Supplemental Manuscript for Unlabeled Principal Component Analysis (NeurIPS 2021)

Anonymous Author(s)
Affiliation
Address
email

1 Proofs of Theorems

1.1 Proof of Theorem 1

We first prove the theorem over $\mathbb{C}$, then we transfer the statement over $\mathbb{R}$. We note here that there is nothing special about $\mathbb{R}$ and $\mathbb{C}$ with regards to the problem. Indeed, the same proof applies if one replaces $\mathbb{R}$ with any infinite field $\mathbb{F}$ and $\mathbb{C}$ with the algebraic closure $\bar{\mathbb{F}}$ of $\mathbb{F}$. Set

$$\mathcal{M}_{\mathbb{C}} = \{\, X \in \mathbb{C}^{m\times n} | \operatorname{rank}_{\mathbb{C}} X \leq r \,\}$$

and note that since $\mathcal{M}_{\mathbb{C}}$ is irreducible, the intersection of finitely many non-empty open sets in $\mathcal{M}_{\mathbb{C}}$ is itself non-empty and open, and thus dense. Here irreducibility means that $\mathcal{M}_{\mathbb{C}}$ can not be decomposed as the union of two proper subvarieties of $\mathcal{M}_{\mathbb{C}}$.

Lemma 1. *There is an open dense set $\mathcal{U}_1$ in $\mathcal{M}_{\mathbb{C}}$ such that for any $X \in \mathcal{U}_1$ and any $\underline{\pi} = (\Pi_1, \dots, \Pi_n) \in \prod_{i\in[n]} \mathcal{P}_m$, every $m \times r$ submatrix of $\underline{\pi}(X)$ has rank r.*

Proof. First, fix some $\underline{\pi} = (\Pi_1, \dots, \Pi_n) \in \prod_{i\in[n]} \mathcal{P}_m$ and then some index set $\mathcal{J} = \{j_1, \dots, j_r\} \subset [n]$. The submatrix $\underline{\pi}(X)_{\mathcal{J}} := [\Pi_{j_1} x_{j_1}, \cdots, \Pi_{j_r} x_{j_r}]$ of $\underline{\pi}(X)$ has rank less than r if and only if all of its $r \times r$ minors are zero. For each subset $\mathcal{I} = \{i_1, \dots, i_r\} \subset [m]$ we have a polynomial $\det \underline{\pi}(Z)_{\mathcal{I},\mathcal{J}} \in \mathbb{C}[Z]$ where $\underline{\pi}(Z)_{\mathcal{I},\mathcal{J}}$ is the row-submatrix of $\underline{\pi}(Z)_{\mathcal{J}}$ obtained by selecting the rows with index in $\mathcal{I}$. The set of matrices in $\mathbb{C}^{m\times n}$ for which the evaluation of this polynomial is non-zero is an open set, call it $\mathcal{U}_{\underline{\pi},\mathcal{I},\mathcal{J}}$. Then $\underline{\pi}(X)_{\mathcal{J}}$ has rank r if and only if $X \in \mathcal{U}_{\underline{\pi},\mathcal{J}} := \bigcup_{\mathcal{I}} \mathcal{U}_{\underline{\pi},\mathcal{I},\mathcal{J}}$, where $\mathcal{I}$ ranges over all subsets of $[m]$ of cardinality r. As a union of finitely many open sets, $\mathcal{U}_{\underline{\pi},\mathcal{J}}$ is open. Moreover, every $m \times r$ submatrix of $\underline{\pi}(X)$ has rank r if and only if $X \in \mathcal{U}_{\underline{\pi}} := \bigcap_{\mathcal{J}} \mathcal{U}_{\underline{\pi},\mathcal{J}}$, where now $\mathcal{J}$ ranges over all subsets of $[n]$ of cardinality r. $\mathcal{U}_{\underline{\pi}}$ is open because it is the finite intersection of open sets. Finally, every $m \times r$ submatrix of $\underline{\pi}(X)$ has rank r for any $\underline{\pi}$ if and only if X is in the open set $\mathcal{U}_1 := \bigcap_{\underline{\pi}} \mathcal{U}_{\underline{\pi}}$, where the intersection is taken over all $\underline{\pi}$'s.

