

NoisyCUR: An algorithm for two-cost budgeted matrix completion

Dong Hu¹, Alex Gittens¹, and Malik Magdon-Ismail¹

Rensselaer Polytechnic Institute, Troy, NY 12180, USA
{hud3,gittea}@rpi.edu, magdon@cs.rpi.edu

Abstract. Matrix completion is a ubiquitous tool in machine learning and data analysis. Most work in this area has focused on the number of observations necessary to obtain an accurate low-rank approximation. In practice, however, the cost of observations is an important limiting factor, and experimentalists may have on hand multiple modes of observation with differing noise-vs-cost trade-offs. This paper considers matrix completion subject to such constraints: a budget is imposed and the experimentalist’s goal is to allocate this budget between two sampling modalities in order to recover an accurate low-rank approximation. Specifically, we consider that it is possible to obtain low noise, high cost observations of individual entries or high noise, low cost observations of entire columns. We introduce a regression-based completion algorithm for this setting and experimentally verify the performance of our approach on both synthetic and real data sets. When the budget is low, our algorithm outperforms standard completion algorithms. When the budget is high, our algorithm has comparable error to standard nuclear norm completion algorithms and requires much less computational effort.

Keywords: Matrix Completion · Low-rank Approximation · Nuclear Norm Minimization.

1 Introduction

Matrix completion (MC) is a powerful and widely used tool in machine learning, finding applications in information retrieval, collaborative filtering, recommendation systems, and computer vision. The goal is to recover a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ from only a few, potentially noisy, observations $\mathbf{y} \in \mathbb{R}^d$, where $d \ll mn$.

In general, the MC problem is ill-posed, as many matrices may give rise to the same set of observations. Typically the inversion problem is made feasible by assuming that the matrix from which the observations were generated is in fact low-rank, $\text{rank}(\mathbf{A}) = r \ll \min\{m, n\}$. In this case, the number of degrees of freedom in the matrix is $(n + m)r$, so if the observations are sufficiently diverse, then the inversion process is well-posed.

In the majority of the MC literature, the mapping from the matrix to its observations, although random, is given to the user, and the aim is to design algorithms that minimize the sample complexity under these observation models.

Some works have considered modifications of this paradigm, where the user designs the observation mapping themselves in order to minimize the number of measurements needed [6,18,12].

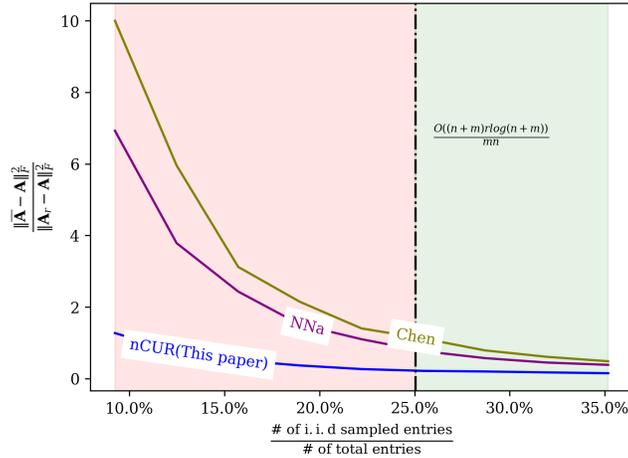


Fig. 1: noisyCUR and nuclear norm MC algorithms applied to the budgeted MC problem on a synthetic incoherent matrix $\mathbf{A} \in \mathbb{R}^{80 \times 60}$. See Section 4 for more details on the experimental setup. The performance of each method at each budget level is averaged over 10 runs, with hyperparameters selected at each value of d by cross-validation. Within the red region the budget is small enough that not enough entries can be sampled for nuclear norm methods to have theoretical support. In this regime, noisyCUR significantly outperforms the baseline algorithms.

This paper considers a budgeted variation of noisy matrix completion with competing observation models: the goal is to allocate a finite budget between the observation models to maximize the accuracy of the recovery. This setup is natural in the context of experimental design: where, given competing options, the natural goal of an experimenter is to spend their budget in a way that maximizes the accuracy of their imputations. Specifically, this work considers a two-cost model where either entire columns can be observed with high per-entry error but low amortized per-entry cost, or entries can be observed with low per-entry error but high per-entry cost.

The noisyCUR algorithm is introduced for this setting, and guarantees are given on its recovery accuracy. Figure 1 is illustrative of the usefulness of noisyCUR in this setting, compared to the standard nuclear norm MC approach. Even in low-budget settings where nuclear norm matrix completion is not theoretically justifiable, noisyCUR satisfies relative-error approximation guarantees. Empirical comparisons of the performance of noisyCUR to that of nuclear norm MC

approaches on a synthetic dataset and on the Jester and MovieLens data sets confirm that noisyCUR can have significantly lower recovery error than standard nuclear norm approaches in the budgeted setting. Additionally, noisyCUR tolerates the presence of coherence in the row-space of the matrix, and is fast as its core computation routine is a ridge regression.

The rest of the paper is organized as follows. In Section 2 we introduce the two-cost model and discuss related works. Section 3 introduces the noisyCUR algorithm and provides performance guarantees; most proofs are deferred to the supplement. Section 4 provides an empirical comparison of the performance of noisyCUR and competing nuclear norm completion algorithms that demonstrates the superior performance of noisyCUR in the limited budget setting. Section 5 concludes the work.

