for properties of bounded sequences, Proposition 9 does not assume that the iterates remain bounded. In fact, (11) also holds true for the diverging sequence one obtains in the setting of Example 6.

Remark 8. Proposition 9 can roughly be phrased as follows: If the sequence $(\gamma^{(i)})_{i\in\mathbb{N}_0}$ is decreased to $\gamma^* = 0$ slowly enough, then $X^{(i)}$ can only converge to a limit of stationary points of $f_\gamma|_{\mathcal{L}^{-1}(y)}$ for $\gamma \searrow 0$.

Proposition 9. *(See Remark 8) Let $(X^{(i)})_{i\in\mathbb{N}_0}$ be generated by an algorithm that satisfies Assumption 7. Let $\Gamma = \{\nu^n \mid n \in \mathbb{N}_0\} \subset \mathbb{R}_{>0}$ for some $\nu \in (0,1)$, and let*

$$\delta_i := \inf_{S\in\mathcal{S}^*_{\gamma^{(i)}}} \|X^{(i)} - S\|_F, \quad i \in \mathbb{N}.$$

For $\gamma^{(0)} := \max(\Gamma) = 1$, we recursively define

$$\gamma^{(i+1)} := \begin{cases} \theta_i & \text{if } \delta_i < \theta_i \\ \gamma^{(i)} & \text{otherwise} \end{cases}, \quad \theta_i := \max\{z \in \Gamma \mid z < \gamma^{(i)}\} = \nu\gamma^{(i)}, \quad i \in \mathbb{N}_0.$$

Then $\lim_{i\to\infty} \delta_i = \lim_{i\to\infty} \gamma^{(i)} = 0$ and for at least one subsequence $(X^{(i_\ell)})_{\ell\in\mathbb{N}_0}$, there exists a sequence of stationary points $(S_\ell)_{\ell\in\mathbb{N}_0}$, $S_\ell \in \mathcal{S}^*_{\gamma^{(i_\ell)}}$, with $\|S_\ell - X^{(i_\ell)}\| \to 0$.

We note that for instance the set Γ could be modified as long as it remains bounded with its one accumulation point at $\inf(\Gamma) = 0$. Likewise, one may choose various other rules for the adaption of $(\gamma^{(i)})_{i\in\mathbb{N}_0}$ that adhere to the necessary relation between δ_i and θ_i.

Proof We first assume that $\delta_{\inf} := \liminf_{i\to\infty} \delta_i > 0$. There are hence only finitely many steps with $\delta_i < \frac{1}{2}\delta_{\inf}$ and it follows that $\gamma^* := \lim_{i\to\infty} \gamma^{(i)} > 0$. Further, as $(\gamma^{(i)})_{i\in\mathbb{N}_0} \subset \Gamma$, there necessarily exists an $N \in \mathbb{N}$ such that $\gamma^{(i)} = \gamma^*$ for all $i \geq N$. Thus, there exists a subsequence $(X^{(i_\ell)})_{\ell\in\mathbb{N}_0}$, $i_\ell \geq N$, $\ell \in \mathbb{N}_0$, for which

$$\|X^{(i_\ell)} - S\|_F \geq \frac{1}{2}\delta_{\inf} > 0, \tag{12}$$

for all $\ell \in \mathbb{N}_0$ and all $S \in S^*_{\gamma^*}$. By Assumption 7 however, $X^{(i)}$ remains bounded and $(X^{(i_\ell)})_{\ell\in\mathbb{N}_0}$ must have an accumulation point, which is within $S^*_{\gamma^*}$. This is in direct contradiction to (12). Thus, $\delta_{\inf} = 0$ must instead hold true, and thereby $\gamma^* = 0$. By construction, there hence exist subsequences $(X^{(i_\ell)})_{\ell\in\mathbb{N}_0}$ (given through the steps in which $\gamma^{(i_\ell+1)} < \gamma^{(i_\ell)}$) as well as $\{S^{(\ell)}\}_{\ell\in\mathbb{N}_0}$, $S_\ell \in \mathcal{S}^*_{\gamma^{(i_\ell)}}$, such that

$$\|X^{(i_\ell)} - S_\ell\|_F \leq 2\delta_{i_\ell} \leq 2\theta_{i_\ell}, \quad \ell \in \mathbb{N}_0.$$

As $\theta_{i_\ell} \to 0$ follows by $\gamma^* = 0$, we obtain $\|X^{(i_\ell)} - S_\ell\| \to 0$. This was to be shown. ■

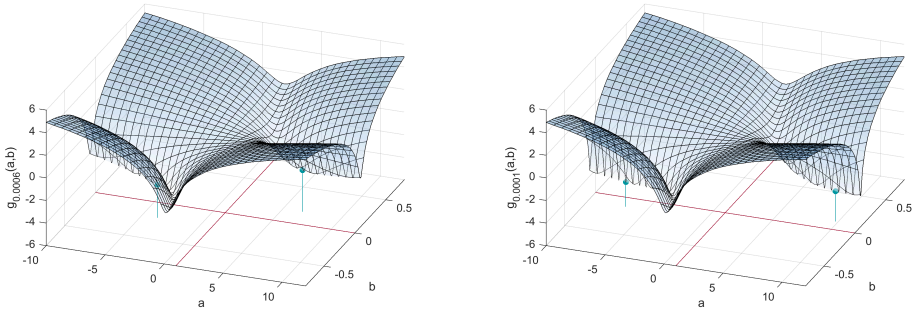

Figure 3: The functions $g_\gamma(a,b) := f_\gamma(X(a,b))$ for $\gamma \in \{0.0006, 0.0001\}$ as defined through the setting in Lemma 10 together with their respective minima.

3.2 Divergent Iterates in Presence of Rank Minimizers at Infinity

While it empirically seems neglectable, as further supported by Section 2.2, the following Lemma 10 shows that the expected conditions in Assumption 7 suffice to imply the existence of settings in which not only the global minimizers but also $(X^{(i)})_{i\in\mathbb{N}_0}$ diverges for $\lim_{i\to\infty}\gamma^{(i)} = 0$. This is different for $p > 0$, as in that case, the iterates and $(S_{\gamma^{(i)},p}(X^{(i)}))_{i\in\mathbb{N}_0}$ always remain bounded. However, as highlighted in Proposition 5 and Example 6, the log-det minimizing iterates are still successful in minimizing the product of singular values and thereby, while diverging, approach even lower-rank matrices.

Lemma 10. *Let, as in Example 6, $\mathcal{L} : \mathbb{R}^{2\times 3} \to \mathbb{R}^4$ be a linear operator and $y \in \mathbb{R}^4$ with*

$$\mathcal{L}^{-1}(y) = \{X(a,b) \mid a,b \in \mathbb{R}\}, \quad X(a,b) := \begin{pmatrix} a & 1 & a+b \\ b+1 & b & b+1 \end{pmatrix}.$$

Then for every weakly decreasing sequence $(\gamma^{(i)})_{i\in\mathbb{N}_0} \subset \mathbb{R}_{\geq 0}$ with $\gamma^* = \lim_{i\to\infty}\gamma^{(i)} = 0$ (and $f_{\gamma^{(0)}}(X^{(0)}) < \log(2)$),[3] the sequence $(X^{(i)})_{i\in\mathbb{N}_0}$ defined by an algorithm that satisfies Assumption 7 diverges with $\|X^{(i)}\| \to \infty$ while $f_{\gamma^{(i)}}(X^{(i)}) \to -\infty$ (cf. Figure 3).

Since for $\gamma > 0$, we have that $f_\gamma(X) \to \infty$ for $\|X\| \to \infty$, the function f_γ must have global minima. While these are difficult to state algebraically, we know that the (real valued) stationary points other than those converging to $X(1,0)$ must likewise diverge along the valleys of f_0 as visualized in Figure 3.

3.3 Artificial Fixed Points through Singular Value Based Adaption

A common, singular value based strategy for the decrease of γ as for instance by Daubechies et al. (2010) and Fornasier et al. (2011) here translates to[4]

$$\gamma^{(i)} = \min(\gamma^{(i-1)}, \alpha\sigma_{K+1}(X^{(i)})^2),$$

3. This assumption just ensures that the iterate does not accidentally fall into the one saddle point of f_0.
4. Note that the parameter γ indeed corresponds to squared singular values as opposed to smoothings such as $\sigma^2 + \varepsilon^2$.

for some $\alpha > 0$ and a sufficiently small bound $K \in \mathbb{N}$ on the minimal rank $r^* \leq K$. While of both practical and theoretical interest, there are some aspects to be aware of. No matter which value α takes, the strategy may introduce artificial fixed points, or more accurately saddle points, which are not intrinsic to the problem setting and independent of the applied algorithm. Lemma 11 considers this in detail for the same case as in Lemma 4. It can thus be advantageous to not only passively adapt γ as with this strategy, but to also actively ensure that the parameter converges to zero. Opposingly, a too small value of α can cause local minima to appear too early due to increasing non-convexity for $\gamma \searrow 0$ (for $p \neq 1$) or slow down the optimization too much (though the latter aspect may depend on the chosen algorithm, cf. Appendix B). Further, while knowing or guessing r^* in advance may yet be considered realistic in the matrix case, this would constitute a major level of prior knowledge for tensors (cf. Krämer, 2021a), where the generalized rank has multiple components (cf. Grasedyck et al., 2013; Lathauwer, 2009).

Lemma 11. *Given the setting in Lemma 4, let Λ be the set of tuples $(\Delta_\delta, \alpha, \gamma, a^*) \in (0,1)^4$ for which $\alpha\sigma_2(X_\delta(a^*))^2 = \gamma$, while $X_\delta(a^*)$ minimizes $f_\gamma|_{\text{image}(X_\delta)}$. Then*

$$\kappa^* := \max\left\{\kappa \geq 0 \;\middle|\; \nexists(\Delta_\delta, \alpha, \gamma, a^*) \in \Lambda : \alpha \leq \kappa\sqrt{\Delta_\delta}\right\} = \frac{3\sqrt{3}}{2}.$$

Further, for all $\varepsilon > 0$, it is

$$\nu^* := \max\left\{\nu \geq 0 \;\middle|\; \nexists(\Delta_\delta, \alpha, \gamma, a^*) \in \Lambda : \alpha \leq \nu \;\wedge\; |1 - a^*| \leq \frac{1}{1+\varepsilon}\right\} = \varepsilon.$$

In particular, these are well defined, and do not depend on δ, but only $\Delta_\delta := \|\delta\|_F < 1$.

In plain, $\alpha \leq \kappa^*\sqrt{\Delta_\delta}$ is the weakest bound (within the class $\kappa\Delta_\delta^s$ for $\kappa, s \geq 0$) to guarantee that the given strategy to adapt γ does not introduce artificial fixed points. However, if only matrices with distance $\|X_\delta(a) - X_\delta(1)\| \leq (1+\varepsilon)^{-1}$ are considered, then this condition becomes the less restrictive $\alpha \leq \nu^*\Delta_\delta^0 = \varepsilon$.

4. Log-det Iteratively Reweighted Least Squares (IRLS-0) and IRLS-p

A second step to take in the numerical realization of asymptotic log-det minimization is to determine a feasible updating function Ψ_γ. While there are various to choose from, we here consider the most basic (one-sided) IRLS-0 method (Algorithm 2 for $p = 0$), because it is simple yet qualitatively representative. We first recapitulate a derivation in Section 4.1 and then state the algorithm in Section 4.2. This method has further along with its mirrored version arguably been most influential in literature. Section 4.3 considers a specific example to highlight the fundamental role of the sequence $(\gamma^{(i)})_{i\in\mathbb{N}_0}$ in connection to the two questions posed in Section 1.3. A further case is analyzed in Appendix B.

4.1 Derivation of IRLS-0 through Majorizing Linearization

The IRLS-0 algorithm can naturally be derived (cf. Candès et al., 2013; Fazel et al., 2003) based on the observation that on the set of positive-definite matrices $M \in \mathbb{R}^{n\times n}$, $M \succ 0$, the function $\log\det(M + \gamma I)$ is strictly concave for every $\gamma > 0$ (cf. Davis, 1957; Rivin,

2002). The thus majorizing linearization in explicit yields that for every $N \in \mathbb{R}^{n\times n}$ with $M \neq N \succ 0$, it holds (see for instance the thesis by Kümmerle, 2019)

$$\begin{aligned} g_\gamma(M) &< g_\gamma(N) + \langle \nabla g_\gamma(N), M - N\rangle \\ &= \operatorname{trace}((N+\gamma I)^{-1}(M+\gamma I)) - \log\det(N+\gamma I)^{-1} - n. \end{aligned}$$

Substitution of $M := XX^T$ as well as $W := (N+\gamma I)^{-1}$ then directly leads to

$$f_\gamma(X) < J_\gamma(X, W) := \|W^{1/2}X\|_F^2 + \gamma\|W^{1/2}\|_F^2 - \log\det(W) - n,$$

whereas equality holds if and only if $W = W_{\gamma,X} := XX^T + \gamma I$. The IRLS-0 algorithm (5) is on that basis frequently viewed as alternating minimization of J_γ, whereas the minimizer in X is given as (cf. Fornasier et al. (2011); Mohan and Fazel (2012))

$$X_W := \operatorname*{argmin}_{X\in\mathcal{L}^{-1}(y)} \|W^{1/2}X\|_F = \mathcal{W}^{-1}\circ\mathcal{L}^*\circ(\mathcal{L}\circ\mathcal{W}^{-1}\circ\mathcal{L}^*)^{-1}(y), \tag{13}$$

for $\mathcal{W}^{-1}(X) := W^{-1}X$, where $\mathcal{L}^*$ is the adjoint of $\mathcal{L}$. While this perspective is convenient, the thesis by Kümmerle (2019) and a work by both Kümmerle and Verdun (2021) have shown that next to the opposing choice $M := X^TX$, a generalization of the derivation above can be used to establish both-sided weights (cf. Section 5.1.2). These allow for a drastically faster decrease of the sequence $(\gamma^{(i)})_{i\in\mathbb{N}_0}$, and under certain conditions even exhibit a local, quadratic convergence rate. Due to simplifications that for instance follow for tensor network based alternating IRLS-p (Krämer, 2021a), one-sided weights can nevertheless be of interest. While we consider only $p = 0$, the following Corollary 12 is easily shown to hold true for general $p \in [0, 1]$. It also holds true for the aforementioned both-sided weights.

Corollary 12. *Let $X^+ := X_{W_{\gamma,X}}$. For each $t \in [0, 2]$ and $0 < \gamma^+ \leq \gamma$, it holds true that $f_{\gamma^+}((1-t)X + tX^+) \leq f_\gamma(X)$.*

Proof It is $f_{\gamma^+}((1-t)X + tX^+) \leq f_\gamma((1-t)X + tX^+) \leq J_\gamma((1-t)X + tX^+, W_X) \leq J_\gamma(X, W_X) = f_\gamma(X)$, where the last inequality follows from (13). ■

In other words, the entire line between the initial and updated matrix (which is contained in $\mathcal{L}^{-1}(y)$) is an improvement. Considering this localized optimization property helps to further understand why slow progress can be observed when traversing a curved, steep valley of the objective function, in particular for ones-sided weights (see Example 20).

4.2 Summarized Algorithm

Similar to the indirect optimization concerning f_γ for $p = 0$ as described in Section 4.1, a different weight strength $p \in (0, 1]$ can be applied to instead minimize the smoothed Schatten-p functions $S_{\gamma,p}$ defined in (6). We include general IRLS-p (Algorithm 2) for later comparison in the numerical experiments (cf. Section 5.2.1).

Algorithm 2 (one-sided, matrix) IRLS-p

1: set $p \in [0,1]$, $X^{(0)} \in \mathcal{L}^{-1}(y)$ and $\gamma^{(0)} > 0$ (cf. Proposition 14)
2: **for** $i = 1, 2, \ldots$ **do**
3: $\quad W^{(i-1)} := (X^{(i-1)} X^{(i-1)^T} + \gamma^{(i-1)} I)^{p/2-1}$
4: $\quad X^{(i)} := \underset{X \in \mathcal{L}^{-1}(y)}{\operatorname{argmin}} \|W^{(i-1)^{1/2}} X\|_F$ (cf. (13))
5: $\quad$ set $\gamma^{(i)} \leq \gamma^{(i-1)}$ according to chosen strategy
6: **end for**

Depending on a weakly decreasing sequence $(\gamma^{(i)})_{i \in \mathbb{N}_0} \subset \mathbb{R}_{>0}$, Algorithm 2 defines a sequence of matrix iterates and according weights $\big((X^{(i)}, W^{(i)})\big)_{i \in \mathbb{N}_0}$. This sequence fulfills, as is well-known, the properties posed by Assumption 7.

Remark 13. As directly presented in the works by Fornasier et al. (2011) as well as Mohan and Fazel (2012) or contained within the proofs of such, the statements in Assumption 7 hold true for IRLS-0 (and similar ones do for IRLS-p). Thereby, also the statements in Proposition 9 and Lemma 10 concern that algorithm. The same holds true for Lemma 11.

Contrary to the asymptotic minimization based on $\gamma \searrow 0$, a uniquely determined, canonical starting value is given through the limit case $\gamma \to \infty$ in Proposition 14 that serves as initial, auxiliary loss function. In particular, a large enough starting value for γ avoids to fall into one of the local minima that are expected to arise eventually for small values.

Proposition 14. *Independently of $X^{(0)} \in \mathcal{L}^{-1}(y)$, it holds*

$$\lim_{\gamma \to \infty} \operatorname*{argmin}_{X \in \mathcal{L}^{-1}(y)} f_\gamma(X) = \lim_{\gamma \to \infty} X_{W_{\gamma, X^{(0)}}} = \operatorname*{argmin}_{X \in \mathcal{L}^{-1}(y)} \|X\|_F,$$

where the first limit is possibly a set based convergence.