The proof will be complete once we show that $\mathcal{U}_1$ is non-empty. By what we said above about intersections of finitely many non-empty open sets in an irreducible variety, it is enough to show that each $\mathcal{U}_{\underline{\pi},\mathcal{J}}$ is non-empty. We do this by constructing a specific $X \in \mathcal{U}_{\underline{\pi},\mathcal{J}}$. Recall here that any $\Pi \in \mathcal{P}_m$ is diagonalizable over $\mathbb{C}$ with non-zero eigenvalues. It is an elementary fact in linear algebra that there exists a choice of eigenvector v_k of Π_{j_k} for every $k \in [r]$ such that $v_1, \dots, v_r$ are linearly independent. Now our X is taken to be the matrix with v_k at column j_k for every $k \in [r]$ and zero everywhere else. Clearly $X \in \mathcal{M}_{\mathbb{C}}$ and moreover $\underline{\pi}(X)_{\mathcal{J}} = [\Pi_{j_1} x_{j_1} \cdots, \Pi_{j_r} x_{j_r}] = [\Pi_{j_1} v_1 \cdots, \Pi_{j_r} v_r] = [\lambda_1 v_1 \cdots \lambda_r v_r]$, where λ_k is the corresponding eigenvalue of v_k. Since none of the λ_k's is zero, this matrix has rank r, that is $X \in \mathcal{U}_{\underline{\pi},\mathcal{J}}$. $\square$

Denote by $\mathcal{C}(X)$ the column-space of X and I_m the identity matrix of size $m \times m$. Note also that whenever p is a non-zero polynomial in ν variables with coefficients in $\mathbb{C}$, there is always some $\xi \in \mathbb{C}^\nu$ such that $p(\xi) \neq 0$.

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

Lemma 2. *There is an open dense set $\mathcal{U}_2$ in $\mathcal{M}_{\mathbb{C}}$ such that for any $X \in \mathcal{U}_2$, we have that $\Pi x_j \notin \mathcal{C}(X)$ for any $\Pi \in \mathcal{P}_m \setminus \{I_m\}$ and any $j \in [n]$.*

Proof. $\Pi x_j \notin \mathcal{C}(X)$ if and only if $\operatorname{rank}[X\ \Pi x_j] = r+1$. As in the proof of Lemma 1, this condition is met on an open set $\mathcal{U}_{\Pi,j}$ of $\mathcal{M}_{\mathbb{C}}$ where some $(r+1)\times(r+1)$ determinant of $[X\ \Pi x_j]$ is non-zero. Then the statement of the theorem is true on the open set $\mathcal{U}_2 = \bigcap_{\Pi \in \mathcal{P}_m,\, j \in [n]} \mathcal{U}_{\Pi,j}$. As in the proof of Lemma 1, to show that $\mathcal{U}_2$ is non-empty it suffices to show that each $\mathcal{U}_{\Pi,j}$ is non-empty. We show the existence of an $X \in \mathcal{U}_{\Pi,j}$. Let $Z = (z_{ik})$ be an $m \times r$ matrix of variables over $\mathbb{C}$ and consider the polynomial ring $\mathbb{C}[Z]$. Let us write z_k for the kth column of Z. Since Π is not the identity, there exists some $i \in [m]$ such that z_{i1} is different from the ith element of Πz_1, where z_1 is the first column of Z. Instead, suppose that the variable z_{i1} appears in the i'th coordinate of Πz_1 with $i' \neq i$. Now take any $\mathcal{I} \subset [m]$ with cardinality $r+1$ such that $i, i' \in \mathcal{I}$ and consider $\det[Z\ \Pi z_1]_{\mathcal{I}}$ where $[Z\ \Pi z_1]_{\mathcal{I}}$ is the submatrix of $[Z\ \Pi z_1]$ obtained by selecting the rows with index in $\mathcal{I}$. This is a polynomial of $\mathbb{C}[Z]$ that has the form $\pm z_{i1}^2 \det[z_2 \cdots z_r]_{\mathcal{I}\setminus\{i,i'\}} + \cdots$ where the remaining terms do not involve z_{i1}^{ν} for $\nu > 1$. Since the entries of Z are algebraically independent, $\det[z_2 \cdots z_r]_{\mathcal{I}\setminus\{i,i'\}}$ is a non-zero polynomial. We conclude that $\det[Z\ \Pi z_1]_{\mathcal{I}}$ is also a non-zero polynomial. Hence there exists some $Z' \in \mathbb{C}^{m\times r}$ such that $\det[Z'\ \Pi z_1']_{\mathcal{I}} \neq 0$. Now define X by setting $x_j = z_1'$, $x_{j_k} = z_k'$, $k \in [r]$ for any choice of j_k's distinct from j, and zeros everywhere else. By construction $X \in \mathcal{U}_{\Pi,j}$. □