2 Problem Statement

2.1 Notation.

Throughout this paper, scalars are denoted by lowercase letters (n), vectors by bolded lowercase letters (\mathbf{x}), and matrices by bolded uppercase letters \mathbf{A} . The spectral, Frobenius, and nuclear norms of $\mathbf{A} \in \mathbb{R}^{m \times n}$ are written $\|\mathbf{A}\|_2$, $\|\mathbf{A}\|_F$, and $\|\mathbf{A}\|_*$ respectively, and its singular values are ordered in a decreasing fashion: $\sigma_1(\mathbf{A}) \geq \dots \geq \sigma_{\min\{m,n\}}$. The smallest nonzero singular value of \mathbf{A} is denoted by $\sigma_{\min}(\mathbf{A})$. The orthogonal projection onto the range of \mathbf{A} is denoted by $\mathbf{P}_{\mathbf{A}}$. Semidefinite inequalities between the positive-semidefinite matrices \mathbf{A} and \mathbf{B} are written, e.g., as $\mathbf{A} \preceq \mathbf{B}$. The standard Euclidean basis vectors are written as \mathbf{e}_1 , \mathbf{e}_2 and so on.

2.2 Problem formulation

Given a limited budget B with which we can pay to noisily observe a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$, our goal is to construct a low-rank matrix $\bar{\mathbf{A}}$ that approximates \mathbf{A} well. There are two modes of observation: very noisy and cheap samples of entire columns of \mathbf{A} , or much less noisy but expensive samples of individual entries of \mathbf{A} .

The following parameters quantify this two-cost observation model:

- Each low-noise sample of an individual entry costs p_e .
- Each high-noise sample of a column costs $p_c > 0$. Because columns are cheap to sample, $p_c < mp_e$.
- The low-noise samples are each observed with additive $\mathcal{N}(0, \sigma_e^2)$ noise.
- Each entry of the high-noise column samples is observed with additive $\mathcal{N}(0, \sigma_c^2)$ noise. Because sampling columns is noisier than sampling entries, $\sigma_c^2 > \sigma_e^2$.

Algorithm 1 Noisy CUR algorithm for completion of a low-rank matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$

Require: d , the number of column samples; s , the number of row samples; σ_e , the noise level of the column samples; σ_c , the noise level of the row samples; and λ , the regularization parameter

Ensure: $\bar{\mathbf{A}}$, approximation to \mathbf{A}

- 1: $\tilde{\mathbf{C}} \leftarrow \mathbf{C} + \mathbf{E}_c$, where the d columns of \mathbf{C} are sampled uniformly at random¹ from \mathbf{A} , and the entries of \mathbf{E}_c are i.i.d. $\mathcal{N}(0, \sigma_c^2)$.
 - 2: $\mathbf{U} \leftarrow$ an orthonormal basis for $\tilde{\mathbf{C}}$
 - 3: $\ell_i \leftarrow \frac{1}{2} \|\mathbf{e}_i^T \mathbf{U}\|_2^2 / \|\mathbf{U}\|_F^2 + \frac{1}{2m}$ for $i = 1, \dots, m$
 - 4: $\mathbf{S} \leftarrow \text{SamplingMatrix}(\ell, m, s)$
 - 5: $\mathbf{Y} = \mathbf{S}^T \mathbf{A} + \mathbf{E}_e$, where the entries of \mathbf{E}_e are i.i.d. $\mathcal{N}(0, \sigma_e^2)$.
 - 6: $\mathbf{X} \leftarrow \arg \min_{\mathbf{Z}} \|\mathbf{Y} - \mathbf{S}^T \tilde{\mathbf{C}} \mathbf{Z}\|_F^2 + \lambda \|\mathbf{Z}\|_F^2$
 - 7: $\bar{\mathbf{A}} \leftarrow \tilde{\mathbf{C}} \mathbf{X}$
 - 8: **return** $\bar{\mathbf{A}}$
-

2.3 Related Work

The MC literature is vast, so we focus on the works most directly related to the problem setup of this work. To the best of the authors' knowledge, there is no prior work on budgeted low-rank MC. The most related works are adaptive sampling models which attempt to identify more informative observations to reduce the sample complexity [11,3]. One of these works is [6], which estimates non-uniform sampling probabilities for the entries in \mathbf{A} and then samples according to these to reduce the sample complexity. The work [12] similarly estimates the norms of the columns of \mathbf{A} and then samples entries according to these to reduce the sample complexity. The work [18] proposes a similar regression-based algorithm to noisyCUR for noiseless matrix completion.

3 The noisyCUR algorithm

The noisyCUR (nCUR) algorithm, introduced in Algorithm 1, is a regression-based MC algorithm. noisyCUR samples d noisy columns from the low-rank matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$, with entry-wise error that has variance σ_c^2 and collects these columns in the matrix $\tilde{\mathbf{C}} \in \mathbb{R}^{m \times d}$. It $\tilde{\mathbf{C}}$ to approximate the span of \mathbf{A} and returns an approximation of the form $\bar{\mathbf{A}} = \tilde{\mathbf{C}} \mathbf{X}$.

The optimal coefficient matrix \mathbf{X} would require \mathbf{A} to be fully observed. The sample complexity is reduced by instead sampling s rows nosily from \mathbf{A} , with entry-wise error that has variance σ_e^2 , to form the matrix \mathbf{Y} . These rows are then used to estimate the coefficient matrix \mathbf{X} by solving a ridge-regression problem; ridge regression is used instead of least-squares regression because in practice it better mitigates the noisiness of the observations. The rows are sampled according to the shrunk leverage scores [13,15] of the rows of $\tilde{\mathbf{C}}$.

The cost of observing the entries of \mathbf{A} needed to form the approximation $\bar{\mathbf{A}}$ is $dp_c + snp_e$, so the number of column observations d and the number of row observations s are selected to use the entire budget, $B = dp_c + snp_e$.

Our main theoretical result is that when \mathbf{A} is column-incoherent, has low condition number and is sufficiently dense, this algorithm can exploit the two sampling modalities to achieve additive error guarantees in budget regimes where the relative error guarantees of nuclear norm completion do not apply.

First we recall the definition of column-incoherence.