Proof The second equality follows since the weight matrix takes the role of a scaled identity for $\gamma \to \infty$. That is, for the angle it holds $\measuredangle(W_{\gamma, X^{(0)}}, I_n) \to 0$. Let therefore $X_\infty = \operatorname{argmin} \|X\|_F$ subject to $\mathcal{L}(X) = y$. For $\gamma \geq 1$, let further $X_\gamma = \operatorname{argmin} f_\gamma(X)$ subject to $\mathcal{L}(X) = y$. For any fixed X, it holds true that $\gamma^n + \gamma^{n-1}\|X\|_F^2 \leq \exp(f_\gamma(X)) \leq \gamma^n + \gamma^{n-1}\|X\|_F^2 + \gamma^{n-2} c_X$ (cf. (4)). Thus, for some fixed $c_{X_\infty} \geq 0$, it follows that $\|X_\gamma\|_F^2 \leq \gamma^{1-n} \exp(f_\gamma(X_\gamma)) - \gamma \leq \|X_\infty\|_F^2 + c_{X_\infty} \gamma^{-1}$. Since both $X_\gamma - X_\infty \in \operatorname{kernel}(\mathcal{L}) \perp X_\infty$, we then obtain $\langle X_\gamma - X_\infty, X_\gamma - X_\infty \rangle = \langle X_\gamma - X_\infty + 2X_\infty, X_\gamma - X_\infty \rangle = \langle X_\gamma, X_\gamma \rangle - \langle X_\infty, X_\infty \rangle = \|X_\gamma\|_F^2 - \|X_\infty\|_F^2 \leq c_{X_\infty} \gamma^{-1}$. It thereby follows that $\lim_{\gamma \to \infty} X_\gamma = X_\infty$. ∎

4.3 Stagnation and Convergence to Arbitrary Points

The following Example 15 discusses in detail, with regard to the questions posed in Section 1.3, to which far reaching extent an unsuitable discretization of the asymptotic process $\gamma \searrow 0$ can disguise itself. Namely, by a specific, too fast decrease of that parameter, the iterates can be caused to converge to a selected but essentially arbitrary point, which in that

sense may appear like a local minimum. A further setting for IRLS-0 is empirically studied in Appendix B. Apart from the specific observations considered here, it seems challenging if not impossible in general practice to definitely determine if one is near an asymptotic minimizer or even stationary point. Empirically, some insight can be gained though through sensitivity tests (cf. Section 5.1.3).

Example 15. Let $\mathcal{L} : \mathbb{R}^{2\times 2} \to \mathbb{R}^2$ be a linear operator and $y \in \mathbb{R}^2$ such that

$$\mathcal{L}^{-1}(y) = \{X(a,b) \mid a,b \in \mathbb{R}\}, \quad X(a,b) := \begin{pmatrix} a & 1 \\ 1 & b \end{pmatrix}.$$

The functions $g_\gamma : \mathbb{R}^2 \to \mathbb{R}$, $g_\gamma(a,b) := f_\gamma(X(a,b))$, are visualized in Figure 4. The matrix $X(a,b)$ has rank 1 if and only if $ab = 1$. While there is not a unique solution to the rank minimization problem, we have

$$\mathcal{X}^* = \operatorname*{argmin}_{X\in\mathcal{L}^{-1}(y),\ \operatorname{rank}(X)=1} \sigma_1(X) = \left\{ \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}, \begin{pmatrix} -1 & 1 \\ 1 & -1 \end{pmatrix} \right\}.$$

The only stationary points of f_γ, $\gamma \in (0,1]$, in turn are given by

$$\mathcal{S}^*_\gamma := \left\{ \begin{pmatrix} -\sqrt{1-\gamma} & 1 \\ 1 & -\sqrt{1-\gamma} \end{pmatrix}, \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \begin{pmatrix} \sqrt{1-\gamma} & 1 \\ 1 & \sqrt{1-\gamma} \end{pmatrix} \right\},$$

where the second one is repellent. The relation $X^{(i)} = X(a_i, b_i) = X_{W_{\gamma, X^{(i-1)}}}$, $i \in \mathbb{N}$, is satisfied by rational functions $a_i = q_1(a_{i-1}, b_{i-1})$ and $b_i = q_2(a_{i-1}, b_{i-1})$ for

$$q_1(a,b) = \frac{a+b}{1+\gamma+b^2}, \quad q_2(a,b) = \frac{a+b}{1+\gamma+a^2}.$$

Short calculations then show that $0 < q_1(a,b) \cdot q_2(a,b) \leq 1$ is equivalent to $a + b \neq 0$, whereas both $q_1(a,b) = 0$ and $q_2(a,b) = 0$ are equivalent to $a + b = 0$. Thus, the stationary point $a = b = 0$ is either reached directly or never. Further, for any $a, b \in \mathbb{R}$, it holds true that

$$|q_1(a,b) - q_2(a,b)| \leq |a - b|, \tag{14}$$

where equality requires $a = b$, or both $\gamma = 0$ and $ab = 1$. The only attracting fixed points $(a,b) = (q_1(a,b), q_2(a,b))$ are consequently given through the corresponding[5] stationary points $\mathcal{S}^*_\gamma$ for $\gamma \in (0,1]$, as well as through all singular matrices $X(a,b)$, $ab = 1$, for $\gamma = 0$. The behavior of $(X^{(i)})_{i\in\mathbb{N}_0}$ now depends on the sequence $(\gamma^{(i)})_{i\in\mathbb{N}_0}$. Without loss of generality, we assume $0 < a_1 b_1 \leq 1$ and $a_1 \geq b_1$ in the following three cases:

1) For $\gamma^* := \lim_{i\to\infty} \gamma^{(i)} > 0$, the fact that Assumption 7 holds true yields that $X^{(i)}$ converges to one of the two attracting fixed points in $\mathcal{S}^*_{\gamma^*}$.

5. The fact that here the lowest norm solution (cf. Proposition 14) is a local maximum of f_γ, $0 < \gamma < 1$, and the global minimum for $\gamma \geq 1$, is not representative of the general situation.

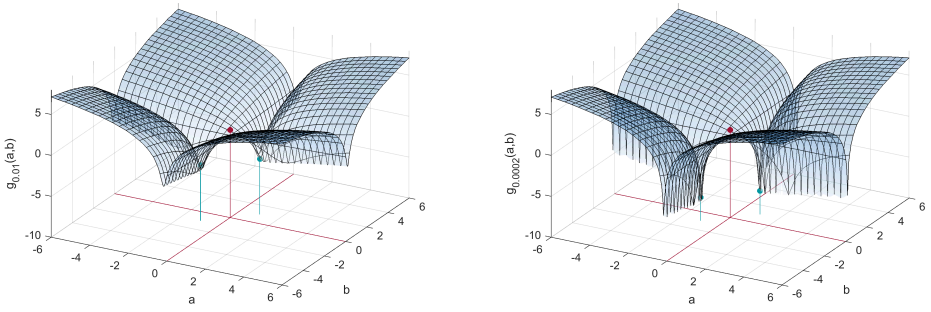

Figure 4: The functions $g_\gamma(a,b) := f_\gamma(X(a,b))$ for $\gamma \in \{0.01, 0.0002\}$ as defined through the setting in Example 15 together with their common saddle point and the respective, numerically calculated global minima.

2) For $\gamma = 0$, we have that $0 < ab < 1$ implies both $|q_1(a,b)| > |a|$ and $|q_2(a,b)| > |b|$. Thus, for $\gamma^{(i)} \equiv 0$, the sequence $(X^{(i)})_{i\in\mathbb{N}_0}$ converges from below to a rank-one matrix with

$$a_1^2 + b_1^2 + 2 \leq \left\| \lim_{i\to\infty} X^{(i)} \right\|_F^2 \leq (a_1 - b_1)^2 + 4,$$

where the second inequality follows due to (14) and $a_i b_i \leq 1$, $i \in \mathbb{N}$. The norm of the limit can thus be arbitrarily large, yet here a single sequence never diverges. Further, (a_1, b_1) already determines if the limit may be in $\mathcal{X}^*$.

3) Given $\gamma^* := \lim_{i\to\infty} \gamma^{(i)} = 0$, the rate of decrease is deciding.

i) Let $1 < t < a$ (or for that matter, $0 < b < \frac{1}{t}$) as well as $\mu > 0$. Then given

$$0 < \gamma = \Gamma(a,b) := \mu \cdot (a/t - 1 + b/t - b^2),$$

the inequality $q_1(a,b) > t$ is equivalent to $\mu < 1$. Thus, for fixed $0 < \mu < 1 < t < a_1$ and $\gamma^{(i)} := \min(\gamma^{(i-1)}, \Gamma(a_i, b_i))$, it follows with (14) that $a_i \to t$ and $b_i \to \frac{1}{t}$ as well as $\gamma^{(i)} \searrow 0$. In particular, this happens independently so of (a_1, b_1).

ii) On the other hand,

$$\gamma^{(i)} := \min(\gamma^{(i-1)}, \sqrt{1 - \gamma^{(i-1)}} - |b_{i-1}|)$$

will always yield a sequence for which $X^{(i)}$ converges to a point in $\mathcal{X}^*$ (cf. Proposition 9). Certainly, this exact strategy would require knowledge of this set in the first place.

5. Numerical Experiments

As mentioned in the introduction, Section 1, we are mainly interested in the ARM problem itself, for which the criteria for success are as laid out in Section 5.1.3. For simplicity, the

mode sizes n and m are chosen equal as $\overline{n} \in \mathbb{N}$. Further, we briefly consider the analogous ACM problem (cf. Section 1.4) in Appendix C. The MATLAB code behind all results is available as public repository under the name `a-irls`. A comparison with other approaches to ARM and ACM is intentionally not made here. Likewise, the emphasis does not lie on computational demands but a further investigation into the properties of asymptotic log-det minimization and its feasibility via IRLS-0. An overview over the numerical experiments as well as corresponding results and further visualizations can be found in Appendix F.

5.1 Benchmark Design

The following subsections clarify several aspects surrounding the numerical experiments.

5.1.1 Reference Solutions, Measurements Vectors and Operators

Each measurement vector is constructed via a (not necessarily sought for) rank $r_{\mathrm{rs}} \in \mathbb{N}$ reference solution, which in turn relies on a randomly generated low-rank decomposition,

$$y = \mathcal{L}(X^{(\mathrm{rs})}) \in \mathbb{R}^{\ell}, \quad X^{(\mathrm{rs})} = Y^{(\mathrm{rs})} Z^{(\mathrm{rs})} \in \mathbb{R}^{n \times m}, \quad Y^{(\mathrm{rs})} \in \mathbb{R}^{n \times r_{\mathrm{rs}}},\ Z^{(\mathrm{rs})} \in \mathbb{R}^{r_{\mathrm{rs}} \times m}.$$

All entries of the two components $Y^{(\mathrm{rs})}$ and $Z^{(\mathrm{rs})}$ are assigned independent, normally distributed entries. For the operator $\mathcal{L}$, we distinguish between two types:

Gaussian Measurements With Gaussian measurements, we refer to $\mathcal{L}(X) := L\ \mathrm{vec}(X)$, generated via a random matrix $L \in \mathbb{R}^{\ell \times nm}$ with independent, normally distributed entries.

Random Sampling Operator As sampling operator, we denote $\mathcal{L}(X) := (X_{p_i})_{i=1}^{\ell}$, for uniformly drawn indices $\{p_1, \dots, p_\ell\} \subset \{1, \dots, n\} \times \{1, \dots, m\}$. We however repeat the random number generation until at least r_{rs} entries of each column and each row of the corresponding matrix structure are observed.

5.1.2 Solution Methods

Based on a sufficiently large starting value $\gamma^{(0)} > 0$, we by default choose $\gamma^{(i)} = \nu \gamma^{(i-1)}$, $i \in \mathbb{N}$, where $\nu < 1$ remains constant throughout each single run of an algorithm. If not otherwise specified, the default weight strength, as it is our main interest, is given through $p = 0$. Note that due to the convex nature for $p = 1$, IRLS-1 always finds the one nuclear norm minimizer $\operatorname{argmin}_{X \in \mathcal{L}^{-1}(y)} \|X\|_*$ if only the sequence $(\gamma^{(i)})_{i \in \mathbb{N}_0}$ declines to zero reasonably slow. We consider the following two types of optimization.

*One-sided Weights (*IRLS-p*)* As in (13), the full matrix is optimized based on the (literally interpreted) image update formula. When instability threatens to occur, an equivalent kernel based, though more time consuming update formula is used.

Both-sided Weights (only stated for $p = 0$) Both-sided weights (Kümmerle and Verdun, 2021) are instead applied. For $p = 0$, this corresponds to updates

$$X^{(i)} := \operatorname*{argmin}_{X \in \mathcal{L}^{-1}(y)} \|W_L^{(i-1)^{1/2}} X W_R^{(i-1)^{1/2}}\|_F,$$

for $W_L^{(i-1)} := (X^{(i-1)}X^{(i-1)^T} + \gamma^{(i-1)}I)^{-1/2}$ and $W_R^{(i-1)} := (X^{(i-1)^T}X^{(i-1)} + \gamma^{(i-1)}I)^{-1/2}$. Though the work by Kümmerle and Verdun (2021) also states a more efficient, data-sparse calculation for the sampling case, we here settle for the simplified version as we do not compare computational demands. Likewise, it should be noted that while a singular value based adaption is one of the prerequisites for the local quadratic convergence of this modification, it is still beneficial for other strategies of how to adapt γ (cf. Figure 5).

5.1.3 Experimental Setup and Evaluation

The reference solution $X^{(\mathrm{rs})}$ is not necessarily the sought for solution. So in order to evaluate the output $X^{(\mathrm{alg})}$, we compare their *non-neglectable* singular values. We therefore define (cf. Theorem 1)

$$\det\nolimits^2_{n,\gamma,\varepsilon}(X) := \gamma^{n-\mathrm{rank}_\varepsilon(X)} \prod_{j=1}^{\mathrm{rank}_\varepsilon(X)} (\sigma_j(X)^2 + \gamma),$$

for $\mathrm{rank}_\varepsilon(X) := \max\{j \in \{1,\dots,n\} \mid \sigma_j(X) > \varepsilon \cdot \|X\|_F\}$, $\varepsilon := 10^{-6}$. We firstly compare the approximate ranks, and secondly compare the products of singular values. These two aspects are reflected by the limit

$$\mathcal{Q}_\varepsilon(X^{(\mathrm{alg})}, X^{(\mathrm{rs})}) := \lim_{\gamma \searrow 0} \frac{\det_{n,\gamma,\varepsilon}(X^{(\mathrm{alg})})}{\det_{n,\gamma,\varepsilon}(X^{(\mathrm{rs})})} \in [0, 0.98] \cup (0.98, 1.005) \cup [1.005, \infty].$$

The three intervals are related to the categorization into improvements, successes or failures as outlined below, whereas the limits 0 or ∞ are reached if and only if $\mathrm{rank}_\varepsilon(X^{(\mathrm{alg})})$ and $\mathrm{rank}_\varepsilon(X^{(\mathrm{rs})})$ differ.

Post-Processing As the truncation of minor singular values is a numerical necessity, additional care has been taken when this process may falsify the results. Therefore, we alternatingly and sufficiently often project any $X^{(\mathrm{alg})}$ in question to the sets $\mathcal{L}^{-1}(y)$ and $V_{\leq r_{\mathrm{rs}}}$, where r_{rs} is then chosen as rank of $X^{(\mathrm{rs})}$. This usually allows us to reduce the parameter ε to machine precision given the close[6] starting value $X^{(\mathrm{alg})}$. The insight following this process is not included in the presented tables but each in the discussion thereafter.

Details of Comparison If the quotient satisfies $\mathcal{Q}_\varepsilon(X^{(\mathrm{alg})}, X^{(\mathrm{rs})}) = \infty$, then the result is considered a *strong failure*, whereas $1.005 \leq \mathcal{Q}_\varepsilon(X^{(\mathrm{alg})}, X^{(\mathrm{rs})}) < \infty$ is referred to as *weak failure*. Otherwise, for $0.98 < \mathcal{Q}_\varepsilon(X^{(\mathrm{alg})}, X^{(\mathrm{rs})}) < 1.005$, we consider the result *successful*, while for $\mathcal{Q}_\varepsilon(X^{(\mathrm{alg})}, X^{(\mathrm{rs})}) \leq 0.98$, we say the result is an *improvement*, subject to the consideration above. A more distinguished illustration is presented as explained in Section 5.1.4.

Sensitivity Analysis For each instance, we lower the meta parameter $\nu = \nu_k = \sqrt{\nu_{k-1}}$ (cf. Section 5.1.2), starting with $\nu_0 = 1.2$, and rerun the respective algorithm from the start until the result is not a *failure*, or, if after too many reruns $k > k_{\max}$, we give up and thus either achieve a *weak* or *strong failure* depending on the result for $k = k_{\max}$. Considering

6. While this process is suitable for already approximately low-rank matrices, it does not work well as a standalone algorithm, apart from the fact that the rank is generally unknown.

the canonical, equal starting values, instances that failed only for some $k > k_{\max}$ thus did so most likely due to a too fast decrease of $\gamma \searrow 0$, and not due to local minima.

All other meta parameters for each algorithm are common to all respective experiments.

5.1.4 Presentation of Results

Each experiment is reflected upon in three different ways as summarized in Table 6.

ACM/ARM/Recovery Tables For each instance, we list the percentual numbers of ACM or ARM improvements, successes or failures as defined in Section 5.1.3. Successes are further distinguished regarding recoveries, that is whether $\|X^{(\mathrm{alg})} - X^{(\mathrm{rs})}\|_F \leq 10^{-4}\|X^{(\mathrm{rs})}\|_F$. In nearly all cases where this is fulfilled, the relative residual even falls below 10^{-6} (see Appendix F), in which case the algorithm stops automatically.[7]

ACM/ARM/Recovery Figures (Button Plots) We display the following points as clustered scatter plot (see *button plot* by Krämer, 2021b as well as Appendix E) logarithmic in the y-axis. Given the i-th result $X^{(\mathrm{alg})}$ as well as reference solution $X^{(\mathrm{rs})}$, the initial (x, y)-value of the i-th disk is given by (cf. Section 5.1.3)

$$x_i = \max\left(0.9, \min\left(\mathcal{Q}_\varepsilon(X^{(\mathrm{alg})}, X^{(\mathrm{rs})}), 1.05\right)\right), \quad y_i = \min\left(\frac{\|X^{(\mathrm{alg})} - X^{(\mathrm{rs})}\|_F}{\|X^{(\mathrm{rs})}\|_F}, 1\right).$$

Note that the algorithm stops automatically if the latter value remains below 10^{-6}.