Let $f : \mathbb{C}^{m\times r} \times \mathbb{C}^{r\times n} \to \mathcal{M}_{\mathbb{C}}$ be the surjective map given by $f(B', C') = B'C'$.

Lemma 3. *There is an open dense set $\mathcal{U}_3$ in $\mathcal{M}_{\mathbb{C}}$ such that for any $X \in \mathcal{U}_3$, we have that for any $j \in [n]$, any $\mathcal{J} = \{j_1, \dots, j_r\} \subset [n]$ with $j \notin \mathcal{J}$ and any $\Pi_1, \dots, \Pi_r \in \mathcal{P}_m$ not all identities, it holds that* $\operatorname{rank}[x_j\ \Pi_1 x_{j_1} \cdots \Pi_r x_{j_r}] = r+1$.

Proof. With $j, \mathcal{J}$ and Π_k's fixed, the set $\mathcal{U}_{j,\mathcal{J},\Pi_1,\dots,\Pi_r}$ of X's in $\mathcal{M}_{\mathbb{C}}$ for which the rank of $[x_j\ \Pi_1 x_{j_1} \cdots \Pi_r x_{j_r}]$ is $r+1$, is open. Indeed, this is defined by the non-simultaneous vanishing of all $(r+1)\times(r+1)$ minors of $[z_j\ \Pi_1 z_{j_1} \cdots \Pi_r z_{j_r}]$, where z_k is the kth column of the matrix of variables Z from the proof of Lemma 2. We note that these are polynomials in Z with integer coefficients. Set $\mathcal{U}_3 = \bigcap_{j,\mathcal{J},\Pi_1,\dots,\Pi_r} \mathcal{U}_{j,\mathcal{J},\Pi_1,\dots,\Pi_r}$ where the intersection is taken over all choices of $j, \mathcal{J}, \Pi_1, \dots, \Pi_r$ as in the statement of the lemma. As in the proof of Lemma 1, the set $\mathcal{U}_3$ is open and to show that it is non-empty is suffices to show that each $\mathcal{U}_{j,\mathcal{J},\Pi_1,\dots,\Pi_r}$ is non-empty.