Definition 1 *If $\mathbf{A} \in \mathbb{R}^{m \times n}$ and $\mathbf{V} \in \mathbb{R}^{m \times r}$ is an orthonormal basis for the row space of \mathbf{A} , then the column leverage scores of \mathbf{A} are given by*

$$\ell_i = \|\mathbf{e}_i^T \mathbf{V}\|_2^2 \quad \text{for } i = 1, \dots, n.$$

The column coherence of \mathbf{A} is the maximum of the column leverage scores of \mathbf{A} , and \mathbf{A} is said to have a β -incoherent column space if its column coherence is smaller than $\beta \frac{r}{n}$.

Theorem 1 *Let $\mathbf{A} \in \mathbb{R}^{m \times n}$ be a rank- r matrix with β -incoherent column space. Assume that \mathbf{A} is dense: there is a $c > 0$ such that at least half² of the entries of \mathbf{A} satisfy $|a_{ij}| \geq c$.*

Fix a precision parameter $\varepsilon \in (0, 1)$ and invoke Algorithm 1 with

$$d \geq \max \left\{ \frac{6 + 2\varepsilon}{3\varepsilon^2} \beta r \log \frac{r}{\delta}, \frac{4(1 + \delta)}{c^2(1 - \varepsilon)\varepsilon} r \kappa_2(\mathbf{A})^2 \sigma_c^2 \right\} \quad \text{and } s \geq \frac{6 + 2\varepsilon}{3\varepsilon^2} 2d \log \frac{d}{\delta}.$$

The returned approximation satisfies

$$\|\mathbf{A} - \bar{\mathbf{A}}\|_F^2 \leq \left(\gamma + \varepsilon + 40 \frac{\varepsilon}{1 - \varepsilon} \right) \|\mathbf{A}\|_F^2 + 12\varepsilon \left(\frac{\sigma_e^2}{\sigma_c^2} \right) d \sigma_r^2(\mathbf{A})$$

with probability at least $0.9 - 2\delta - 2 \exp(-\frac{(m-r)\delta^2}{2}) - \exp(-\frac{sn}{32})$. Here,

$$\gamma \leq 2 \left(\frac{1 + \varepsilon}{1 - \varepsilon} \right) \left[\frac{\lambda}{(1 + \varepsilon) \left(\frac{1}{2} \sqrt{m - r} - \sqrt{d} \right)^2 \sigma_c^2 + \lambda} \right]^2.$$

The proof of this result is deferred.

Comparison of nuclear norm and noisy CUR approximation guarantees. Theorem 1 implies that, if \mathbf{A} has low condition number, is dense and column-incoherent, and the regularization parameter λ is selected to be $o((\sqrt{m - r} - \sqrt{d})^2 \sigma_c)$, then a mixed relative-additive bound of the form

$$\begin{aligned} \|\mathbf{A} - \bar{\mathbf{A}}\|_F^2 &\leq \varepsilon' \|\mathbf{A}\|_F^2 + \varepsilon'' \tilde{O}(r) \sigma_r^2(\mathbf{A}) \\ &= \varepsilon' \|\mathbf{A}\|_F^2 + \varepsilon'' \tilde{O}(\|\mathbf{A}\|_F^2) \end{aligned} \quad (1)$$

² This fraction can be changed, with corresponding modification to the sample complexities d and s .

holds with high probability for the approximation returned by Algorithm 1, where ε' and ε'' are $o(1)$, when d and s are $\tilde{\Omega}(r)$.

By way of comparison, the current best guarantees for noisy matrix completion using nuclear norm formulations state that if $d = \Omega((n+m)r \log(n+m))$ entries are noisily observed with noise level σ_e^2 , then a nuclear norm formulation of the matrix completion problem yields an approximation $\bar{\mathbf{A}}$ that satisfies

$$\|\mathbf{A} - \bar{\mathbf{A}}\|_F^2 = O\left(\frac{\sigma_e^2}{\sigma_r^2(\mathbf{A})} \frac{nm}{r}\right) \|\mathbf{A}\|_F^2$$

with high probability [7,2]. The conditions imposed to obtain this guarantee [7] are that \mathbf{A} has low condition number and that both its row and column spaces are incoherent. If we additionally require that \mathbf{A} is dense, so that the assumptions applied to both algorithms are comparable, then $\|\mathbf{A}\|_F^2 = \Omega(mn)$ and the guarantee for nuclear norm completion becomes

$$\|\mathbf{A} - \bar{\mathbf{A}}\|_F^2 = O(\sigma_e^2 \kappa_2^2(\mathbf{A})) \|\mathbf{A}\|_F^2. \quad (2)$$

Comparison of budget requirements for nuclear norm completion and noisy CUR. For nuclear norm completion approaches to assure guarantees of the form (2) it is necessary to obtain $\Omega((n+m)r \log(n+m))$ high precision noisy samples [4], so the budget required is

$$B_{NN} = \Omega((n+m)r \log(n+m)p_e).$$

When B_{NN} exceeds the budget B , there is no theory supporting the use of a mix of more expensive high and cheaper low precision noisy measurements.

The noisyCUR algorithm allows exactly such a mix: the cost of obtaining the necessary samples is

$$B_{nCUR} = dp_c + snp_e = \tilde{\Omega}(r)p_c + \tilde{\Omega}_r(nr)p_e,$$

where the notation $\tilde{\Omega}_r(\cdot)$ is used to indicate that the omitted logarithmic factors depend only on r . It is evident that

$$B_{nCUR} < B_{NN},$$

so the noisy CUR algorithm is applicable in budget regimes where nuclear norm completion is not.

3.1 Proof of Theorem 1

Theorem 1 is a consequence of two structural results that are established in the supplement.

The first result states that if $\text{rank}(\mathbf{C}) = \text{rank}(\mathbf{A})$ and the bottom singular value of \mathbf{C} is large compared to σ_c , then the span of $\tilde{\mathbf{C}}$ will contain a good approximation to \mathbf{A} .