γ-decrease Sensitivity To each single trial that did not yield a failure, we assign the one index k for which the parameter $\nu = \nu_k$ first led to a successful or improving run as described in Section 5.1.3. The frequencies of these indices as well as failures are then plotted as bars, where improvements are plotted below the x-axis.

5.2 Benchmarks and Selection of Results

The following subsections as well as Appendices C and D present larger scale benchmarks as well as, mainly, corresponding tabular results. Further illustrations of results can be found in subsequent sections.

5.2.1 Affine Rank Minimization Depending on p

Experiment 16. For $\overline{n} = 12$, $r_{\mathrm{rs}} = 3$ and $\ell \in \{63, 72, 108\}$, we consider the ARM problem based on *Gaussian measurements* under varying choice of the weight strength $p \in \{0, 0.2, 0.4, 0.6, 0.8, 1\}$. Each constellation is repeated 1000 times for $k_{\max} = 12$. The results are covered in Table 1 and Figures 10 and 11.

The degrees of freedom of at most rank-r_{rs} square matrices of size $\overline{n}$ is $\dim(V_{\leq r_{\mathrm{rs}}}) = 2\overline{n}r_{\mathrm{rs}} - r_{\mathrm{rs}}^2$, which in Experiment 16 results in $\dim(V_{\leq r_{\mathrm{rs}}}) = 63$. For $\ell = 63$, IRLS-p (at least for small p) quite often yields improvements (which is surprisingly different to the vector

7. This is the only point at which the reference solution itself is used within the algorithm, and only done in order to save computation time.

instance	ARM: / recovery:	$\mathcal{Q}_\varepsilon \in [0, 0.98]$ no	$\mathcal{Q}_\varepsilon \in (0.98, 1.005)$ no		yes	$\mathcal{Q}_\varepsilon \in [1.005, \infty)$ no	$\mathcal{Q}_\varepsilon = \infty$
$\ell = 63$	• $p = 0.0$	15.9	4.9	+	31.9	11.2	36.1
	• $p = 0.2$	16.2	5.1	+	34	8.8	35.9
	• $p = 0.4$	15.9	4.6	+	32.2	5.6	41.7
	• $p = 0.6$	9.3	2.8	+	22.4	1.6	63.9
	• $p = 0.8$	0.3	0	+	0.2	0	99.5
	• $p = 1.0$	0	0	+	0	0	100
$\ell = 72$	• $p = 0.0$	0	0	+	99	0	1
	• $p = 0.2$	0	0	+	99	0	1
	• $p = 0.4$	0	0	+	98.5	0	1.5
	• $p = 0.6$	0	0	+	96.5	0	3.5
	• $p = 0.8$	0	0	+	49.8	0	50.2
	• $p = 1.0$	0	0	+	0	0	100
$\ell = 108$	• $p = 0.0$	0	0	+	100	0	0
	• $p = 0.2$	0	0	+	100	0	0
	• $p = 0.4$	0	0	+	100	0	0
	• $p = 0.6$	0	0	+	100	0	0
	• $p = 0.8$	0	0	+	100	0	0
	• $p = 1.0$	0	0	+	98.8	0	1.2

Table 1: matrix IRLS-p ($\overline{n} = 12$, $r_{\text{rs}} = 3$) – table as specified in Section 5.1.4 for Experiment 16 (see sensitivity analysis, Figure 10, and button plot, Figure 11, for more details).

instance	ARM: / recovery:	$\mathcal{Q}_\varepsilon \in [0, 0.98]$ no	$\mathcal{Q}_\varepsilon \in (0.98, 1.005)$ no		yes	$\mathcal{Q}_\varepsilon \in [1.005, \infty)$ no	$\mathcal{Q}_\varepsilon = \infty$
$\ell = 62$	• gaussian	46.2	26.5	+	0	20.8	6.5
	• gaussian (both sided)	46.2	32	+	0	17	4.8
	• sampling	20.7	10.3	+	0	7.2	61.8
	• sampling (both sided)	29.5	16.2	+	0	11	43.3
$\ell = 63$	• gaussian	16.7	7.1	+	34	11.6	30.6
	• gaussian (both sided)	18.1	5.9	+	43.2	10	22.8
	• sampling	10	6	+	10.1	1.4	72.5
	• sampling (both sided)	13.3	7.3	+	24.1	3.1	52.2
$\ell = 64$	• gaussian	0	0.2	+	58.2	0.1	41.5
	• gaussian (both sided)	0	0	+	62.9	0.1	37
	• sampling	7.4	7.2	+	22.9	0.6	61.9
	• sampling (both sided)	8.7	7.7	+	35.9	1.1	46.6

Table 2: matrix IRLS-0 ($\overline{n} = 12$, $r_{\text{rs}} = 3$) – table as specified in Section 5.1.4 for Experiment 17 (see sensitivity analysis, Figure 5, and button plot, Figure 6, for more details).

case as discussed in Appendix C), but despite the minimal number of measurements also often manages to reconstruct the reference solution. In particular the sensitivity analysis for $\ell = 72$ demonstrates (in accordance with results by Mohan and Fazel (2012)) that $p = 0$ seems to overall be the best choice in this setting, whereas the most significant differences can be observed between $p = 0.6$ and $p = 1$.

5.2.2 Observing the Theoretical Phase Transition for Matrix Recovery

Experiment 17. For $\overline{n} = 12$, $r_{\text{rs}} = 3$ and deviations $\ell \in \{\dim(V_{\leq r_{\text{rs}}}) + \Delta\ell \mid \Delta\ell \in \{-1, 0, 1\}\} = \{62, 63, 64\}$, we consider the ARM problem based on *samples* or *Gaussian measurements* for the weight strength $p = 0$ for *one-sided* or *both-sided* weights. Each constellation is repeated 1000 times for the increased value $k_{\max} = 14$ (with $\nu_{14} \approx 1.00001^{-1}$). The results are covered in Table 2 and Figures 5 and 6.

As in Experiment 16, we have $\dim(V_{\leq r_{\text{rs}}}) = 63$. From Table 2, we can clearly observe that $\ell = 62$ are as expected too few measurements to allow for any recovery.[8] In turn, the value

8. This is because no reference solution happens to minimize the product of singular values.

instance	ARM: recovery:	$\mathcal{Q}_\varepsilon \in [0, 0.98]$ no	$\mathcal{Q}_\varepsilon \in (0.98, 1.005)$ no		yes	$\mathcal{Q}_\varepsilon \in [1.005, \infty)$ no	$\mathcal{Q}_\varepsilon = \infty$
$\ell = 63$	• gaussian	15.9	4.9	+	31.9	11.2	36.1
	• gaussian (both sided)	18.1	5.9	+	43.1	10	22.9
	• sampling	8	4.3	+	5.6	0.9	81.2
	• sampling (both sided)	13.3	7.3	+	24	3	52.4
$\ell = 72$	• gaussian	0	0	+	99	0	1
	• gaussian (both sided)	0	0	+	99	0	1
	• sampling	2.4	2	+	77.1	0.2	18.3
	• sampling (both sided)	2.4	1.8	+	82.2	0.1	13.5
$\ell = 108$	• gaussian	0	0	+	100	0	0
	• gaussian (both sided)	0	0	+	100	0	0
	• sampling	0	0	+	100	0	0
	• sampling (both sided)	0	0	+	100	0	0

Table 3: matrix IRLS-0 ($\overline{n} = 12$, $r_{\mathrm{rs}} = 3$) – table as specified in Section 5.1.4 for Experiment 18 (see sensitivity analysis, Figure 12, and button plot, Figure 13, for more details).

$\ell = \dim(V_{\leq r_{\mathrm{rs}}}) + 1$ (which here is $\ell = 64$) is the minimal sufficient amount of *generic*[9] measurements (thus not including sampling) to provide $\mathcal{L}^{-1}(\mathcal{L}(X^{(\mathrm{rs})})) \cap V_{\leq r_{\mathrm{rs}}} = \{X^{(\mathrm{rs})}\}$ for *generic* $X^{(\mathrm{rs})} \in V_{\leq r_{\mathrm{rs}}}$, as more generally proven by Breiding et al. (2023). Indeed, there are only multiple solutions for $\ell \leq 63$, while for $\ell = 64$, the exceptions displayed in Table 2 for Gaussian measurements fail to withstand the post-processing (cf. Section 5.1.3) and thus do in fact not approximate truly rank-r_{rs} matrices. Note that the stationary points for one- and both-sided weights are equal, such that the differences in results originate from issues with the decrease $\gamma \searrow 0$ as observable in Figure 6.

5.2.3 Affine Rank Minimization Depending on Number of Measurements

Experiment 18. For $\overline{n} = 12$, $r_{\mathrm{rs}} = 3$ and $\ell \in \{63, 72, 108\}$, we consider the ARM problem based on *samples* or *Gaussian measurements* for the weight strength $p = 0$. Each constellation is repeated 1000 times for $k_{\max} = 12$. The results are covered in Table 3 and Figures 12 and 13.

Each the number of measurements corresponds to an oversampling factor $\ell / \dim(V_{\leq r_{\mathrm{rs}}})$ of roughly $1, 1.15$ and 1.72, respectively. Although the stationary points for one- and both-sided weights are equal, for fewer, in particular pointwise measurements, it can be observed that the rate of success for both-sided weights is higher. A closer inspection of Figure 13 shows that for the one-sided version, a yet slower decrease of γ is necessary. This can also be seen in Figure 6, which visualizes an overlap with Experiment 17 for $\ell = 63$, but allows for $k_{\max} = 14$. In case of sampling, it can further be observed that an oversampling factor of 1.15 is not yet enough to allow for a guaranteed recovery in the case of successful rank minimization.

6. Conclusions and Outlook

We have shown that asymptotic (global) minimization of the log-det functions for a smoothing parameter $\gamma \searrow 0$ solves the affine rank minimization problem. Though if the minimal rank is only attainable at infinity, then also IRLS-0 may correspondingly produce diverging

9. To be more precise, *generic* in that context is an algebraic property that is stronger than the ones that stem from analysis or probability theory, but roughly similar as detailed by Breiding et al. (2023).

iterates. Comparable solution statements do in turn not hold true for Schatten-p minimization, which is also reflected by the overall better numerical results for IRLS-0.

Multiple examples as well as numerical experiments suggest how an unfavorable numerical realization, especially the discretization of the asymptotic process $\gamma \searrow 0$, can produce misinterpretable results. Even under close observation, such may lead to an overestimation of the number of local minima inherent to the log-det approach. Whereas the adaption of γ can always be proceeded as to avoid these issues, it is not clear whether there also is a practically feasible method to do so. For instance, the common and quite successful singular value based control of the smoothing parameter can yet introduce artificial fixed points.

When care is taken to the aforementioned issues, IRLS-0 can exhibit a phase transition for generic recovery at the theoretical minimum. In suitable settings, and in particular given a sufficient number of measurements, the iterates then also recover a sought ground truth. While the approach with both-sided weights overall drastically reduces the necessity of a slow decrease of γ, it is not devoid of concerns. In particular with generalizations to tensor recovery in mind, this remains an important point to handle properly.

We have separated affine rank minimization from its application to matrix recovery to improve upon the overall interpretability of the data produced by numerical experiments as well as its underlying theoretical framework. A rigorous understanding of when one indeed finds the desired global minimizer however remains an interesting subject to future work.

Acknowledgments

The author would like to thank Tim A. Werthmann and Lars Grasedyck for fruitful discussions, as well as Maren Klever and Henrik Eisenmann in particular for proof reading, Morten Mørup for an early exchange on reweighted least squares, Christian Kümmerle for conversations on optimal weights, as well as Paul Breiding and Nick Vannieuwenhoven for exchanges on generic recoverability. The work of the author was supported by the DFG SPP-1886 under grant GR3179/5-2.

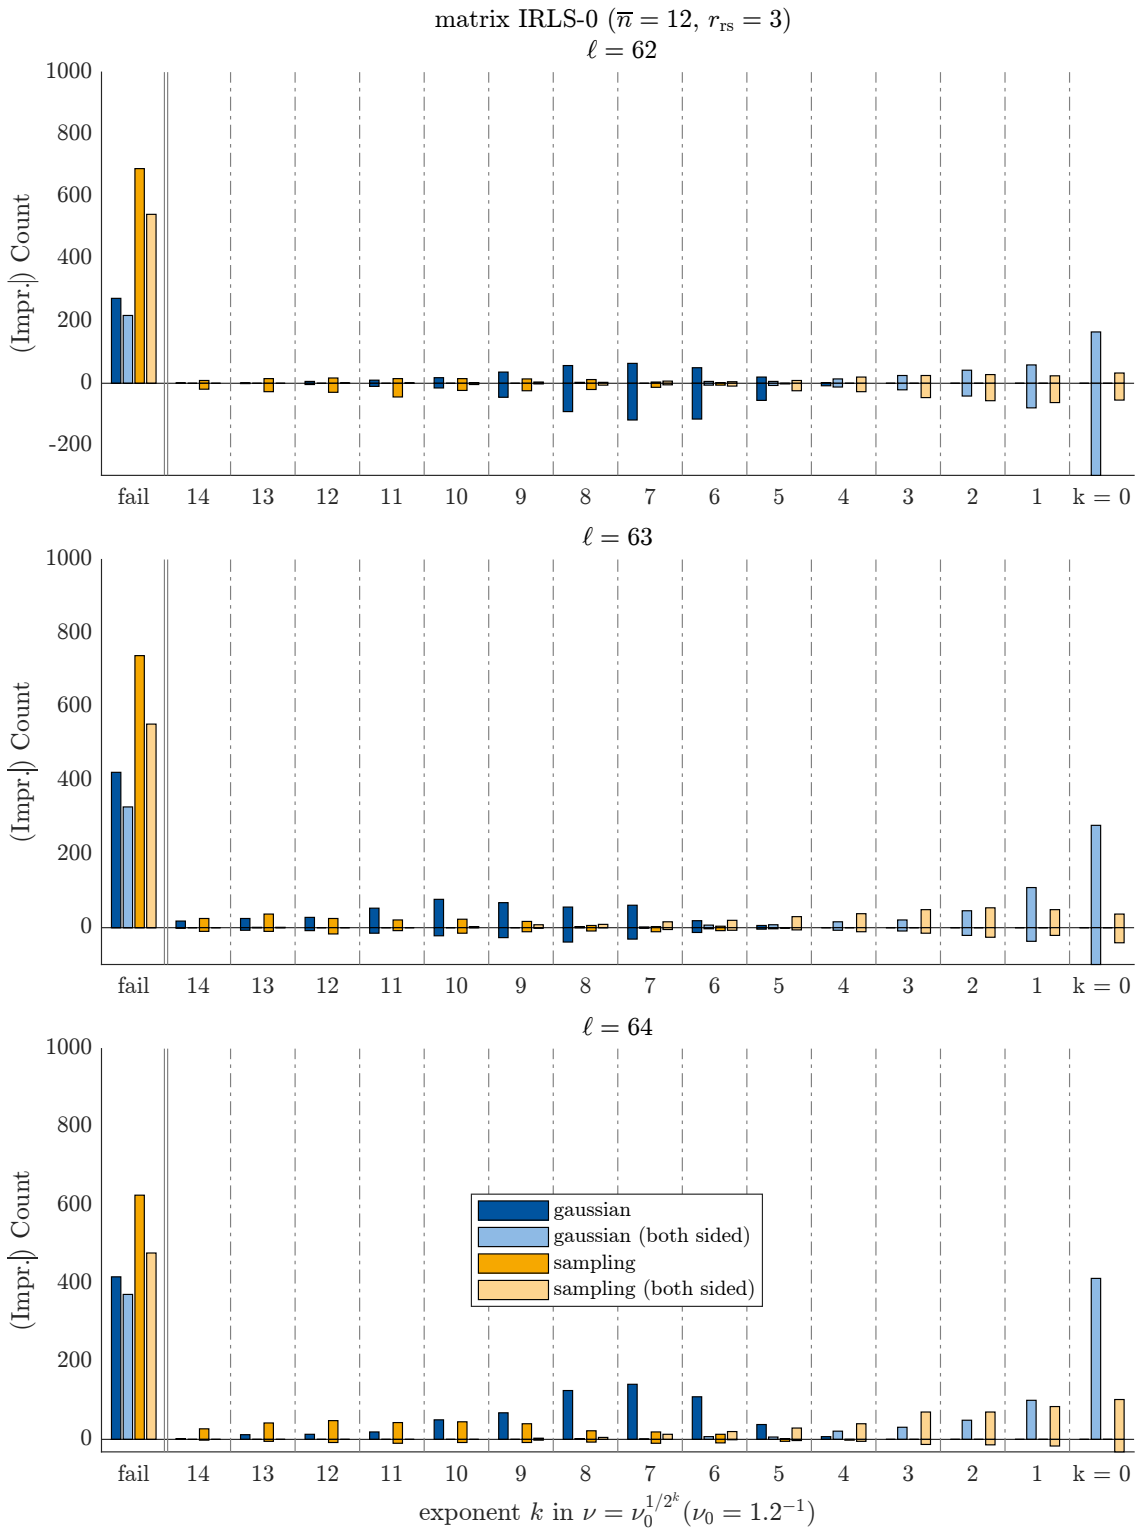

Figure 5: Results for Experiment 17 as described in Section 5.1.4.