Let $\mathcal{U}_1, \mathcal{U}_2$ be the open sets of Lemmas 1 and 2. Since $\mathcal{M}_{\mathbb{C}}$ is irreducible and $\mathcal{U}_1, \mathcal{U}_2$ are open and non-empty, we have that $\mathcal{U}_1 \cap \mathcal{U}_2$ is non-empty. Since f is surjective, $f^{-1}(\mathcal{U}_1 \cap \mathcal{U}_2)$ is also non-empty. Take any $(B', C') \in f^{-1}(\mathcal{U}_1 \cap \mathcal{U}_2)$. By definition, the rank of $[\Pi_1 B' c_{j_1}' \cdots \Pi_r B' c_{j_r}']$ is r. By hypothesis, there is some $k \in [r]$ such that Π_k is not the identity and thus again by definition we have $\operatorname{rank}[B'\ \Pi_k B' c_{j_k}'] = r+1$. Consequently, $\Pi_k B' c_{j_k}' \notin \mathcal{C}(B')$ and so the two r-dimensional subspaces $\mathcal{C}(B')$ and $\mathcal{C}\big([\Pi_1 B' c_{j_1}' \cdots \Pi_r B' c_{j_r}']\big)$ are distinct. Thus there exists some $c'' \in \mathbb{C}^r$ such that $B'c'' \notin \mathcal{C}\big([\Pi_1 B' c_{j_1}' \cdots \Pi_r B' c_{j_r}']\big)$. Define $C'' \in \mathbb{C}^{r\times n}$ by setting $c_\nu'' = c_\nu'$ for every $\nu \neq j$ and $c_j'' = c''$. Then by construction $B'C'' \in \mathcal{U}_{j,\mathcal{J},\Pi_1,\dots,\Pi_r}$. □

Take $X^* = [x_1^* \cdots x_n^*] \in \mathcal{U}_3$ and let $\tilde{X} = [\tilde{\Pi}_1 x_1^* \cdots \tilde{\Pi}_n x_n^*]$. Now $\operatorname{rank} \tilde{X} = \operatorname{rank} \tilde{\Pi}_1^{-1}\tilde{X} = \operatorname{rank}[x_1^*\ \tilde{\Pi}_1^{-1}\tilde{\Pi}_2 x_2^* \cdots \tilde{\Pi}_1^{-1}\tilde{\Pi}_n x_n^*]$. If there is some $k \geq 2$ such that $\tilde{\Pi}_1 \neq \tilde{\Pi}_k$, by Lemma 3 any $m \times (r+1)$ submatrix of $\tilde{\Pi}_1^{-1}\tilde{X}$ that contains columns 1 and k will have rank $r+1$. On the other hand, when all $\tilde{\Pi}_k$'s are equal for $k \in [n]$, the rank of $\tilde{X}$ is r by Lemma 1. This concludes the proof of the theorem over $\mathbb{C}$ with the claimed open set being $\mathcal{U}_3$, which we denote in the sequel by $\mathcal{U}_{\mathbb{C}}$.

Set $\mathcal{M}_{\mathbb{R}} = \{X \in \mathbb{R}^{m\times n} | \operatorname{rank}_{\mathbb{R}} X \leq r\}$. There is an inclusion of sets $i : \mathcal{M}_{\mathbb{R}} \hookrightarrow \mathcal{M}_{\mathbb{C}}$ where for $X \in \mathcal{M}_{\mathbb{R}}$ we view $i(X)$ as the complex matrix associated to X. The reason for this inclusion is that if the columns of X generate an r-dimensional subspace over $\mathbb{R}$, then they generate an r-dimensional subspace over $\mathbb{C}$. To finish the proof, it suffices to show the existence of a non-empty open set $\mathcal{U}_{\mathbb{R}}$ in $\mathcal{M}_{\mathbb{R}}$ such that $i(\mathcal{U}_{\mathbb{R}}) \subset \mathcal{U}_{\mathbb{C}}$. This comes from two key ingredients. The first one is the observation that the polynomials that induce $\mathcal{U}_{\mathbb{C}}$, i.e. the polynomials of $\mathbb{C}[Z]$ whose non-simultaneous vanishing indicates membership of a point $X \in \mathcal{M}_{\mathbb{C}}$ in $\mathcal{U}_{\mathbb{C}}$, they have integer and thus real coefficients. This can be seen by inspecting the proof of Lemma 3. Call the set of these polynomials $\mathfrak{p}_{\mathcal{U}} \subset \mathbb{Z}[Z]$. For the second ingredient, let $\mathfrak{p}_{\mathcal{M}} \subset \mathbb{Z}[Z]$ be the set of all $(r+1)\times(r+1)$ minors of the matrix of variables Z. It is a matter of linear algebra that $\mathcal{M}_{\mathbb{R}}$ and $\mathcal{M}_{\mathbb{C}}$ are the common roots of the polynomial system $\mathfrak{p}_{\mathcal{M}}$ over $\mathbb{R}^{m\times n}$ and $\mathbb{C}^{m\times n}$ respectively. What is instead a difficult theorem in commutative algebra is that the following algebraic converse is true; see section 2.6 in [7]: a polynomial $q \in \mathbb{R}[Z]$ vanishes on every point of $\mathcal{M}_{\mathbb{R}}$ if and only if it is a polynomial combination of elements of $\mathfrak{p}_{\mathcal{M}}$, that is if and only if $q = \sum_{p\in\mathfrak{p}_{\mathcal{M}}} c_p\, p$ for some c_p's in $\mathbb{R}[Z]$. This statement also holds true if we replace $\mathbb{R}$ with $\mathbb{C}$. Now the set $\mathcal{U}_{\mathbb{C}}$ consists of those points of $\mathcal{M}_{\mathbb{C}}$ that are roots of the polynomial system $\mathfrak{p}_{\mathcal{M}}$ but not of $\mathfrak{p}_{\mathcal{U}}$. Since $\mathcal{U}_{\mathbb{C}}$ is non-empty, not all polynomials in $\mathfrak{p}_{\mathcal{U}}$ are polynomial combinations of $\mathfrak{p}_{\mathcal{M}}$. But then, by what we just said, not all points of $\mathcal{M}_{\mathbb{R}}$ are common roots of $\mathfrak{p}_{\mathcal{U}}$. This means that the open set of $\mathcal{M}_{\mathbb{R}}$ defined by the non-simultaneous vanishing of all polynonials in $\mathfrak{p}_{\mathcal{U}}$ is non-empty. This open set is the claimed $\mathcal{U}$.