Lemma 1 Fix an orthonormal basis $\mathbf{U} \in \mathbb{R}^{m \times r}$ and consider $\mathbf{A} \in \mathbb{R}^{m \times n}$ and $\mathbf{C} \in \mathbb{R}^{m \times d}$ with factorizations $\mathbf{A} = \mathbf{U}\mathbf{M}$ and $\mathbf{C} = \mathbf{U}\mathbf{W}$, where both \mathbf{M} and \mathbf{W} have full row rank. Further, let $\tilde{\mathbf{C}}$ be a noisy observation of \mathbf{C} , that is, let $\tilde{\mathbf{C}} = \mathbf{C} + \mathbf{G}$ where the entries of \mathbf{G} are i.i.d. $\mathcal{N}(0, \sigma_c^2)$. If $\sigma_{\min}(\mathbf{C}) \geq 2(1 + \delta)\sigma_c\sqrt{m/\varepsilon}$, then

$$\|(\mathbf{I} - \mathbf{P}_{\tilde{\mathbf{C}}})\mathbf{A}\|_F^2 \leq \varepsilon\|\mathbf{A}\|_F^2$$

with probability at least $1 - \exp\left(\frac{-m\delta^2}{2}\right)$.

Recall the definition of a $(1 \pm \varepsilon)$ -subspace embedding.

Definition 2 (Subspace embedding [17]) Let $\mathbf{A} \in \mathbb{R}^{m \times n}$ and fix $\varepsilon \in (0, 1)$. A matrix $\mathbf{S} \in \mathbb{R}^{s \times m}$ is a $(1 \pm \varepsilon)$ -subspace embedding for \mathbf{A} if

$$(1 - \varepsilon)\|\mathbf{x}\|_2^2 \leq \|\mathbf{S}^T \mathbf{x}\|_2^2 \leq (1 + \varepsilon)\|\mathbf{x}\|_2^2$$

for all vectors \mathbf{x} in the span of \mathbf{A} , or equivalently, if

$$(1 - \varepsilon)\mathbf{A}^T \mathbf{A} \preceq \mathbf{A}^T \mathbf{S} \mathbf{S}^T \mathbf{A} \preceq (1 + \varepsilon)\mathbf{A}^T \mathbf{A}.$$

Often we will use the shorthand “subspace embedding” for $(1 \pm \varepsilon)$ -subspace embedding.

The second structural result is a novel bound on the error of sketching using a subspace embedding to reduce the cost of ridge regression, when the target is noisy.

Corollary 1 Let $\tilde{\mathbf{C}} \in \mathbb{R}^{m \times d}$, where $d \leq m$, and $\tilde{\mathbf{A}} = \mathbf{A} + \mathbf{E}$ be matrices, and let \mathbf{S} be an $(1 \pm \varepsilon)$ -subspace embedding for $\tilde{\mathbf{C}}$. If

$$\mathbf{X} = \arg \min_{\mathbf{Z}} \|\mathbf{S}^T (\tilde{\mathbf{A}} - \tilde{\mathbf{C}}\mathbf{Z})\|_F^2 + \lambda\|\mathbf{Z}\|_F^2,$$

then

$$\|\mathbf{A} - \tilde{\mathbf{C}}\mathbf{X}\|_F^2 \leq \|(\mathbf{I} - \mathbf{P}_{\tilde{\mathbf{C}}})\mathbf{A}\|_F^2 + \gamma\|\mathbf{P}_{\tilde{\mathbf{C}}}\mathbf{A}\|_F^2 + \frac{4}{1 - \varepsilon}\|\mathbf{S}^T \mathbf{E}\|_F^2 + \frac{4}{1 - \varepsilon}\|\mathbf{S}^T (\mathbf{I} - \mathbf{P}_{\tilde{\mathbf{C}}})\mathbf{A}\|_F^2,$$

where $\gamma = 2\left(\frac{1+\varepsilon}{1-\varepsilon}\right)\left(\frac{\lambda}{(1+\varepsilon)\sigma_d(\tilde{\mathbf{C}})^2 + \lambda}\right)^2$.

Corollary 1 differs significantly from prior results on the error in sketched ridge regression, e.g. [1,15], in that: (1) it bounds the *reconstruction error* rather than the *ridge regression objective*, and (2) it considers the impact of noise in the target. This result follows from a more general result on sketched noisy proximally regularized least squares problems, stated as Theorem 2 in the supplement.

Together with standard properties of Gaussian noise and subspace embeddings, these two results deliver Theorem 1.

Proof (Proof of Theorem 1). The NoisyCUR algorithm first forms the noisy column samples $\tilde{\mathbf{C}} = \mathbf{C} + \mathbf{E}_c$, where $\mathbf{C} = \mathbf{A}\mathbf{M}$. The random matrix $\mathbf{M} \in \mathbb{R}^{n \times d}$ selects d columns uniformly at random with replacement from the columns of \mathbf{A} , and the entries of $\mathbf{E}_c \in \mathbb{R}^{m \times d}$ are i.i.d. $\mathcal{N}(0, \sigma_c^2)$. It then solves the sketched regression problem

$$\mathbf{X} = \arg \min_{\mathbf{Z}} \|\mathbf{S}^T(\tilde{\mathbf{A}} - \tilde{\mathbf{C}}\mathbf{Z})\|_F^2 + \lambda \|\mathbf{Z}\|_F^2,$$

and returns the approximation $\bar{\mathbf{A}} = \tilde{\mathbf{C}}\mathbf{X}$. Here $\tilde{\mathbf{A}} = \mathbf{A} + \mathbf{E}_e$, where $\mathbf{E}_e \in \mathbb{R}^{m \times n}$ comprises i.i.d. $\mathcal{N}(0, \sigma_e^2)$ entries, and the sketching matrix $\mathbf{S} \in \mathbb{R}^{m \times s}$ samples s rows using the shrunk leverage scores of $\tilde{\mathbf{C}}$.