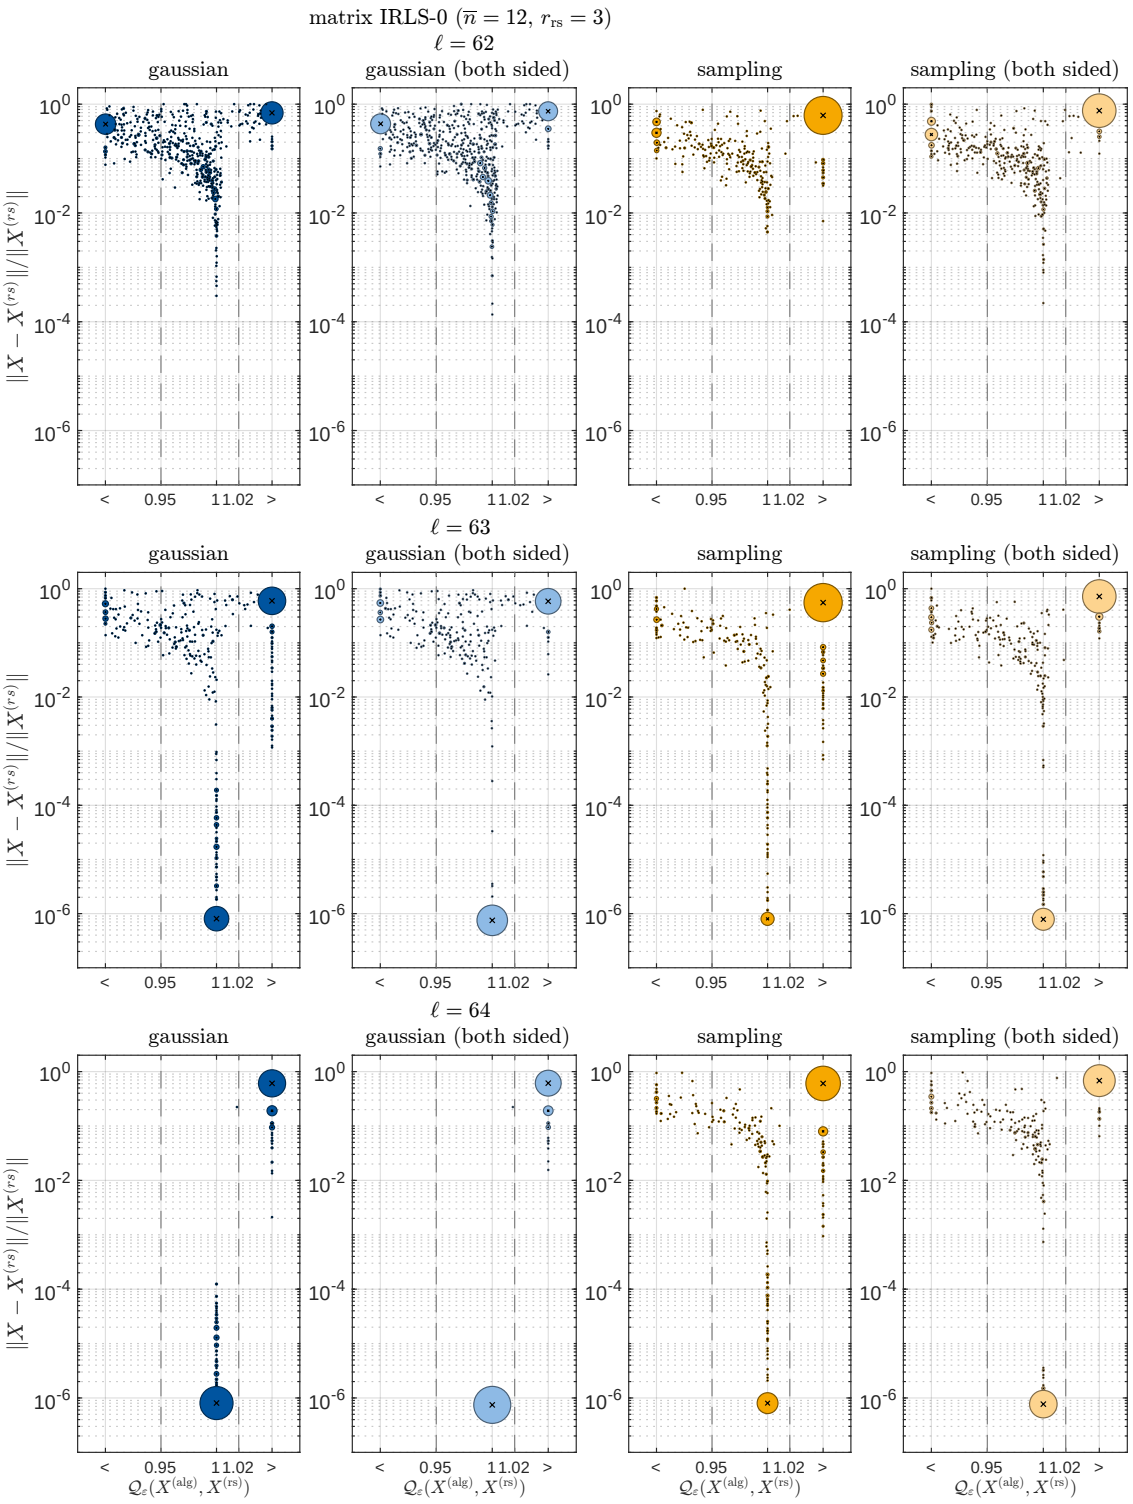

Figure 6: Results for Experiment 17 as described in Section 5.1.4. Note that the algorithm is set to automatically stop after the relative residual falls below 10^{-6} and certain other criteria hold true.

Appendix A. Proofs of Prior Lemmas in Deductive Order

Lemma 19 (Required for Lemma 4)**.** *For all $a, b \geq 0$ and $p \in (0, 1]$, it holds*

$$a^p + b^p \leq 2\left(ab + \frac{p(a-b)^2}{4}\right)^{p/2}.$$

Proof Without loss of generality, let $b \geq a$. The case $a = 0$ follows directly due to $2^p \leq 2p^{p/2}$. Otherwise, dividing the inequality by a^p and substituting $x \equiv b/a \geq 1$ yields the equivalent $f(x)^{p/2} \leq h(x)^{p/2}$, for

$$f(x)^{p/2} := \frac{1+x^p}{2},\ h(x) := x + \frac{p(1-x)^2}{4} = f(1) + f'(1)(x-1) + f''(1)\frac{(x-1)^2}{2}.$$

Direct calculation shows that the third derivative fulfills $f^{(3)}(x) \leq 0$ for $x \geq 1$. A Taylor series then yields $f(x) = h(x) + f^{(3)}(\xi(x))(x-1)^3 \leq h(x)$ since $\xi(x) \geq 1$. ■

Proof of Lemma 4 The matrix $X_\delta(a)$ is at most rank 2. The product of its two first eigenvalues $\{\lambda_1, \lambda_2\}$ is given through the 2×2-minors (cf. (4))

$$\lambda_1\lambda_2 = p_1(X_\delta(a)) = \det\left(\begin{pmatrix} x_\delta & \sqrt{a-1}e_n \end{pmatrix}^T \begin{pmatrix} x_\delta & \sqrt{a-1}e_n \end{pmatrix}\right) = (a-1)\Delta_\delta, \quad x_\delta = \begin{pmatrix}\delta \\ 1\end{pmatrix}.$$

It further holds $\lambda_1 + \lambda_2 = \operatorname{trace}(X_\delta(a)) = \Delta_\delta + a$, and thus $(\Delta_\delta + a - \lambda_2)\lambda_2 = \Delta_\delta(a-1)$. This equation is solved by $a = \overline{a}_\delta(\lambda_2)$ within the domain of this bijective function,

$$\overline{a}_\delta : (-\sqrt{\Delta_\delta + \Delta_\delta^2}, \Delta_\delta) \to (-\Delta_\delta, \infty), \quad \overline{a}_\delta(t) := \frac{\Delta_\delta + \Delta_\delta t - t^2}{\Delta_\delta - t}. \tag{15}$$

Inserted in the trace equation, this yields $\lambda_1 = \overline{\lambda}_{1,\delta}(\lambda_2)$ for $\overline{\lambda}_{1,\delta}(t) := \frac{\Delta_\delta(\Delta_\delta - t + 1)}{\Delta_\delta - t} > 0$, $t < \Delta_\delta$. The singular values of $X_\delta(\overline{a}_\delta(t))$ in turn are given as $\sigma_1(t) = \overline{\lambda}_{1,\delta}(t) > \sigma_2(t) = |t|$, since the matrix is symmetric and $|\overline{\lambda}_{1,\delta}(t)| = |t|$, $t < \Delta_\delta$, is solved only by $t = -\sqrt{\Delta_\delta + \Delta_\delta^2}$. For the derivative of the Schatten-p function, we thus have

$$\frac{\partial}{\partial t}(S_{\gamma,p}(X_\delta(\overline{a}_\delta(t))))|_{t=0} = \frac{\partial}{\partial t}(\overline{\lambda}_{1,\delta}(t)^2 + \gamma)^{p/2}|_{t=0}, \quad \frac{\partial}{\partial t}(\overline{a}_\delta(t))|_{t=0} = 1 + \frac{1}{\Delta_\delta}.$$

Plain evaluation and consideration of the chain rule then yields the first sought result. For the second part, setting $t^* := \Delta_\delta - \sqrt{\Delta_\delta} \leq 0$ such that $\Delta_\delta = \overline{a}_\delta(t^*)$, we have $-\overline{\lambda}_{1,\delta}(t^*)t^* = -\det(X_\delta(\Delta_\delta)) = (1-\Delta_\delta)\Delta_\delta$ and $\overline{\lambda}_{1,\delta}(t^*) + t^* = \operatorname{trace}(X_\delta(\Delta_\delta)) = 2\Delta_\delta$. With Lemma 19,

$$S_{0,p}(X_\delta(\overline{a}_\delta(t^*))) = \overline{\lambda}_{1,\delta}(t^*)^p + (-t^*)^p \leq 2((1-\Delta_\delta)\Delta_\delta + \frac{p(2\Delta_\delta)^2}{4})^{p/2} \leq 2\Delta^{p/2}. \tag{16}$$

Further, $S_{0,p}(X_\delta(1)) = \overline{\lambda}_{1,\delta}(0)^p = (1+\Delta_\delta)^p$. The inequality $S_{0,p}(X_\delta(\Delta_\delta)) < S_{0,p}(X_\delta(1))$ is thus implied by $2\Delta^{p/2} < (1+\Delta_\delta)^p$. Further, this bound implies that the minimizers of $S_{\gamma,p}|_{\mathcal{L}^{-1}(y_\delta)}$ can not converge to $X_\delta(1)$ for $\gamma \searrow 0$ due to continuity in $\gamma \geq 0$. Lastly, $S_{0,1}(X_\delta(a))^2 = 4(1-a)\Delta_\delta + (\Delta_\delta + a)^2$, $a \in [0,1]$, is minimized by $a = \Delta_\delta = \overline{a}_\delta(t^*)$. ■

Proof of Lemma 10 Assume to the contrary that $\{X^{(i)}\}_{i\in\mathbb{N}_0}$ remains bounded. As there is no singular matrix within $\mathcal{L}^{-1}(y)$, this implies that also $\{|f_{\gamma^{(i)}}(X^{(i)})|\}_{i\in\mathbb{N}_0}$ remains bounded. By Assumption 7, part (iii), the sequence $\{X^{(i)}\}_{i\in\mathbb{N}_0}$ must have an accumulation point in S_0^*, which as described in Example 6 only consists of the saddle point $X(1,0)$. As by assumption $f_{\gamma^{(0)}}(X^{(0)}) < \log(2) = f_0(X(1,0))$, part (i) of Assumption 7 ensures that the only stationary point can not be an accumulation point. By contradiction, the sequence $\{X^{(i)}\}_{i\in\mathbb{N}_0}$ must thus diverge, necessarily along one of the valleys of f_0. ■

Proof of Lemma 11 Throughout the following, we substitute $\tau \equiv \sqrt{\gamma} > 0$ as well as $\beta \equiv \alpha^{-1/2} > 0$. Firstly, with (10), it follows that $a_\delta^*(\tau) = \frac{\Delta_\delta^2}{\Delta_\delta^2+\tau^2} \in (0,1)$ is the sole stationary point. Since $b_\delta : [\overline{a}_\delta^{-1}(0), 0] \to [0,1]$, $b_\delta(t) := \overline{a}_\delta(t)$ (cf. Lemma 4) is bijective with $\sigma_2(X_\delta(b_\delta(-t))) = t$, the fixed point condition $(\Delta_\delta, \beta^{-2}, \tau^2, a^*) \in \Lambda$ is equivalent to $b_\delta(-\tau\beta) = a_\delta^*(\tau)$. By simple rearrangement, this can be written as the polynomial condition $H_{\delta,\tau}(\beta) := (\Delta_\delta^2\tau + \tau^3)\beta^2 + (\Delta_\delta^3 + \tau^2\Delta_\delta + \Delta_\delta^2)\beta - \Delta_\delta\tau = 0$, for which there exists one positive root. For the first part, by further substituting $\beta \equiv z\Delta_\delta^{-1/4}$ and $\tau \equiv \zeta\Delta_\delta^{3/4}$, we define

$$Y_{\delta,\zeta}(z) := \Delta_\delta^{-7/4} H_{\delta,\zeta\Delta_\delta^{3/4}}(z\Delta_\delta^{-1/4}) = (\zeta^3 + \sqrt{\Delta_\delta}\zeta)z^2 + (\sqrt{\Delta_\delta}\zeta^2 + \Delta_\delta + 1)z - \zeta,$$

as well as its positive root $z_0(\Delta_\delta, \zeta)$. Considering the coefficients of $Y_{\delta,\zeta}(z)$, the root z_0 is continuous and monotonically decreasing in $\Delta_\delta \geq 0$, whereby $\sup_{\Delta_\delta>0} z_0(\Delta_\delta, \zeta) = z_0(0,\zeta)$. The maximum z^* of $z_0(0,\zeta)$ over $\zeta \geq 0$ is then determined as $z^* \equiv (\kappa^*)^{-1/2}$ at $\zeta^* = 3^{1/4}\cdot 2^{-1/2}$. In other words, z^* is the minimal constant such that $Y_{\delta,\zeta}(z) \neq 0$ for all $z \geq z^*$ and $\Delta_\delta, \zeta > 0$. Resubstituting, this thus means that κ^* is the well defined maximal constant such that there does not exist $(\Delta_\delta, \alpha, \gamma, a^*) \in \Lambda$ with $\alpha \equiv \beta^{-2} \equiv z^{-2}\sqrt{\Delta_\delta} \leq (z^*)^{-2}\sqrt{\Delta_\delta} \equiv \kappa^*\sqrt{\Delta_\delta}$. For the second part, via substitution of $\beta \equiv z\Delta_\delta^0 = z$ and $\tau \equiv \omega\Delta_\delta$, we define

$$Y_{\delta,\omega}(z) := \Delta_\delta^{-2} H_{\delta,\omega\Delta_\delta}(z) = (\Delta_\delta\omega^3 + \omega\Delta_\delta)z^2 + (\Delta_\delta\omega^2 + \Delta_\delta + 1)z - \omega,$$

as well as its positive root $z_0(\Delta_\delta, \omega)$. With analogous argumentation, it then follows by $z_0(0,\omega) = \omega$, that $z^* \equiv (\nu^*)^{-1/2} = \varepsilon^{-1/2}$ is the minimal constant such that $Y_{\delta,\omega}(z) \neq 0$ for all $z \geq z^*$, $\omega \leq \varepsilon^{-1/2}$ and $\Delta_\delta > 0$. Resubstituting, this thus means that $\nu^* \equiv \varepsilon$ is the well defined maximal constant such that there does not exist $(\Delta_\delta, \alpha, \gamma, a^*) \in \Lambda$ with $\alpha \equiv \beta^{-2} \equiv z^{-2} \leq (z^*)^{-2} \equiv \nu^*$ and $\sqrt{\gamma} \equiv \tau \equiv \omega\Delta_\delta \leq \varepsilon^{-1/2}\Delta_\delta$. Since the latter condition is equivalent to $|1 - a_\delta^*(\tau)| \leq \frac{1}{1+\varepsilon}$, this finishes the proof. ■

Appendix B. Apparent Stagnation through Singular Value Based Control

Similar to Section 4.3, the following Example 20 empirically illustrates a drastic instance of the more general problem that steep valleys form along a lower-dimensional curve around the sought matrix for $\gamma \searrow 0$. Therein, solely the chosen strategy to adapt γ can determine the same random, non-stationary point to which the iterates appear to converge, near independently of starting values. Yet the problem only lies within the discretization of $\gamma \searrow 0$ along with the chosen updates of IRLS-0.