1.2 Proof of Theorem 2

Let $\mathcal{U}_1$ be the open set of Theorem 1. Let $\mathcal{U}_2$ be the set of X's for which $\mathcal{C}(X)$ does not drop dimension under projection onto any r coordinates. This set is open in $\mathcal{M}$ because $X \in \mathcal{U}_2$ if and only if for any $\mathcal{I} \subset [m]$ of cardinality r not all $r\times r$ minors of $X_{\mathcal{I}}$ are zero, $X_{\mathcal{I}}$ being the row-submatrix of X obtained by selecting the rows with index in $\mathcal{I}$. Set $\mathcal{U} = \mathcal{U}_1 \cap \mathcal{U}_2$. Then for any $X^* \in \mathcal{U}$ and any $\Pi \in \mathcal{P}_m$ there exists a unique factorization $\Pi X^* = B^*_\Pi C^*_\Pi$ with the top $r \times r$ block of $B^*_\Pi \in \mathbb{R}^{m\times r}$ being the identity. Since $\bar{p}_{\ell,j}(\tilde{X}) = \bar{p}_{\ell,j}(X^*) = \bar{p}_{\ell,j}(\Pi X^*) = \bar{p}_{\ell,j}(B^*_\Pi C^*_\Pi)$ we have that $(B^*_\Pi, C^*_\Pi) \in \mathcal{Y}_{X^*}$ for every $\Pi \in \mathcal{P}_m$. For the reverse direction we recall a fundamental fact:

Lemma 4. *Fix any* $j \in [n]$. *Suppose that* $\xi_1, \xi_2 \in \mathbb{R}^m$ *are such that* $\bar{p}_{\ell,j}(\xi_1) = \bar{p}_{\ell,j}(\xi_2)$ *for every* $\ell \in [m]$. *Then* $\xi_1 = \Pi\xi_2$ *for some* $\Pi \in \mathcal{P}_m$.