By [15][Appendix A.1.1], \mathbf{S} is a subspace embedding for $\tilde{\mathbf{C}}$ with failure probability at most δ when s is as specified. Thus Corollary 1 applies and gives that

$$\begin{aligned} \|\mathbf{A} - \tilde{\mathbf{C}}\mathbf{X}\|_F^2 &\leq \|(\mathbf{I} - \mathbf{P}_{\tilde{\mathbf{C}}})\mathbf{A}\|_F^2 + \gamma' \|\mathbf{P}_{\tilde{\mathbf{C}}}\mathbf{A}\|_F^2 \\ &\quad + \frac{4}{1-\varepsilon} \|\mathbf{S}^T \mathbf{E}\|_F^2 + \frac{4}{1-\varepsilon} \|\mathbf{S}^T(\mathbf{I} - \mathbf{P}_{\tilde{\mathbf{C}}})\mathbf{A}\|_F^2 \\ &= T_1 + T_2 + T_3 + T_4, \end{aligned}$$

where $\gamma' = 2 \left(\frac{1+\varepsilon}{1-\varepsilon} \right) \left(\frac{\lambda}{(1+\varepsilon)\sigma_d(\tilde{\mathbf{C}})^2 + \lambda} \right)^2$. We now bound the four terms T_1, T_2, T_3 , and T_4 .

To bound T_1 , note that by [16][Lemma 13], the matrix $\sqrt{\frac{n}{d}}\mathbf{M}$ is a subspace embedding for \mathbf{A}^T with failure probability at most δ when d is as specified. This gives the semidefinite inequality $\frac{n}{d}\mathbf{C}\mathbf{C}^T = \frac{n}{d}\mathbf{A}\mathbf{M}\mathbf{M}^T\mathbf{A}^T \succeq (1-\varepsilon)\mathbf{A}\mathbf{A}^T$, which in turn gives that

$$\begin{aligned} \sigma_r^2(\mathbf{C}) &\geq (1-\varepsilon) \frac{d}{n} \sigma_r^2(\mathbf{A}) \geq \frac{4(1+\delta)r}{c^2\varepsilon} \frac{r}{n} \|\mathbf{A}\|_2^2 \sigma_c^2 \\ &\geq \frac{4(1+\delta)}{c^2\varepsilon n} \|\mathbf{A}\|_F^2 \sigma_c^2 \geq 2(1+\delta) \frac{m}{\varepsilon} \sigma^2. \end{aligned}$$

The second inequality holds because

$$d \geq \frac{4(1+\delta)}{c^2(1-\varepsilon)\varepsilon} r \kappa_2(\mathbf{A})^2 \sigma_c^2 \quad \text{implies} \quad \sigma_r^2(\mathbf{A}) \geq \frac{4(1+\delta)r}{c^2(1-\varepsilon)\varepsilon} \frac{r}{d} \|\mathbf{A}_2\|^2 \sigma_c^2 \quad (3)$$

The third inequality holds because $r\|\mathbf{A}\|_2^2$ is an overestimate of $\|\mathbf{A}_F\|_2^2$. The final inequality holds because the denseness of \mathbf{A} implies that $\|\mathbf{A}\|_F^2 \geq \frac{1}{2}c^2mn$.

Note also that the span of $\mathbf{C} = \mathbf{A}\mathbf{M}$ is contained in that of \mathbf{A} , and since $\frac{n}{d}\mathbf{C}\mathbf{C}^T \succeq (1-\varepsilon)\mathbf{A}\mathbf{A}^T$, in fact \mathbf{C} and \mathbf{A} have the same rank and therefore span the same space. Thus the necessary conditions to apply Lemma 1 are satisfied, and as a result, we find that

$$T_1 \leq \varepsilon \|\mathbf{A}\|_F^2$$

with failure probability at most $\exp(-\frac{m\delta^2}{2})$.

Next we bound T_2 . Observe that $\|\mathbf{P}_{\tilde{\mathcal{C}}}\mathbf{A}\|_F^2 \leq \|\mathbf{A}\|_F^2$. Further, by Lemma 3 in the supplement,

$$\sigma_d(\tilde{\mathcal{C}}) \geq \left(\frac{1}{2}\sqrt{m-r} - \sqrt{d}\right)^2 \sigma_c^2$$

with failure probability at most $\exp(-\frac{(m-r)\delta^2}{2})$. This allows us to conclude that

$$T_2 \leq \gamma \|\mathbf{A}\|_F^2,$$

where γ is as specified in the statement of this theorem.

To bound T_3 , we write

$$\begin{aligned} T_3 &= \frac{4}{1-\varepsilon} \|\mathbf{S}^T \mathbf{P}_S \mathbf{E}\|_F^2 \leq \frac{4}{1-\varepsilon} \|\mathbf{S}\|_2^2 \|\mathbf{P}_S \mathbf{E}\|_F^2 \\ &\leq \frac{8}{1-\varepsilon} \frac{m}{s} \|\mathbf{Q}^T \mathbf{E}\|_F^2, \end{aligned}$$

where \mathbf{Q} is an orthonormal basis for the span of \mathbf{S} . The last inequality holds because [15][Appendix A.1.2] shows that $\|\mathbf{S}\|_2^2 \leq 2\frac{m}{s}$ always. Finally, note that \mathbf{Q} has at most s columns, so in the worst case $\mathbf{Q}^T \mathbf{E}$ comprises sn i.i.d. $\mathcal{N}(0, \sigma_e^2)$ entries. A standard concentration bound for χ^2 random variables with sn degrees of freedom [14][Example 2.11] guarantees that

$$\|\mathbf{Q}^T \mathbf{E}\|_F^2 \leq \frac{3}{2} sn \sigma_e^2$$

with failure probability at most $\exp(-\frac{sn}{32})$. We conclude that, with the same failure probability,

$$T_3 \leq \frac{12}{1-\varepsilon} mn \sigma_e^2.$$

Now recall (3), which implies that

$$\begin{aligned} \varepsilon(1-\varepsilon) d\sigma_r^2(\mathbf{A}) &\geq \frac{4(1+\delta)}{c^2} r \|\mathbf{A}_2\|^2 \sigma_c^2 \geq \frac{4(1+\delta)}{c^2} \|\mathbf{A}\|_F^2 \sigma_c^2 \\ &\geq 2(1+\delta) mn \sigma_c^2 \geq mn \sigma_c^2. \end{aligned}$$