Example 20. Let $\mathcal{L} : \mathbb{R}^{2\times 3} \to \mathbb{R}^4$ be a linear operator and $y \in \mathbb{R}^4$ such that

$$\mathcal{L}^{-1}(y) = \{X(a,b) \mid a,b \in \mathbb{R}\}, \quad X(a,b) := \begin{pmatrix} a & 1 & 3-b \\ 1 & b & 3-a \end{pmatrix}.$$

The functions $g_\gamma : \mathbb{R}^2 \to \mathbb{R}$, $g_\gamma(a,b) := f_\gamma(X(a,b))$, are visualized in Figure 7, together with the stationary points of $f_\gamma|_{\mathcal{L}^{-1}(y)}$ as derived in Lemma 21. Different Taylor series at the only stationary point $X(1,1)$ for $\gamma = 0$ reveal that

$$\det(X(s,1/s)X(s,1/s)^T) = 2(s-1)^6 + \mathcal{O}(s-1)^7,$$
$$\det(X(s,s)X(s,s)^T) = 12(s-1)^2 + \mathcal{O}(s-1)^3,$$

whereas $\|X(a,b)\|_F^2 = 12 + \mathcal{O}(a-1) + \mathcal{O}(b-1)$. Consequently, the second singular value $\sigma_2(a,b)$ of $X(a,b)$ along $(s,1/s)$ at $s=1$ grows much slower than in the orthogonal direction (s,s) at that point. Vice versa, $\sigma_2(a,b)$ can tend to zero rapidly for $(a,b) \to (1,1)$. Further, the functions g_γ yield a curved and, for $\gamma \to 0$, thus increasingly steep valley. In the face of Corollary 12, the optimization within these valleys can become so slow relatively to the decline of γ that it becomes practically indistinguishable from convergence:

1) Using the adaption $\gamma^{(i)} := \min(\gamma^{(i-1)}, \alpha\sigma_2(X^{(i)})^2)$, $\alpha = 1$, then for most[10] starting values $X^{(0)} = X(a_0,b_0)$, the parameter is decreased so quickly that the iterates *appear* to stagnate far away from the desired matrix $X(1,1)$. Namely, still the ten millionth iterate $X^{(10^7)}$ is close to the same $\widetilde{X}_1 := X(1.01987, 0.98051)$ or $\widetilde{X}_2 := X(0.98051, 1.01987)$, with $\gamma^{(10^7)} \approx \widetilde{\gamma} := 10^{-11}$, each with a deviation of roughly $\|X^{(10^7)} - \widetilde{X}\|_F \approx 10^{-6}$ and a change of $a^{(i)}$ and $b^{(i)}$ in the magnitude of 10^{-9} per iteration. However, the only stationary point for this $\widetilde{\gamma}$ is at $X(t,t) \in \mathcal{S}_{\widetilde{\gamma}}^*$ for $t \approx 1 + 10^{-12}$ (see Lemma 21). Thus, while the iterates appear to converge to some distinct matrix, they neither converge to a stationary point nor an artificial stationary point as in Lemma 11, but at the same iteration pass one of the same two non-stationary matrices solely determined by the strategy to adapt γ. Moreover, these two points are quite robust towards changes in α. For both-sided weights (cf. Section 5.1.2 and Kümmerle and Verdun (2021)), the iterates come much closer to the sought solution, but the same type of effect is still observable.

2) Using the adaption $\gamma^{(i)} := \nu\gamma^{(i-1)}$ for $\gamma^{(0)} = 1$ (cf. Proposition 14) with $\nu = 0.999$ given the same random starting values, then less than 30000 iterations suffice to reach $X(1,1)$ with an accuracy close to machine precision. For moderately larger values of ν, the iterates do again not come close to the stationary points.

Lemma 21. *The only stationary point of f_γ, $\gamma > 0$, given through the setting in Example 20, is $X(a,b)$ for*

$$a = b = \frac{z^2 + 8z - 20 - 8\gamma}{6z}, \quad z^3 = 12\gamma - 64 + 4\sqrt{32\gamma^3 + 249\gamma^2 + 504\gamma + 756}, \tag{17}$$

for $z \in \mathbb{R}$. The only stationary point of $\exp \circ f_0$, that is, the only singular matrix within $\mathcal{L}^{-1}(y)$, is given through $a = b = 1$.

10. This is based on 50 different, random starting values, where a_0 and b_0 are each drawn from an independent, normal distribution, with 3 result not fulfilling the bound.

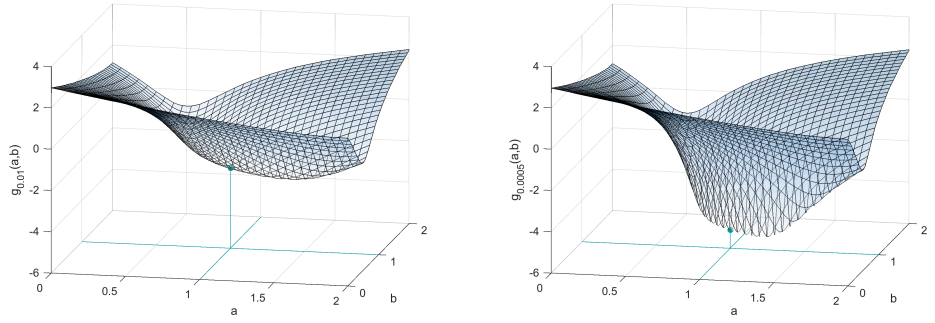

Figure 7: The functions $g_\gamma(a,b) := f_\gamma(X(a,b))$ for $\gamma \in \{0.01, 0.0005\}$ as defined through the setting in Example 20 together with their respective minimum.

Proof Let $h_\gamma(t,s) := \exp(f_\gamma(X(t+s,t-s)))$. To find all stationary points, one may search for $(t,s) \in \mathbb{R}^2$ that fulfill

$$\frac{\partial}{\partial s} h_\gamma(t,s) = \big(12s^2 + p_-(t,\gamma)\big)s \stackrel{!}{=} 0 \text{ and } \frac{\partial}{\partial t} h_\gamma(t,s) = (20t-32)s^2 + p_+(t,\gamma) \stackrel{!}{=} 0,$$

where $p_-(t,\gamma) = 20t^2 - 64t + 8\gamma + 44$ and $p_+(t,\gamma) = 12t^3 - 48t^2 + 8t\gamma + 84t - 12\gamma - 48$. Assume that both equations hold true. Then for fixed $s \neq 0$, we obtain $12s^2 + p_-(t,\gamma) = 0$. The real roots of this polynomial are $t_1 = \frac{8-r}{5}$ and $t_2 = \frac{8+r}{5}$ for $r = \sqrt{9 - 15s^2 - 10\gamma} \in [0,3)$ and according $s \neq 0$ and $\gamma > 0$. Inserting these roots into the other derivative $\frac{\partial}{\partial t} h_\gamma$ yields the each contradictory $125 \frac{\partial}{\partial t} h_\gamma|_{t=t_1} = 12r^3 + 48r^2 + (500s^2 + 200\gamma + 564)r + 100\gamma + 1584 > 0$ and $375 \frac{\partial}{\partial t} h_\gamma|_{t=t_2} = 4(100r+75)\gamma + 4(r-3)^2(16r+132) > 0$. Thus, $s = 0$ must hold true. The points at which $\frac{\partial}{\partial t} p_+(t,\gamma) = 0$ holds true are $\frac{4}{3} \pm \frac{\sqrt{-5-2\gamma}}{3} \in \mathbb{C} \setminus \mathbb{R}$. Hence, $p_+(t,\gamma)$ has exactly one real root, which, when inserted, are confirmed to be $t = a = b$ as in (17) for $\gamma > 0$, as well as $t = 1$ for $\gamma = 0$. ■

Appendix C. Affine Cardinality Minimization

Experiment 22. For $n = 160$, $|\mathrm{supp}(\mathrm{x}^{(\mathrm{ss})})| = 40$ and $\ell \in \{65, 75, 85, 95\}$, we consider the cardinality minimization problem based on *Gaussian measurements* under varying choice of the weight strength $p \in \{0, 0.2, 0.4, 0.6, 0.8, 1\}$. Each constellation is repeated 1000 times. The results are covered in Table 4 and Figures 8 and 9.

The results suggest that $p = 0$ (or at least a very small value) is here the overall best choice as well. This is in apparent contradiction with the results of Daubechies et al. (2010), who have also considered a very similar problem setting. Their results seem to yield that only $p \geq 0.5$ is recommendable. This might however be due to the different and possibly less tolerant strategy chosen for the adaption[11] of γ that does not seem suitable for smaller values of p (cf. Example 15). With ℓ close to $|\mathrm{supp}(\mathrm{x}^{(\mathrm{rs})})| = 40$, it is not surprising that fewer reference solutions are recovered, though it should theoretically be possible as long as

11. In Daubechies et al. (2010), these parameters are denoted τ and ε, respectively.

instance	ACM: recovery:	$\mathcal{Q}_\varepsilon \in [0, 0.98]$ no	$\mathcal{Q}_\varepsilon \in (0.98, 1.005)$ no	+	yes	$\mathcal{Q}_\varepsilon \in [1.005, \infty)$ no	$\mathcal{Q}_\varepsilon = \infty$
$\ell = 50$	$p = 0.0$	0	0	+	0	0	100
	$p = 0.2$	0	0	+	0	0	100
	$p = 0.4$	0	0	+	0	0	100
	$p = 0.6$	0	0	+	0	0	100
	$p = 0.8$	0	0	+	0	0	100
	$p = 1.0$	0	0	+	0	0	100
$\ell = 60$	$p = 0.0$	0	0	+	0.7	0	99.3
	$p = 0.2$	0	0	+	0.6	0	99.4
	$p = 0.4$	0	0	+	0.2	0	99.8
	$p = 0.6$	0	0	+	0	0	100
	$p = 0.8$	0	0	+	0	0	100
	$p = 1.0$	0	0	+	0	0	100
$\ell = 70$	$p = 0.0$	0	0	+	39.6	0	60.4
	$p = 0.2$	0	0	+	39	0	61
	$p = 0.4$	0	0	+	34.7	0	65.3
	$p = 0.6$	0	0	+	24.9	0	75.1
	$p = 0.8$	0	0	+	8.6	0	91.4
	$p = 1.0$	0	0	+	0	0	100
$\ell = 80$	$p = 0.0$	0	0	+	92.1	0	7.9
	$p = 0.2$	0	0	+	93	0	7
	$p = 0.4$	0	0	+	90.3	0	9.7
	$p = 0.6$	0	0	+	86.1	0	13.9
	$p = 0.8$	0	0	+	66.7	0	33.3
	$p = 1.0$	0	0	+	2.3	0	97.7
$\ell = 90$	$p = 0.0$	0	0	+	99.8	0	0.2
	$p = 0.2$	0	0	+	99.8	0	0.2
	$p = 0.4$	0	0	+	99.8	0	0.2
	$p = 0.6$	0	0	+	99.5	0	0.5
	$p = 0.8$	0	0	+	97.9	0	2.1
	$p = 1.0$	0	0	+	35.6	0	64.4
$\ell = 100$	$p = 0.0$	0	0	+	100	0	0
	$p = 0.2$	0	0	+	100	0	0
	$p = 0.4$	0	0	+	100	0	0
	$p = 0.6$	0	0	+	100	0	0
	$p = 0.8$	0	0	+	100	0	0
	$p = 1.0$	0	0	+	88.2	0	11.8

Table 4: vector IRLS-p (gaussian, $n = 160$, $\mathrm{card}(x_{\mathrm{rs}}) = 40$) – table as specified in Section 5.1.4 for Experiment 22 (see sensitivity analysis, Figure 8, and button plot, Figure 9, for more details).

$\ell \geq 41$ (cf. Breiding et al. (2023)). In contrast to ARM in the matrix case, vector IRLS-p here also tends to fail to solve the ACM problem for minimal numbers of measurements, possibly due to the less fortunate, only partial nestedness of its defining varieties. Note that for $\ell = 40$ in turn, finding a solution with cardinality less or equal 40 immediately becomes trivial.

Appendix D. Singular Value Based Adaption

Experiment 23. For $\overline{n} = 12$, $r_{\mathrm{rs}} = 3$ and $\ell \in \{63, 64, 72, 108\}$, we consider the ARM problem based on *samples* or *Gaussian measurements* for the weight strength $p = 0$ with *both-sided* weights. We compare constant decrease $\gamma^{(i)} = \nu\gamma^{(i-1)}$ with the singular based adaption $\gamma^{(i)} = \min(\gamma^{(i-1)}, \alpha\sigma_{K+1}(X^{(i)})^2)$, $\alpha := 1$, for $K \in \{r_{\mathrm{rs}}, r_{\mathrm{rs}} + 1\}$ (cf. Appendix B) as indicated. In the latter case, ν solely corresponds to a breaking criterion to ensure a reasonable comparison. Each constellation is repeated 1000 times for $k_{\max} = 8$. The results are covered in Table 5 and Figures 14 and 15.

Appendix E. Button Plot

With a button plot (with logarithmic scale in y), we refer to a two dimensional, clustered scatter plot. Therein, any circular markers with centers (x_i, y_i) and areas v_i, $i = 1, \ldots, k$,

instance	ARM: / recovery:	$\mathcal{Q}_\varepsilon \in [0, 0.98]$ no	$\mathcal{Q}_\varepsilon \in (0.98, 1.005)$ no		yes	$\mathcal{Q}_\varepsilon \in [1.005, \infty)$ no	$\mathcal{Q}_\varepsilon = \infty$
$\ell = 63$	gaussian	18.1	5.9	+	43.1	10	22.9
	gaussian, $\sigma_{r_{\text{rs}}+1}(X)^2$	12.4	5.7	+	36	17.9	28
	gaussian, $\sigma_{r_{\text{rs}}+2}(X)^2$	0	0	+	0	0	100
	sampling	13.3	7.3	+	24	3	52.4
	sampling, $\sigma_{r_{\text{rs}}+1}(X)^2$	12.7	6.3	+	25.1	6.1	49.8
	sampling, $\sigma_{r_{\text{rs}}+2}(X)^2$	0	0	+	0	0	100
$\ell = 64$	gaussian	0	0	+	62.9	0.1	37
	gaussian, $\sigma_{r_{\text{rs}}+1}(X)^2$	0.1	0	+	55.9	0.1	43.9
	gaussian, $\sigma_{r_{\text{rs}}+2}(X)^2$	0	0	+	0.4	0	99.6
	sampling	8.7	7.7	+	35.9	1.1	46.6
	sampling, $\sigma_{r_{\text{rs}}+1}(X)^2$	7.9	6	+	38.1	3.3	44.7
	sampling, $\sigma_{r_{\text{rs}}+2}(X)^2$	0	0	+	0.1	0	99.9
$\ell = 72$	gaussian	0	0	+	99	0	1
	gaussian, $\sigma_{r_{\text{rs}}+1}(X)^2$	0	0	+	98.8	0	1.2
	gaussian, $\sigma_{r_{\text{rs}}+2}(X)^2$	0	0	+	46.3	0	53.7
	sampling	2.4	1.8	+	82.2	0.1	13.5
	sampling, $\sigma_{r_{\text{rs}}+1}(X)^2$	2.3	1.5	+	81.7	0.4	14.1
	sampling, $\sigma_{r_{\text{rs}}+2}(X)^2$	1	0.4	+	19.1	0	79.5
$\ell = 108$	gaussian	0	0	+	100	0	0
	gaussian, $\sigma_{r_{\text{rs}}+1}(X)^2$	0	0	+	100	0	0
	gaussian, $\sigma_{r_{\text{rs}}+2}(X)^2$	0	0	+	100	0	0
	sampling	0	0	+	100	0	0
	sampling, $\sigma_{r_{\text{rs}}+1}(X)^2$	0	0	+	100	0	0
	sampling, $\sigma_{r_{\text{rs}}+2}(X)^2$	0.4	0	+	96.8	0.1	2.7

Table 5: matrix IRLS-0 (both sided, $\overline{n} = 12$, $r_{\text{rs}} = 3$) – table as specified in Section 5.1.4 for Experiment 23 (see sensitivity analysis, Figure 14, and button plot, Figure 15, for more details).

that would (visually) overlap, are recursively combined to each one larger circle $(\widehat{x}, \widehat{y})$ with area $\widehat{v}$ according to the appropriately weighted means

$$\widehat{x} = \sum_{i=1}^{k} \frac{v_i}{\widehat{v}} x_i, \quad \widehat{y} = \prod_{i=1}^{k} y_i^{v_i/\widehat{v}}, \quad \widehat{v} = \sum_{i=1}^{k} v_i.$$

The centers of all resulting circles are indicated as crosses. Thus, if only one circle remains, then the position of that cross is given by the arithmetic mean of all initial x-coordinates and the geometric mean of all initial y-coordinates. If no disks are combined, then their centers are the initial coordinates and their areas are all equal.

Appendix F. Further Sensitivity and ACM/ARM/Recovery Figures

Each of the following even and odd numbered pair of pages contains two related visualizations of the results of one of Experiments 16 to 18, 22 and 23 as summarized in Table 6. All such visualizations are summarized below by order of appearance of corresponding tables.

short description	experiment	γ-sensitivity	button plot	table
IRLS-p (varying p)	Experiment 16	Figure 10	Figure 11	Table 1
IRLS-0 (phase transition)	Experiment 17	Figure 5	Figure 6	Table 2
IRLS-0 (larger ℓ)	Experiment 18	Figure 12	Figure 13	Table 3
vector IRLS-p	Experiment 22	Figure 8	Figure 9	Table 4
σ_K-based adaption	Experiment 23	Figure 14	Figure 15	Table 5

Table 6: Overview over experiments, related figures and tables.