Proof. See proof of Lemma 2 in [6]. □

Now let $(B', C') \in \mathcal{Y}_{X^*}$ and write c'_j for the jth column of C'. For a fixed $j \in [n]$ the equations $q_{\ell,j}(B', C') = 0$ are equivalent to $\bar{p}_{\ell,j}(B'c'_j) = \bar{p}_{\ell,j}(x^*_j)$ for every $\ell \in [m]$. By Lemma 4 there must exist some $\Pi_j \in \mathcal{P}_m$ such that $B'c'_j = \Pi_j x^*_j$. This is true for every $j \in [n]$ so that $B'C' = [\Pi_1 x^*_1 \cdots \Pi_n x^*_n]$. This implies that $\operatorname{rank}[\Pi_1 x^*_1 \cdots \Pi_n x^*_n] = r$. Since $X^* \in \mathcal{U}$, Theorem 1 gives that all Π_j's must be the same permutation $\Pi \in \mathcal{P}_m$, so that $B'C' = \Pi X^*$. Since by construction for any $(B'', C'') \in \mathcal{Y}_{X^*}$ the top $r \times r$ block of B'' is the identity, we have that $B' = B^*_\Pi$ and thus necessarily $C' = C^*_\Pi$.

1.3 Proof of Theorem 3

We first notice $\#\{\tilde{x}_j \,|\, \tilde{x}_j \in S^*;\, j \in [n]\} \geq \mu(I_m) \geq r + 1$. Now we suppose $\tilde{x}_{j_1}, \dots, \tilde{x}_{j_r}, \tilde{x}_{j_{r+1}}$ are $r + 1$ points in $\tilde{X}$ such that not all $\Pi_{j_1}, \dots, \Pi_{j_r}, \Pi_{j_{r+1}}$ are the identity I_m. Since $\mu(\Pi) < r$ for $\Pi \neq I_m$, it is impossible that $\Pi_{j_1} = \cdots = \Pi_{j_r} = \Pi_{j_{r+1}}$. According to Theorem 1, the points $\tilde{x}_{j_1}, \dots, \tilde{x}_{j_r}, \tilde{x}_{j_{r+1}}$ span a subspace of dimension $r + 1$. Hence, for any subspace $S \neq S^*$ with $\dim(S) \leq r$, we have $\#\{\tilde{x}_j \,|\, \tilde{x}_j \in S;\, j \in [n]\} \leq r$. □

2 Implementation Details in Experiments

Robust-PCA methods in Section 4.1. In Self-Repr and CoP, $\hat{S}$ is taken to be the subspace spanned by the top r $\tilde{x}_j$'s with largest inlier scores. We use the Iteratively-Reweighed-Least-Squares method proposed by [9, 4] for solving the DPCP problem. The output subspace $\hat{S}$ of OP is obtained as the rth principal component subspace of the decomposed low-rank matrix. For Self-Expr we use $\lambda = 0.95$, $\alpha = 10$ and $T = 1000$, see section 5 in [12]. For DPCP we use $T_{\max} = 1000$, $\epsilon = 10^{-9}$ and $\delta = 10^{-15}$, see Algorithm 2 in [10]. Finally, OP uses $\lambda = 0.5$ in Algorithm 1 of [11], while the parameter τ of the augmented Lagrangian is 1.

Unlabeled sensing methods. For AIEM, we use the customized Gröbner basis solvers of [8], developed for $r \leq 4$, which solve the polynomial system in milliseconds, and the maximum number iterations in the alternating minimization procedure is $T_{\max} = 1000$. For $r = 5$, the design of such solvers is an open problem[1], thus we use the generic solver Bertini ([1]), which runs within a few

[1] The fast solver generator of [3] is an improved version of the one used by [8] for $r = 3, 4$. However, we found that for $r = 5$ it suffers from numerical stability issues.

seconds. For $r \geq 6$ though, AIEM remains as of now practically intractable. For CCV-Min the precision is 0.001, $T_{\text{max}} = 50$, and the maximum depth is 12 for $r = 3$ and 14 for $r = 4, 5$. For ℓ_1-RR we use $\lambda = 0.01\sqrt{\log(n)/n}$ in (13) of [5].