It follows from the last two displays that

$$T_3 \leq 12\varepsilon \left(\frac{\sigma_e^2}{\sigma_c^2}\right) d\sigma_r^2(\mathbf{A}).$$

The bound for T_4 is an application of Markov's inequality. In particular, it is readily verifiable that $\mathbb{E}[\mathbf{S}\mathbf{S}^T]$, which implies that

$$\mathbb{E}T_4 = \frac{4}{1-\varepsilon} \|(\mathbf{I} - \mathbf{P}_{\tilde{\mathcal{C}}})\mathbf{A}\|_F^2 = \frac{4}{1-\varepsilon} T_1 \leq \frac{4\varepsilon}{1-\varepsilon} \|\mathbf{A}\|_F^2.$$

The final inequality comes from the bound $T_1 \leq \varepsilon \|\mathbf{A}\|_F^2$ that was shown earlier. Thus, by Markov's inequality,

$$T_4 \leq \frac{40\varepsilon}{1-\varepsilon} \|\mathbf{A}\|_F^2$$

with failure probability at most 0.1.

Collating the bounds for T_1 through T_4 and their corresponding failure probabilities gives the claimed result.

4 Empirical Evaluation

In this section we investigate the performance of the noisyCUR method on a small-scale synthetic data set and on the Jester and MovieLens data sets. We compare with the performance of three nuclear norm-based algorithms in a low and a high-budget regime.

4.1 Specifying the experimental setups.

Four parameters are manipulated to control the experiment setup:

1. The budget, taken to be of the size $B = c_0 m r p_e$ for some constant positive integer c_0 . This choice ensures that the $O((n+m)r \log(n+m))$ high precision samples needed for nuclear norm completion methods cannot be obtained.
2. The ratio of the cost of sampling a column to that of individually sampling each entry in that column, $\alpha = \frac{p_c}{m p_e}$. For all three experiments, we set $\alpha = 0.2$.
3. The entry sampling noise level σ_e^2 .
4. The column sampling noise level σ_c^2 .

Based on the signal-to-noise ratio between the matrix and the noise level of the noisiest observation model, σ_c^2 , we classify an experiment as being high noise or low noise. The entry-wise signal-to-noise ratio is given by

$$SNR = \frac{\|\mathbf{A}\|_F^2}{mn\sigma_c^2}.$$

High SNR experiments are said to be low noise, while those with low SNR are said to be high noise.

4.2 Methodology: noisyCUR and the baselines.

Given the settings of the two-cost model, the noisyCUR algorithm is employed by selecting a value for d , the number of noisy column samples; the remaining budget then determines s , the number of rows that are sampled with high precision. Cross-validation is used to select the regularization parameter λ .

We compare to three nuclear norm-based MC algorithms, as nuclear norm-based approaches are the most widely used and theoretically investigated algorithms for low-rank MC. To explain the baselines, we introduce some notation. The set Ω_s below comprises the indices of entries of \mathbf{A} sampled with high accuracy, while Ω_c comprises the indices of entries of \mathbf{A} sampled using the low accuracy column observation model. Given a set of indices Ω , the operator

$\mathcal{P}_\Omega : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^{m \times n}$ returns a matrix whose values are the same as those of the input matrix on the indices in Ω , and zero on any indices not in Ω .

The first of the baselines is the formulation introduced in [5], which forms the approximation

$$\begin{aligned} \bar{\mathbf{A}} &= \arg \min_{\mathbf{Z}} \|\mathbf{Z}\|_* \\ \text{s.t. } &\|\mathcal{P}_{\Omega_e}(\mathbf{Z} - \mathbf{Y})\|_F \leq \delta, \quad (i, j) \in \Omega_e \end{aligned} \quad (4)$$

Given the settings of the two-cost model, all the budget is spent on sampling entries to form Ω_e . Thus the performance of this model is a constant independent of d in the figures. The hyperparameter δ is selected through cross-validation. This baseline is referred to as NNa (nuclear norm minimization for all entries) in the figures.

The second baseline is a nuclear norm formulation that penalizes the two forms of observations separately, forming the approximation

$$\begin{aligned} \bar{\mathbf{A}} &= \min_{\mathbf{Z}} \|\mathbf{Z}\|_* \\ \text{s.t. } &\|\mathcal{P}_{\Omega_c}(\mathbf{Z} - \mathbf{A}_{obs})\|_F^2 \leq C_1 dm \gamma \sigma_e^2 \\ &\|\mathcal{P}_{\Omega_e}(\mathbf{Z} - \mathbf{A}_{obs})\|_F^2 \leq C_2 f \sigma_e^2 \end{aligned} \quad (5)$$

where C_1 and C_2 are parameters, and \mathbf{A}_{obs} is the partially observed matrix which contains the observed noisy column samples and entry samples and is zero elsewhere. As with the noisyCUR method, given a value of d , the remaining budget is spent on sampling s rows with high precision. The hyperparameters C_1 and C_2 are selected through cross-validation. This baseline is referred to as NNs (nuclear norm split minimization) in the figures.

The final baseline is the two-phase sampling method of [6]. This method spends a portion of the budget to randomly sample entries to estimate leverage scores, then uses the rest of the budget to sample entries according to the leverage score and reconstructs from the union of these samples using the same optimization as NNa. The performance is therefore independent of d . This baseline is referred to as Chen in the figures.

The details of cross-validation of the parameters for the nuclear norm methods are omitted because there are many relevant hyperparameters: in addition to the constraint parameters in the optimization formulations, there are important hyperparameters associated with the ADMM solvers used (e.g., the Lagrangian penalty parameters).