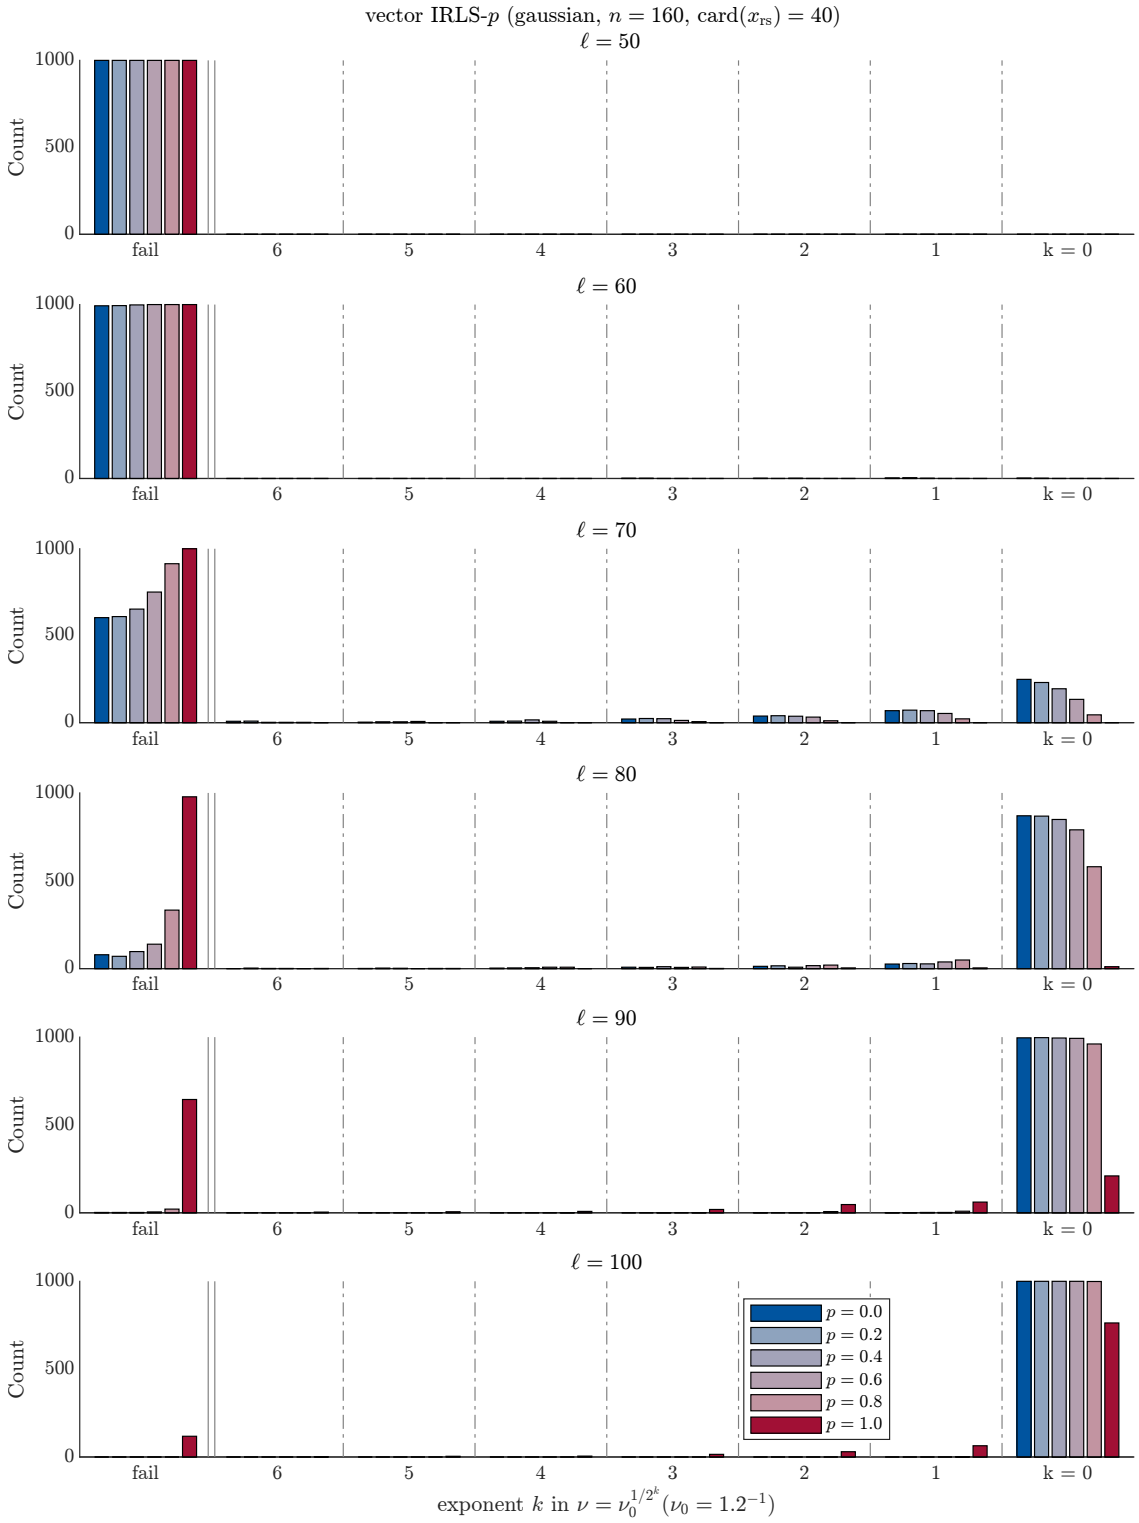

Figure 8: Results for Experiment 22 as described in Section 5.1.4.

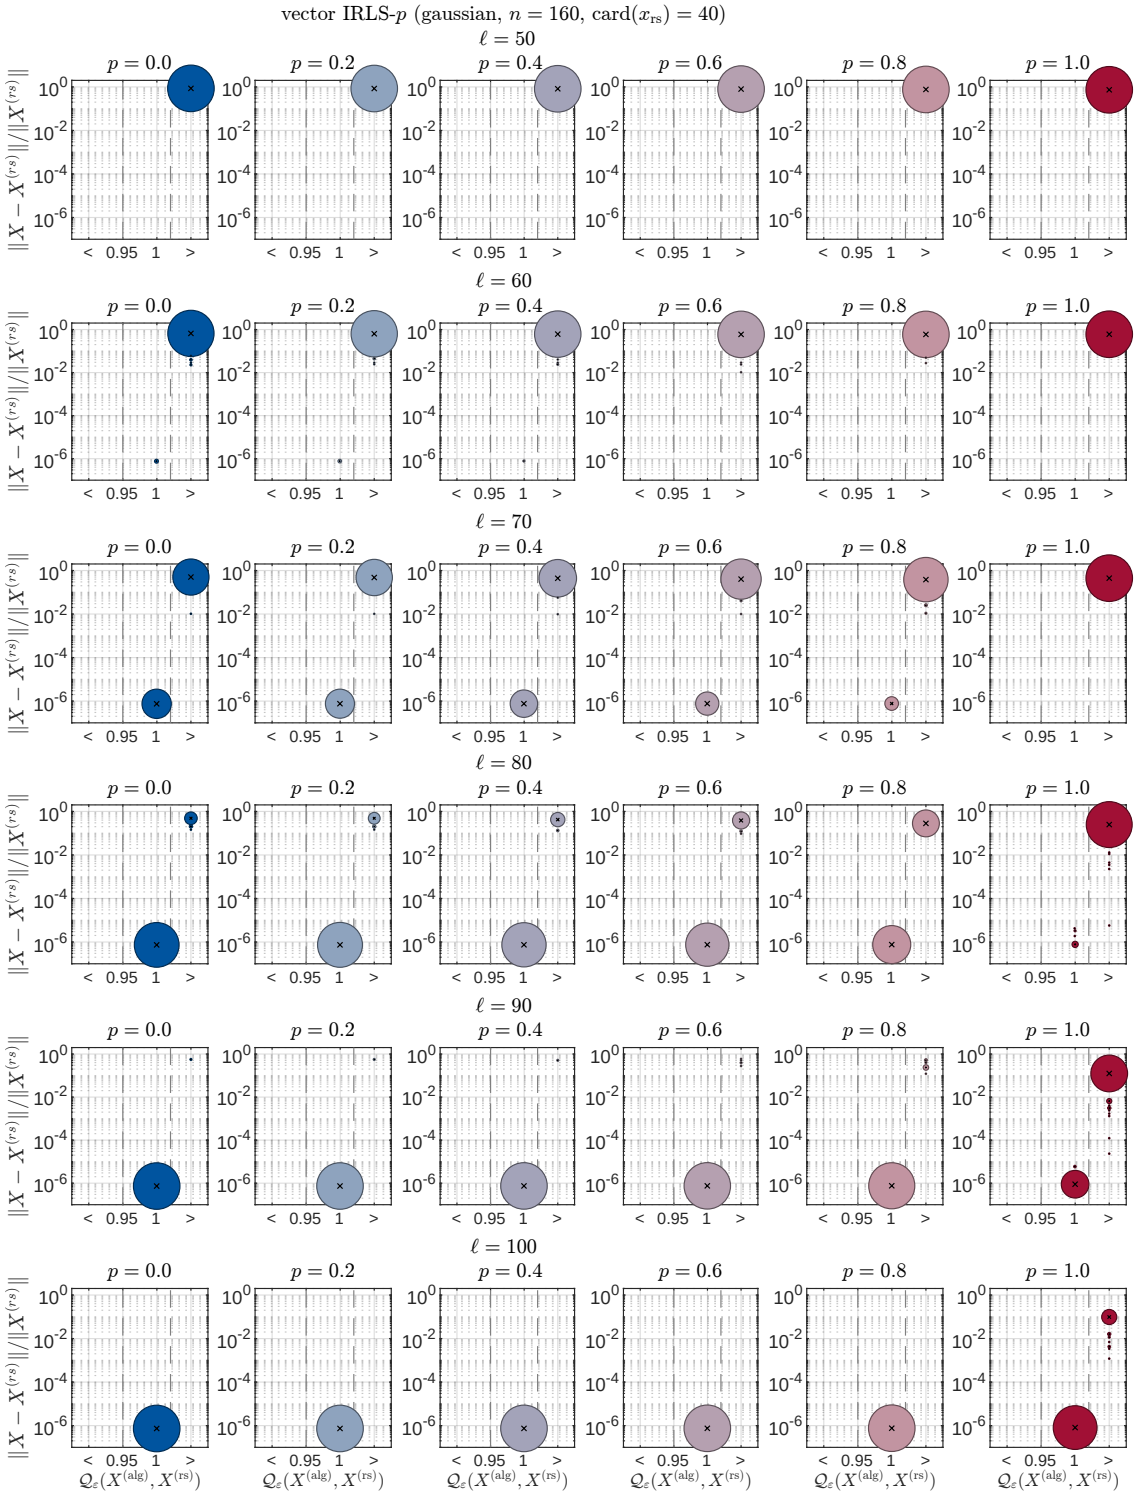

Figure 9: Results for Experiment 22 as described in Section 5.1.4.

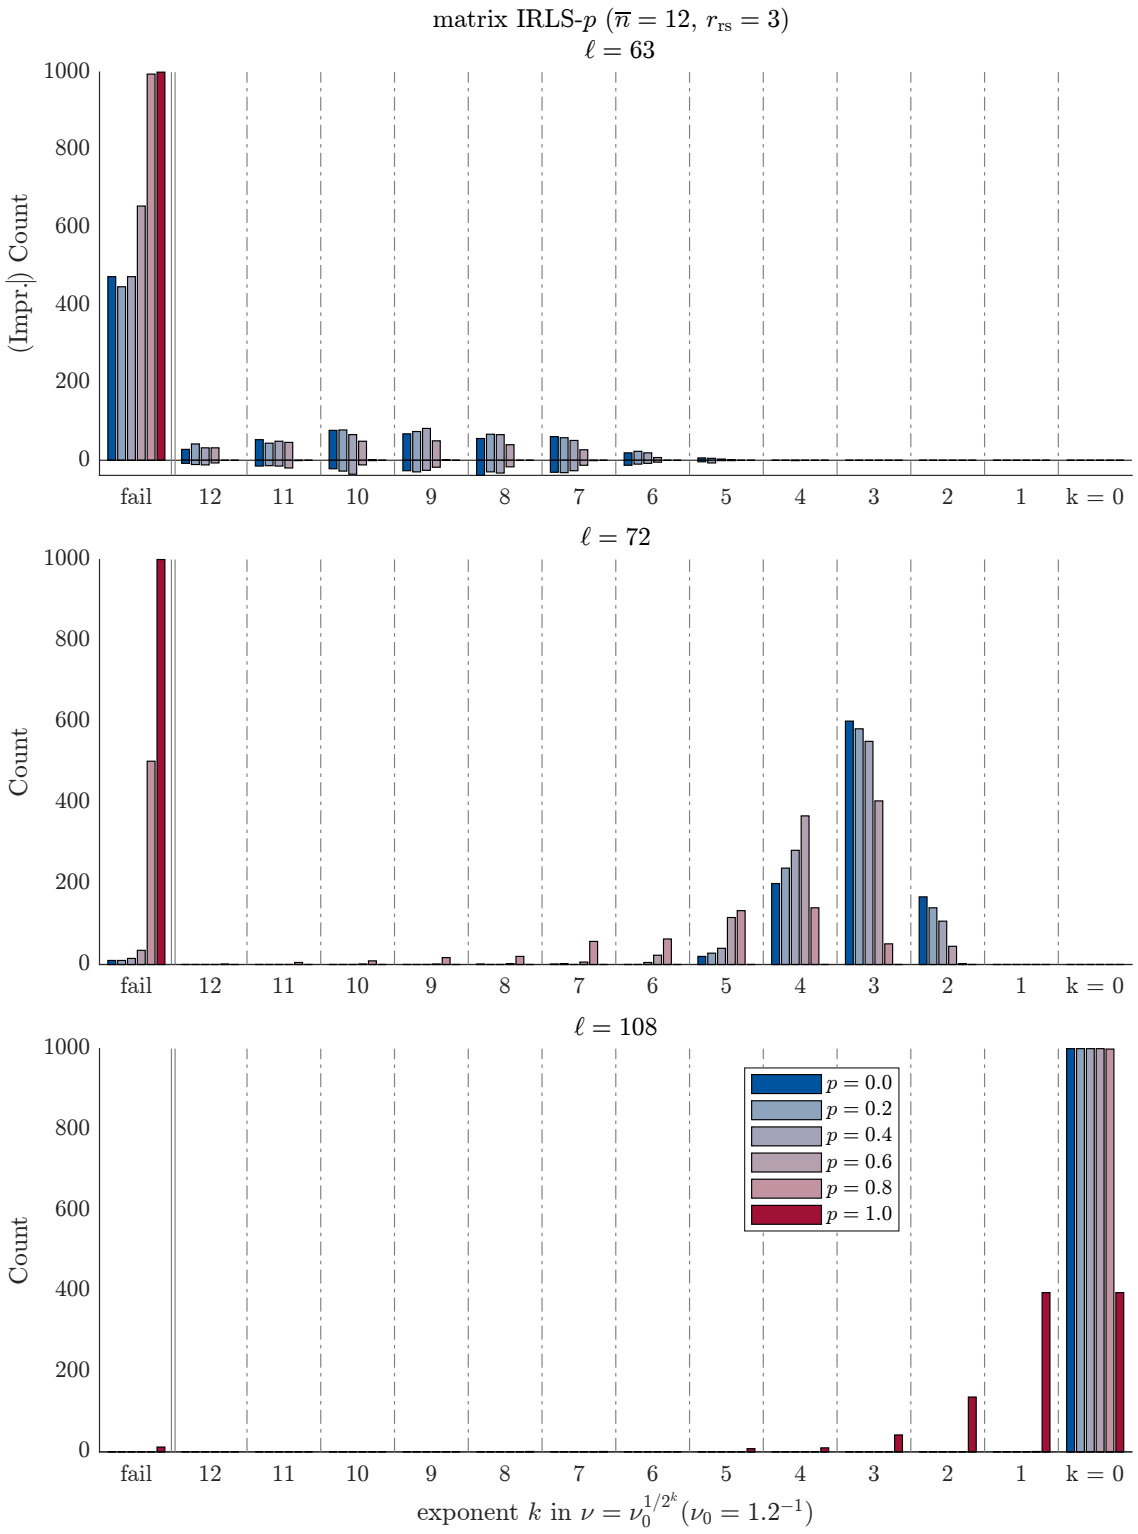

Figure 10: Results for Experiment 16 as described in Section 5.1.4.

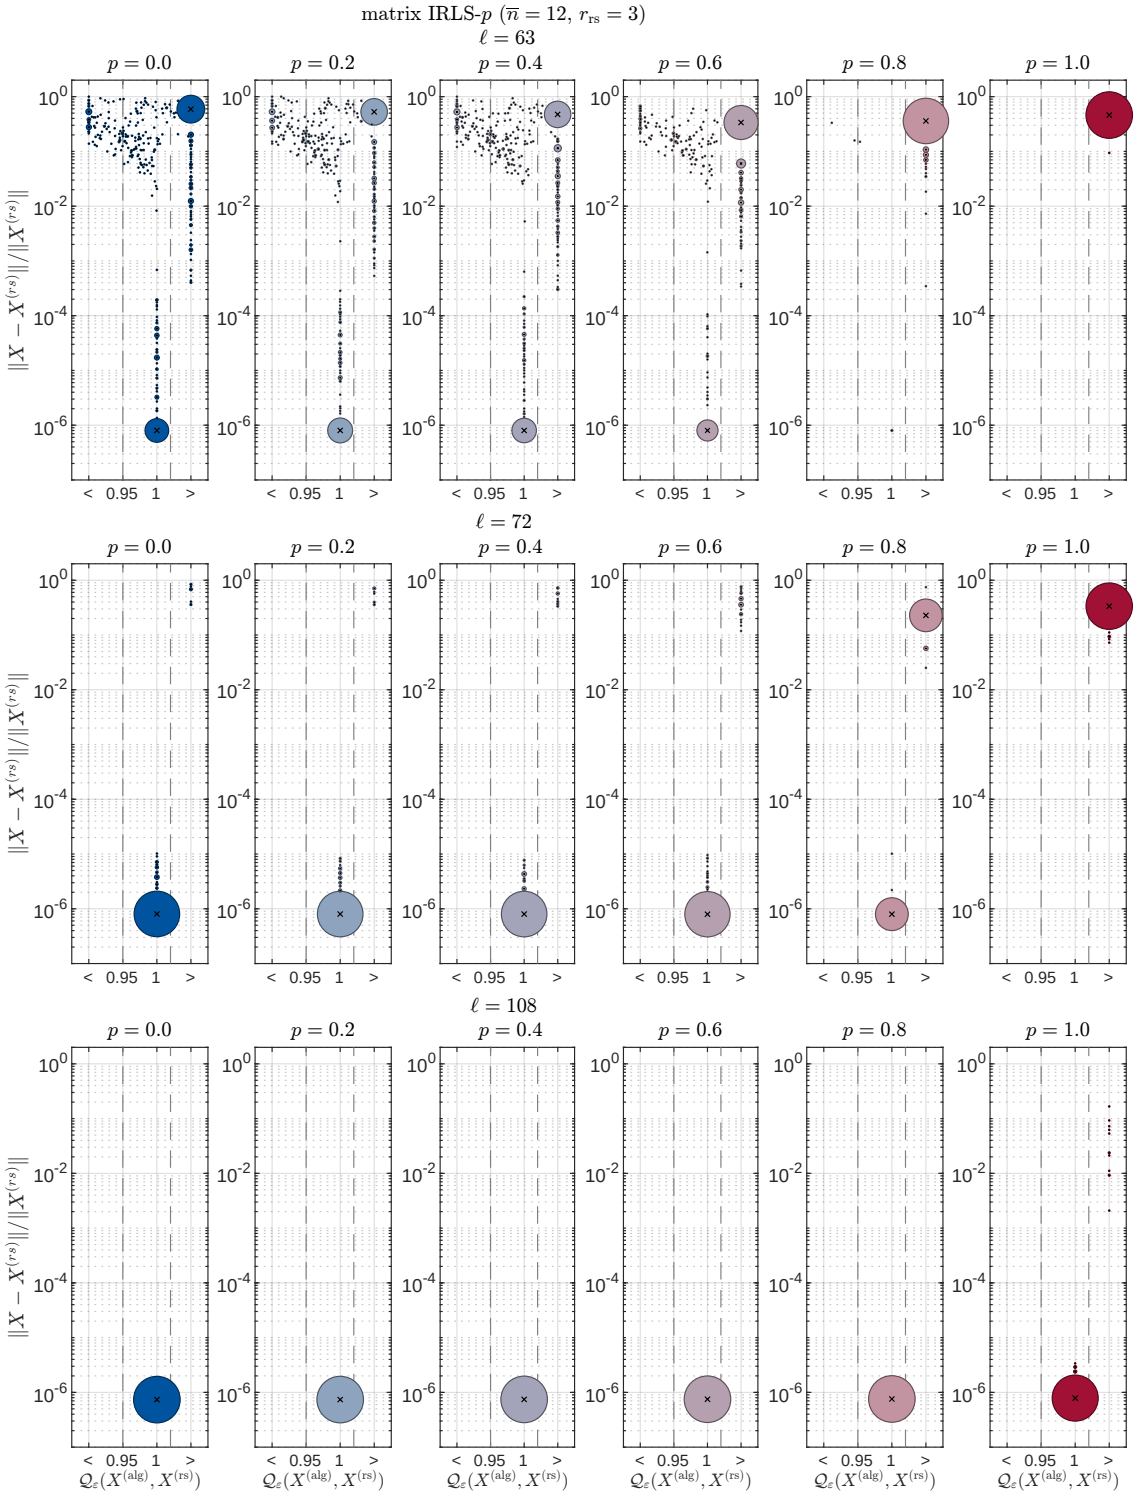

Figure 11: Results for Experiment 16 as described in Section 5.1.4.