Face Images. We compute $\hat{S}$ as follows. With $\tilde{X} = U\Sigma V^\top$ the thin SVD of $\tilde{X}$, where $U \in \mathbb{R}^{32256\times 64}$, DPCP fits a 4-dimensional subspace $\bar{S}$ to the columns of $\bar{X} = U^\top \tilde{X}$, a process which takes about a tenth of a second. Then $\bar{S}$ is embedded back into $\mathbb{R}^{32256}$ via the map $U : \mathbb{R}^{64} \to \mathbb{R}^{32256}$ to yield $\hat{S}$. To compute $\hat{X}$ from $\hat{S}$ and $\tilde{X}$ we use the custom algebraic solver of AIEM as well as ℓ_1-RR, PL, Algorithm 2, with a proximal subgradient implementation of ℓ_1-RR using the toolbox of [2].

3 Additional Figures

First, we provide two more rows of $\alpha = 0.6, 0.2$ of Figure 1 in the paper, see Figure 1 in this supplemental manuscript.

Second, we provide one more evidence that the estimated subspace $\hat{S}$ is satisfactory, which is computed by DPCP with the same settings in Figure 3 of the paper. Figure 2 in this supplemental manuscript additionally shows the estimation error in Figures 2e-2g when S^* is used instead of $\hat{S}$. Evidently, the performance is nearly identical regardless of whether $\hat{S}$ or S^* is used. This is justified by Figure 2a, which shows that the maximal principal angle between $\hat{S}$ and S^* always stays below 2°.

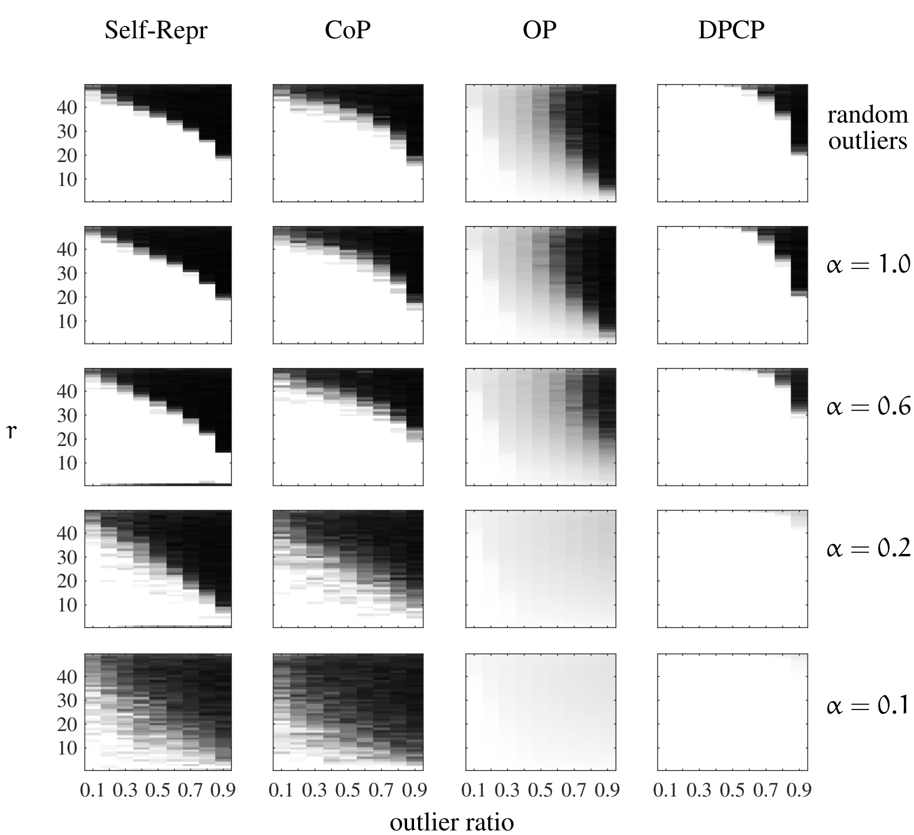

Figure 1: Same setting of Figure 1 in the paper but with an additional row $\alpha = 0.6, 0.2$.