4.3 Synthetic Dataset

Figure 2 compares the performance of the three baseline methods and noisyCUR on an incoherent synthetic data set $\mathbf{A} \in \mathbb{R}^{80 \times 60}$ generated by sampling a matrix with i.i.d. $\mathcal{N}(5, 1)$ entries and taking its best rank four approximation. For each value of d , the regularization parameter λ of noisyCUR is selected via cross-validation from 500 logarithmically spaced points in the interval $(10^{-4}, 10)$.

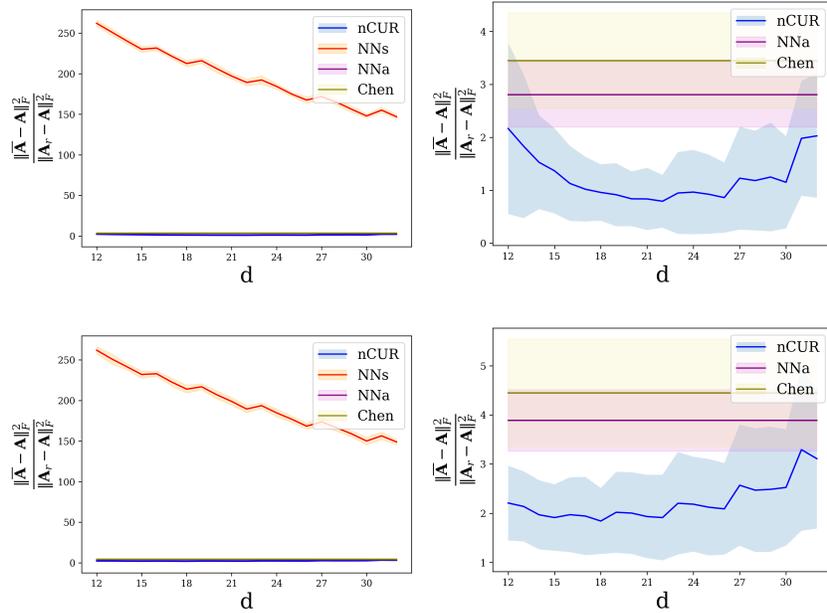


Fig. 2: Performance on the synthetic dataset. $\beta = 15\%$ for all plots. The noise level is low in the top plots: $\sigma_e^2 = 0.01, \sigma_c^2 = 0.05$. The upper left plot shows all methods, while the upper right plot removes the NNs method to facilitate comparison of the better performing methods. In the bottom two plots, the noise level is higher: $\sigma_e^2 = 0.04, \sigma_c^2 = 0.2$. Similarly, the bottom left plot shows all methods, while the bottom right removes the NNs method. Each point in the plots is the average of 100 runs.

4.4 Jester

Figure 3 compares the performance of the baseline methods and noisyCUR on a subset of the Jester dataset of [9]. Specifically the data set was constructed by extracting the submatrix comprising the 7200 users who rated all 100 jokes. For each value of d , the regularization parameter λ of noisyCUR is selected via cross-validation from 200 points logarithmically spaced in the interval $(10, 10^5)$.

4.5 Movielens-100K

Figure 4 compares the performance of the baseline methods and noisyCUR on the Movielens-100K dataset of [10]. The original 1682×943 data matrix is quite sparse, so complete it to obtain a low-rank matrix using the iterative low-rank SVD algorithm of [8] before applying noisyCUR and the three baselines. For each value of d , the parameter λ of noisyCUR is cross-validated among 200 points linearly spaced in the interval $[1, 200]$.

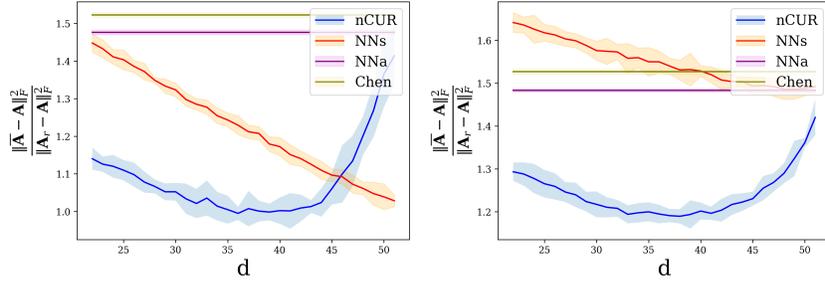


Fig. 3: Performance on the Jester data set. $\beta = 11\%$ for all plots. The noise level is low in the top plots: $\sigma_e^2 = 0.04, \sigma_c^2 = 2$. In the bottom two plots, the noise level is higher: $\sigma_e^2 = 0.25, \sigma_c^2 = 12.5$. As in Figure 2, the plots to the left contain the NNs baseline while those to the right do not. Each point in the plots is the average over 10 runs.