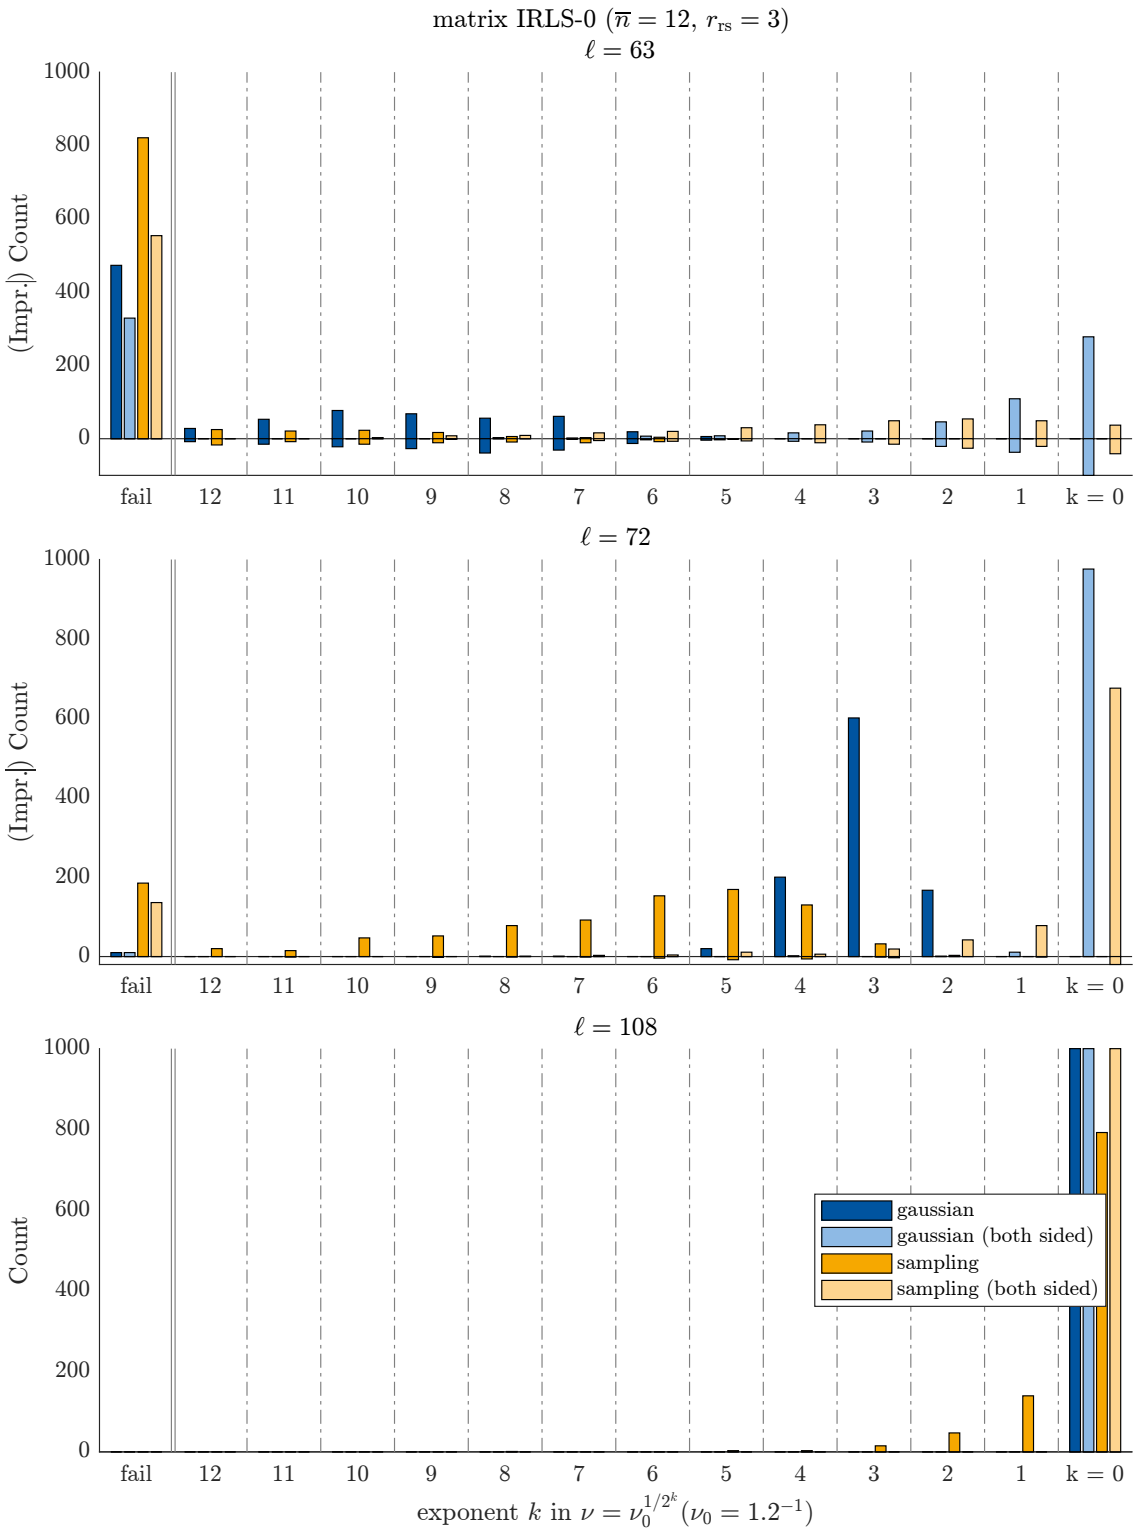

Figure 12: Results for Experiment 18 as described in Section 5.1.4.

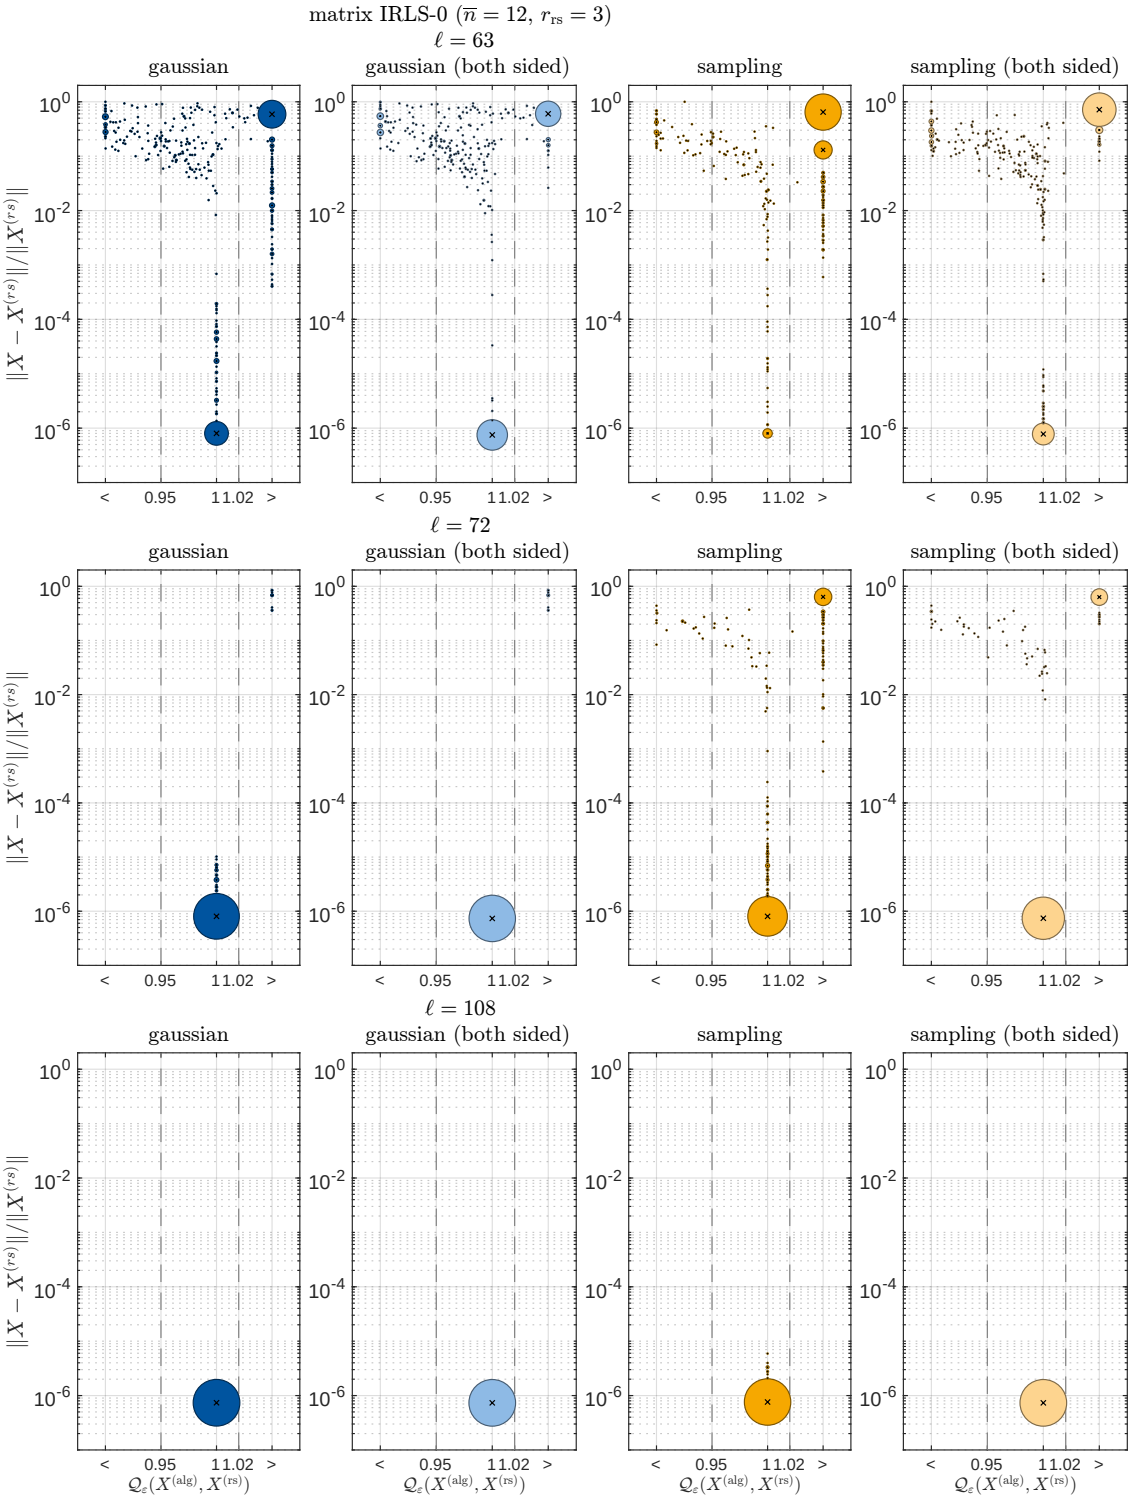

Figure 13: Results for Experiment 18 as described in Section 5.1.4.

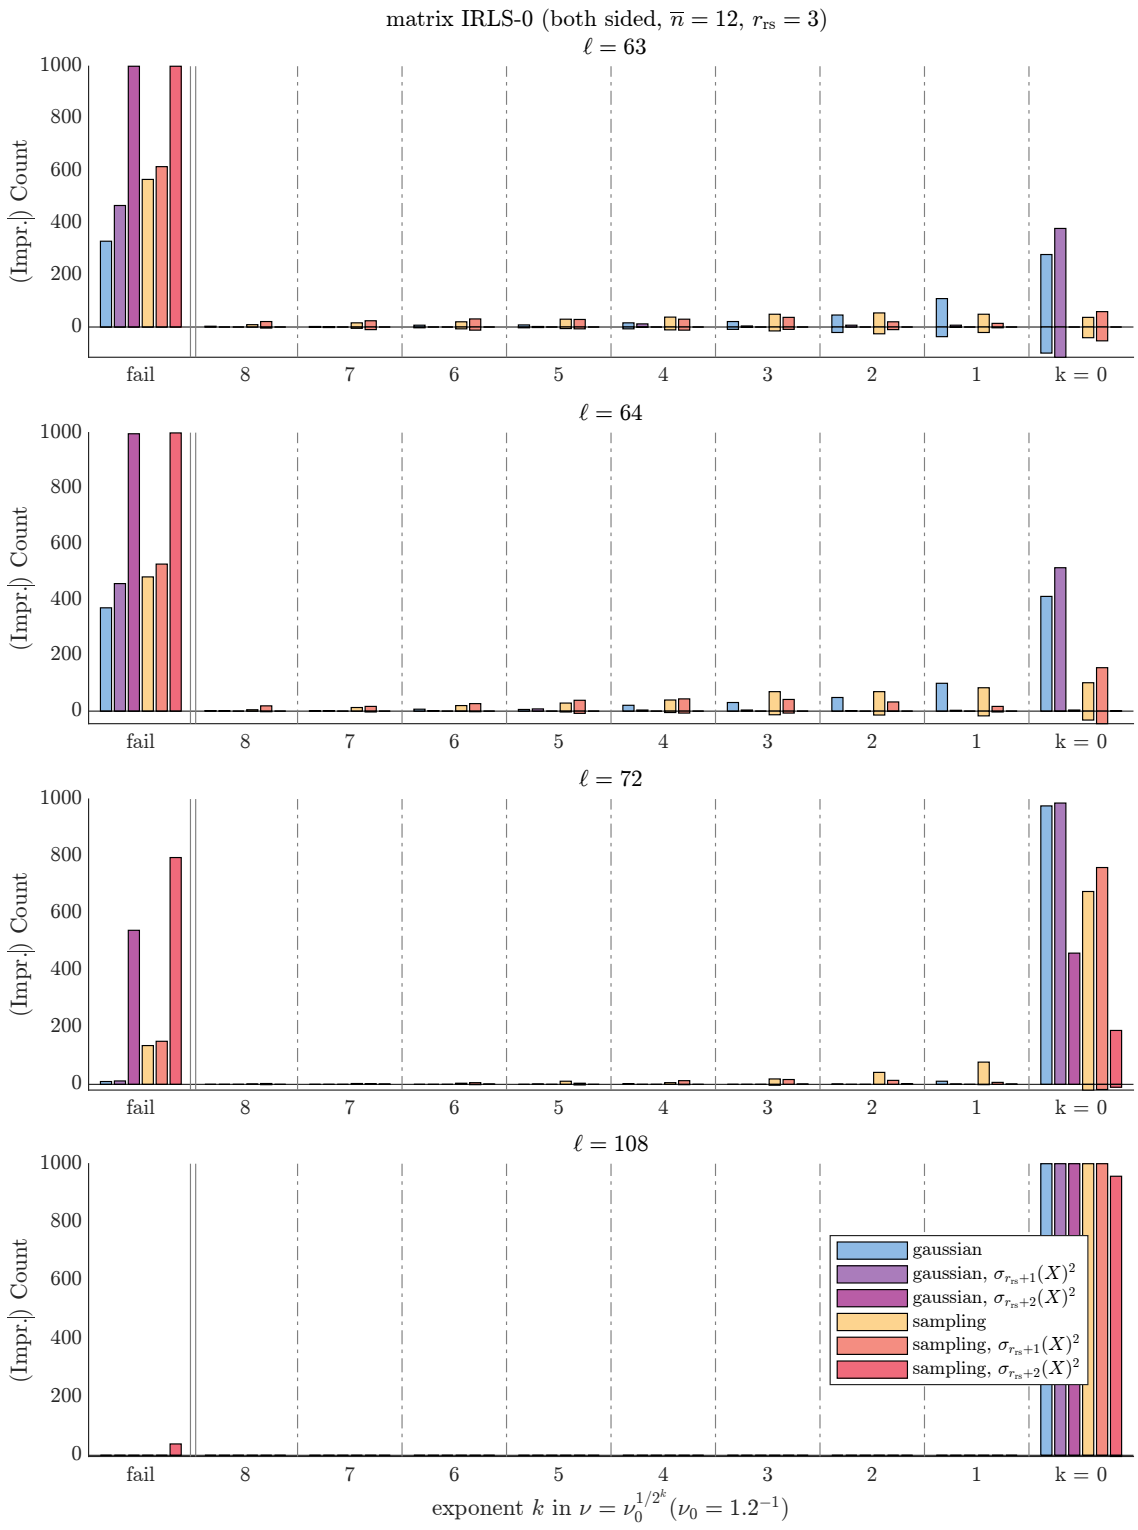

Figure 14: Results for Experiment 23 as described in Section 5.1.4.