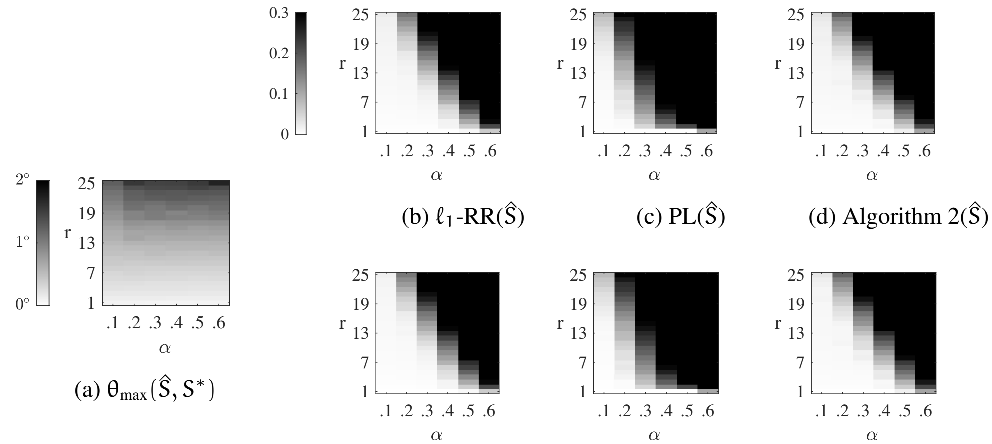

(a) $\theta_{\max}(\hat{S}, S^*)$ (b) ℓ_1-RR($\hat{S}$) (c) PL($\hat{S}$) (d) Algorithm 2($\hat{S}$) (e) ℓ_1-RR(S^*) (f) PL(S^*) (g) Algorithm 2(S^*)

Figure 2: Same settings of Figure 3 in the paper but with 2e-2g: Relative estimation error for the same setting with $\hat{X}$ computed from $\tilde{X}$ and S^*.

References

[1] D. J. Bates, J. D. Hauenstein, S. J. Andrew, and C. W. Wampler. Bertini: Software for numerical algebraic geometry. Available at bertini.nd.edu.

[2] A. Beck and N. Guttmann-Beck. Fom – a matlab toolbox of first-order methods for solving convex optimization problems. *Optimization Methods and Software*, 34(1):172–193, 2019.

[3] V. Larsson, K. Astrom, and M. Oskarsson. Efficient solvers for minimal problems by syzygy-based reduction. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2383–2392, 2017.

[4] G. Lerman and T. Maunu. Fast, robust and non-convex subspace recovery. *Information and Inference: A Journal of the IMA*, 7(2):277–336, 2018.

[5] M. Slawski and E. Ben-David. Linear regression with sparsely permuted data. *Electronic Journal of Statistics*, 13(1):1–36, 2019.

[6] X. Song, H. Choi, and Y. Shi. Permuted linear model for header-free communication via symmetric polynomials. In *IEEE International Symposium on Information Theory*, pages 661–665, 2018.

[7] M. C. Tsakiris. An exposition to the finiteness of fibers in matrix completion via Plücker coordinates. Technical report, arXiv:2004.12430v4 [cs.LG], 2020.

[8] M. C. Tsakiris, L. Peng, A. Conca, L. Kneip, Y. Shi, and H. Choi. An algebraic-geometric approach for linear regression without correspondences. *IEEE Transactions on Information Theory*, 66(8):5130–5144, 2020.

[9] M. C. Tsakiris and R. Vidal. Hyperplane clustering via dual principal component pursuit. In *International Conference on Machine Learning*, pages 3472–3481, 2017.

[10] M. C. Tsakiris and R. Vidal. Dual principal component pursuit. *The Journal of Machine Learning Research*, 19(1):684–732, 2018.

[11] H. Xu, C. Caramanis, and S. Sanghavi. Robust pca via outlier pursuit. *IEEE Transactions on Information Theory*, 58(5):3047–3064, 2012.

[12] C. You, D. P. Robinson, and R. Vidal. Provable self-representation based outlier detection in a union of subspaces. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 3395–3404, 2017.