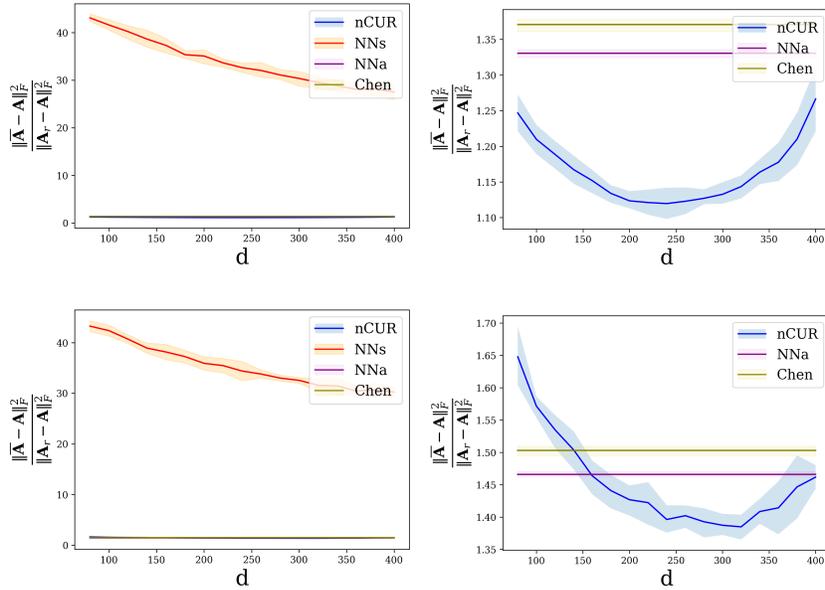


Fig. 4: Performance on the MovieLens-100K data set. $\beta = 10.6\%$ for all plots. The noise level is low in the upper two plots: $\sigma_e^2 = 0.003, \sigma_c^2 = 0.06$. In the bottom two plots, the noise level is higher: $\sigma_e^2 = 0.0225, \sigma_c^2 = 0.65$. As in Figure 2, the plots to the left show the NNs baseline while the plots to the right do not. Each point in the plots is the average of 10 runs.

4.6 Observations

In both the lower and higher noise regimes on all the datasets, noisyCUR exhibits superior performance when compared to the baselines, and exhibits a v-shaped curve, where the optimal approximation error is achieved at the bottom of the v. This convex shape of the performance of noisyCUR with respect to d suggests that there may be a single optimal d given a dataset and the parameters of the two-cost model.

We note a practical consideration that arose in the empirical evaluations: the noisyCUR method is much faster than the nuclear norm algorithms, as they invoke an iterative solver (ADMM was used in these experiments) that computes SVDs of large matrices during each iteration, while noisyCUR solves a single ridge regression problem.

5 Conclusion

This paper introduced the noisyCUR algorithm as an approach to solving the budgeted matrix completion problem where the two observation modalities are high accuracy but expensive entry samples and low accuracy but cheap column samples. We showed that recovery guarantees for noisyCUR hold even in low budget regimes where standard nuclear norm completion approaches are not feasible. The noisyCUR algorithm is fast, as the main computation involved is a ridge regression. Empirically, it was shown that noisyCUR has lower reconstruction error than standard nuclear norm completion baselines in the low-budget setting. It is an interesting open problem to determine optimal or near-optimal values for d and s given the parameters of the two-cost model (B , p_e , p_c , σ_e , and σ_c). An implementation of noisyCUR and the baseline methods is available at anonymized URL, along with code for replicating our experiments.

References

1. Avron, H., Clarkson, K.L., Woodruff, D.P.: Sharper bounds for regularized data fitting. *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques* (2017)
2. Balcan, M.F., Liang, Z., Song, Y., Woodruff, D.P., Zhang, H.: Non-convex matrix completion and related problems via strong duality. *Journal of Machine Learning Research* **20**(102), 1–56 (2019)
3. Balcan, M.F., Zhang, H.: Noise-tolerant life-long matrix completion via adaptive sampling. In: *Advances in Neural Information Processing Systems* 29, pp. 2955–2963. Curran Associates, Inc. (2016)
4. Candés, E.J., Tao, T.: The power of convex relaxation: Near-optimal matrix completion. *IEEE Transactions on Information Theory* **56**(5), 2053–2080 (2010)
5. Candés, E.J., Plan, Y.: Matrix completion with noise. *Proceedings of the IEEE* **98**(6), 925–936 (June 2010)
6. Chen, Y., Bhojanapalli, S., Sanghavi, S., Ward, R.: Coherent matrix completion. *Proceedings of The 31st International Conference on Machine Learning (ICML)* pp. 674–682 (2014)

7. Chen, Y., Chi, Y., Fan, J., Ma, C., Yan, Y.: Noisy matrix completion: Understanding statistical guarantees for convex relaxation via nonconvex optimization. arXiv preprint arXiv:1902.07698 (2019)
8. Cho, K., Reyhani, N.: An iterative algorithm for singular value decomposition on noisy incomplete matrices. International Joint Conference on Neural Networks (IJCNN) pp. 1–6 (2012)
9. Goldberg, K., Roeder, T., Gupta, D., Perkins, C.: An iterative algorithm for singular value decomposition on noisy incomplete matrices. Information Retrieval **6**(2), 133–151 (2001)
10. Harper, F.M., Konstan, J.A.: The movielens datasets: History and context. ACM Transactions on Interactive Intelligent Systems (TiiS) 5 (2015)
11. Krishnamurthy, A., Singh, A.: Low-rank matrix and tensor completion via adaptive sampling. In: Advances in Neural Information Processing Systems 26, pp. 836–844. Curran Associates, Inc. (2013)
12. Krishnamurthy, A., Singh, A.R.: On the power of adaptivity in matrix completion and approximation. ArXiv [abs/1407.3619](https://arxiv.org/abs/1407.3619) (2014)
13. Ma, P., Mahoney, M.W., Yu, B.: A statistical perspective on algorithmic leveraging. The Journal of Machine Learning Research **16**(1), 861–911 (2015)
14. Wainwright, M.J.: High-dimensional statistics: A non-asymptotic viewpoint, vol. 48. Cambridge University Press (2019)
15. Wang, S., Gittens, A., Mahoney, M.W.: Sketched ridge regression: Optimization perspective, statistical perspective, and model averaging. Journal of Machine Learning Research **18**(218), 1–50 (2018)
16. Wang, S., Zhang, Z., Zhang, T.: Towards more efficient spsd matrix approximation and cur matrix decomposition. The Journal of Machine Learning Research **17**(1), 7329–7377 (2016)
17. Woodruff, D.P.: Sketching as a tool for numerical linear algebra. Foundations and Trends® in Theoretical Computer Science **10**(1–2), 1–157 (2014)
18. Xu, M., Jin, R., Zhou, Z.H.: Cur algorithm for partially observed matrices. In: International Conference on Machine Learning. pp. 1412–1421 (2015)