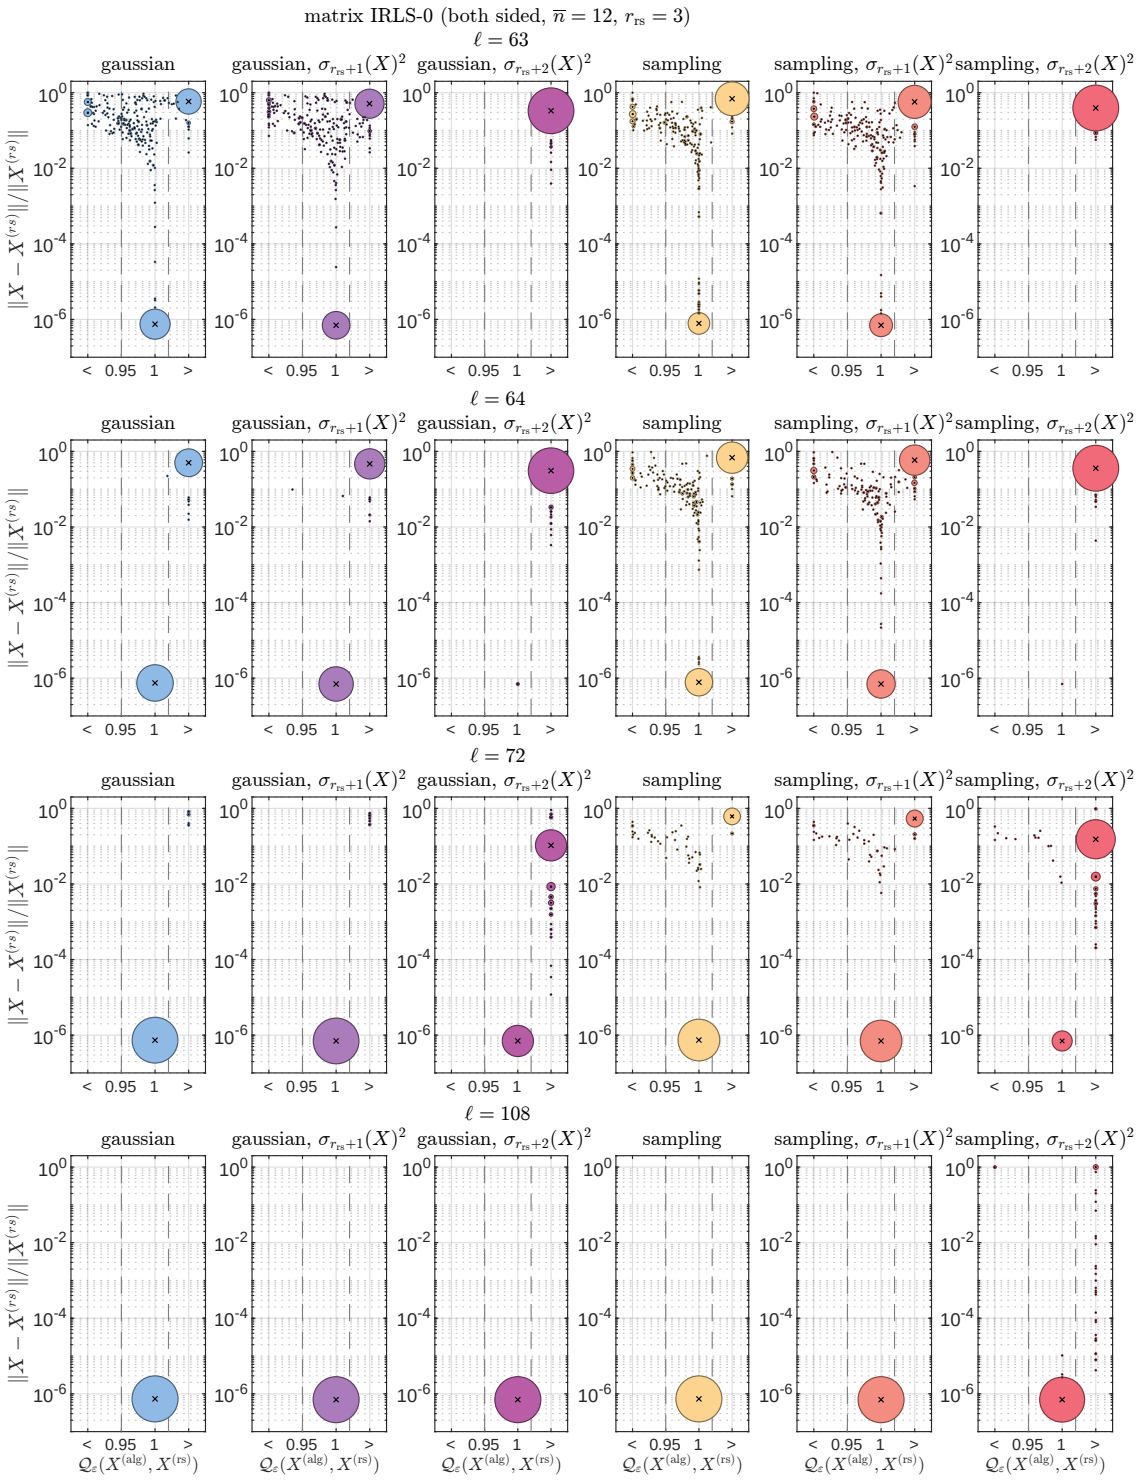

Figure 15: Results for Experiment 23 as described in Section 5.1.4.

References

Laura Balzano, Robert Nowak, and Benjamin Recht. Online identification and tracking of subspaces from highly incomplete information. In *2010 48th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 704–711. IEEE, sep 2010. doi:10.1109/allerton.2010.5706976.

Paul Breiding and Nick Vannieuwenhoven. Sensitivity of low-rank matrix recovery. *Numerische Mathematik*, 152(4):725–759, October 2022. ISSN 0945-3245. doi:10.1007/s00211-022-01327-7.

Paul Breiding, Fulvio Gesmundo, Mateusz Michałek, and Nick Vannieuwenhoven. Algebraic compressed sensing. *Applied and Computational Harmonic Analysis*, 65:374–406, July 2023. ISSN 1063-5203. doi:10.1016/j.acha.2023.03.006.

Jian-Feng Cai, Emmanuel J. Candès, and Zuowei Shen. A singular value thresholding algorithm for matrix completion. *SIAM Journal on Optimization*, 20(4):1956–1982, jan 2010. doi:10.1137/080738970.

Emmanuel J. Candès and Benjamin Recht. Exact matrix completion via convex optimization. *Foundations of Computational Mathematics*, 9(6):717–772, apr 2009. ISSN 1615-3383. doi:10.1007/s10208-009-9045-5.

Emmanuel J. Candes and Terence Tao. The power of convex relaxation: Near-optimal matrix completion. *IEEE Transactions on Information Theory*, 56(5):2053–2080, may 2010. ISSN 0018-9448. doi:10.1109/tit.2010.2044061.

Emmanuel J. Candès, Michael B. Wakin, and Stephen P. Boyd. Enhancing sparsity by reweighted ℓ_1 minimization. *Journal of Fourier Analysis and Applications*, 14(5-6):877–905, oct 2008. ISSN 1531-5851. doi:10.1007/s00041-008-9045-x.

Emmanuel J. Candès, Yonina C. Eldar, Thomas Strohmer, and Vladislav Voroninski. Phase retrieval via matrix completion. *SIAM Journal on Imaging Sciences*, 6(1):199–225, jan 2013. doi:10.1137/110848074.

Rick Chartrand and Valentina Staneva. Restricted isometry properties and nonconvex compressive sensing. *Inverse Problems*, 24(3):035020, may 2008. doi:10.1088/0266-5611/24/3/035020.

Rick Chartrand and Wotao Yin. Iteratively reweighted algorithms for compressive sensing. In *2008 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 3869–3872. IEEE, mar 2008. doi:10.1109/icassp.2008.4518498.

Yuejie Chi, Yue M. Lu, and Yuxin Chen. Nonconvex optimization meets low-rank matrix factorization: An overview. *IEEE Transactions on Signal Processing*, 67(20):5239–5269, oct 2019. ISSN 1053-587X. doi:10.1109/tsp.2019.2937282.

Wei Dai, Olgica Milenkovic, and Ely Kerman. Subspace evolution and transfer (SET) for low-rank matrix completion. *IEEE Transactions on Signal Processing*, 59(7):3120–3132, jul 2011. doi:10.1109/tsp.2011.2144977.

Ingrid Daubechies, Ronald DeVore, Massimo Fornasier, and C. Sinan Güntürk. Iteratively reweighted least squares minimization for sparse recovery. *Communications on Pure and Applied Mathematics*, 63(1):1–38, jan 2010. ISSN 1097-0312. doi:10.1002/cpa.20303.

Mark A. Davenport and Justin Romberg. An overview of low-rank matrix recovery from incomplete observations. *IEEE Journal of Selected Topics in Signal Processing*, 10(4): 608–622, jun 2016. doi:10.1109/jstsp.2016.2539100.

C Davis. On invariant convex functions. *Arch. Math*, 1957.

M. Fazel, H. Hindi, and S. P. Boyd. Log-det heuristic for matrix rank minimization with applications to Hankel and Euclidean distance matrices. In *Proceedings of the 2003 American Control Conference, 2003.*, volume 3, pages 2156–2162 vol.3. IEEE, June 2003. doi:10.1109/acc.2003.1243393.

Maryam Fazel. *Matrix Rank Minimization with Applications*. PhD thesis, Stanford University, Stanford, CA, 2002.

Massimo Fornasier, Holger Rauhut, and Rachel Ward. Low-rank matrix recovery via iteratively reweighted least squares minimization. *SIAM Journal on Optimization*, 21(4): 1614–1640, oct 2011. doi:10.1137/100811404.

Paris Giampouras, Rene Vidal, Athanasios Rontogiannis, and Benjamin Haeffele. A novel variational form of the Schatten-p quasi-norm. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 21453–21463. Curran Associates, Inc., 2020. URL `https://proceedings.neurips.cc/paper/2020/file/f53eb4122d5e2ce81a12093f8f9ce922-Paper.pdf`.

Donald Goldfarb and Shiqian Ma. Convergence of fixed-point continuation algorithms for matrix rank minimization. *Foundations of Computational Mathematics*, 11(2):183–210, feb 2011. ISSN 1615-3383. doi:10.1007/s10208-011-9084-6.

Lars Grasedyck and Sebastian Krämer. Stable ALS approximation in the TT-format for rank-adaptive tensor completion. *Numerische Mathematik*, 143(4):855–904, aug 2019. ISSN 0945-3245. doi:10.1007/s00211-019-01072-4.

Lars Grasedyck, Daniel Kressner, and Christine Tobler. A literature survey of low-rank tensor approximation techniques. *GAMM-Mitteilungen*, 36(1):53–78, aug 2013. ISSN 1522-2608. doi:10.1002/gamm.201310004.

David Gross. Recovering low-rank matrices from few coefficients in any basis. *IEEE Transactions on Information Theory*, 57(3):1548–1566, mar 2011. ISSN 0018-9448. doi:10.1109/tit.2011.2104999.

Trevor Hastie, Rahul Mazumder, Jason D. Lee, and Reza Zadeh. Matrix completion and low-rank SVD via fast alternating least squares. *J. Mach. Learn. Res.*, 16(1):3367–3402, January 2015. ISSN 1532-4435. URL `https://www.jmlr.org/papers/v16/hastie15a.html`.

Raghunandan H. Keshavan, Andrea Montanari, and Sewoong Oh. Matrix completion from noisy entries. *J. Mach. Learn. Res.*, 11:2057–2078, August 2010a. ISSN 1532-4435. URL https://www.jmlr.org/papers/v11/keshavan10a.html.

Raghunandan H. Keshavan, Andrea Montanari, and Sewoong Oh. Matrix completion from a few entries. *IEEE Transactions on Information Theory*, 56(6):2980–2998, jun 2010b. doi:10.1109/tit.2010.2046205.

Sebastian Krämer. Asymptotic log-det sum-of-ranks minimization via tensor (alternating) iteratively reweighted least squares, 2021a. URL https://arxiv.org/abs/2106.15201.

Sebastian Krämer. *MATLAB Central File Exchange*, 9 2021b. URL https://www.mathworks.com/matlabcentral/fileexchange/97697-button-plot.

Christian Kümmerle and Claudio M. Verdun. A scalable second order method for ill-conditioned matrix completion from few samples. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 5872–5883. PMLR, 18–24 Jul 2021. URL https://proceedings.mlr.press/v139/kummerle21a.html.

Christian Kümmerle. *Understanding and Enhancing Data Recovery Algorithms*. Dissertation, Technische Universität München, München, 2019. URL https://mediatum.ub.tum.de/1521436.

Ming-Jun Lai, Yangyang Xu, and Wotao Yin. Improved iteratively reweighted least squares for unconstrained smoothed ℓ_q minimization. *SIAM Journal on Numerical Analysis*, 51 (2):927–957, jan 2013. doi:10.1137/110840364.

Lieven De Lathauwer. A survey of tensor methods. In *2009 IEEE International Symposium on Circuits and Systems*, pages 2773–2776. IEEE, may 2009. doi:10.1109/iscas.2009.5118377.

Kiryung Lee and Yoram Bresler. ADMiRA: Atomic decomposition for minimum rank approximation. *IEEE Transactions on Information Theory*, 56(9):4402–4416, sep 2010. doi:10.1109/tit.2010.2054251.

Dmitry Malioutov and Aleksandr Aravkin. Iterative log thresholding. In *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7198–7202. IEEE, may 2014. doi:10.1109/icassp.2014.6854997.

K Mohan and M Fazel. Reweighted nuclear norm minimization with application to system identification. In *Proceedings of the 2010 American Control Conference*, pages 2953–2959. IEEE, jun 2010. doi:10.1109/acc.2010.5531594.

Karthik Mohan and Maryam Fazel. Iterative reweighted algorithms for matrix rank minimization. *Journal of Machine Learning Research*, 13(Nov):3441–3473, 2012. URL http://www.jmlr.org/papers/v13/mohan12a.html.

B. K. Natarajan. Sparse approximate solutions to linear systems. *SIAM Journal on Computing*, 24(2):227–234, apr 1995. doi:10.1137/s0097539792240406.

Luong Trung Nguyen, Junhan Kim, and Byonghyo Shim. Low-rank matrix completion: A contemporary survey. *IEEE Access*, 7:94215–94237, 2019. doi:10.1109/access.2019.2928130.

Samet Oymak, Karthik Mohan, Maryam Fazel, and Babak Hassibi. A simplified approach to recovery conditions for low rank matrices. In *2011 IEEE International Symposium on Information Theory Proceedings*, pages 2318–2322. IEEE, jul 2011. doi:10.1109/isit.2011.6033976.

B.D. Rao and K. Kreutz-Delgado. An affine scaling methodology for best basis selection. *IEEE Transactions on Signal Processing*, 47(1):187–200, 1999. doi:10.1109/78.738251.

Benjamin Recht. A simpler approach to matrix completion. *Journal of Machine Learning Research*, 12:3413–3430, 2011. URL `http://www.jmlr.org/papers/v12/recht11a.html`.

Benjamin Recht, Maryam Fazel, and Pablo A. Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM Review*, 52(3):471–501, jan 2010a. doi:10.1137/070697835.

Benjamin Recht, Weiyu Xu, and Babak Hassibi. Null space conditions and thresholds for rank minimization. *Mathematical Programming*, 127(1):175–202, oct 2010b. ISSN 1436-4646. doi:10.1007/s10107-010-0422-2.

Igor Rivin. Another simple proof of a theorem of chandler davis, 2002. URL `https://arxiv.org/abs/math/0208223`.

Ruoyu Sun and Zhi-Quan Luo. Guaranteed matrix completion via non-convex factorization. *IEEE Transactions on Information Theory*, 62(11):6535–6579, nov 2016. doi:10.1109/tit.2016.2598574.

Jared Tanner and Ke Wei. Normalized iterative hard thresholding for matrix completion. *SIAM Journal on Scientific Computing*, 35(5):S104–S125, jan 2013. doi:10.1137/120876459.

Jared Tanner, Andrew Thompson, and Simon Vary. Matrix rigidity and the ill-posedness of robust PCA and matrix completion. *SIAM Journal on Mathematics of Data Science*, 1(3):537–554, jan 2019. doi:10.1137/18m1227846.

Bart Vandereycken. Low-rank matrix completion by Riemannian optimization. *SIAM Journal on Optimization*, 23(2):1214–1236, jan 2013. doi:10.1137/110845768.

Yangyang Xu, Wotao Yin, Zaiwen Wen, and Yin Zhang. An alternating direction algorithm for matrix completion with nonnegative factors. *Frontiers of Mathematics in China*, 7 (2):365–384, apr 2012. ISSN 1673-3576. doi:10.1007/s11464-012-0194-5.