PHAST: MODEL-FREE PHASELESS SUBSPACE TRACKING

Seyedehsara Nayer and Namrata Vaswani

Iowa State University, Ames, IA, USA.

ABSTRACT

Phaseless subspace tracking is the problem of recovering a time sequence of discrete signals from magnitude-only measurements of their linear projections, when the true signal lies in a low-dimensional subspace that can change with time. A typical assumption used in a lot of work is that the subspace changes gradually over time. We define this as (i) the maximum principal angle between the old and new subspaces is not too large (less than 90 degrees) or the number directions that changes is few or both; and (ii) the delay between subspace change times is large enough. This paper presents a novel algorithm, that we call PhaST, for model-free, minibatch and fast Phaseless Subspace Tracking. We show via experiments that PhaST is significantly faster, and significantly more memory-efficient, than an existing algorithm for lowrank phase retrieval (which can be interpreted as a batch version of phaseless subspace tracking that does not assume anything about subspace changes). When fewer measurements are available, it also has significantly better recovery performance than both LRPR and single signal phase retrieval methods when its structural assumptions are valid.

Index Terms- Phase retrieval, PCA, low-rank

1. INTRODUCTION

The Phase Retrieval (PR) problem occurs in many applications such as ptychography, crystallography, astronomy. The original PR problem involves recovering an n length signal ℓ from the magnitudes of its Discrete Fourier Transform (DFT) coefficients. Generalized PR (see [1] and [2]) replaces DFT by inner products with any set of measurement vectors, d_i . Thus, the goal is to recover ℓ from $|\langle d_i, \ell \rangle|, i = 1, 2, \dots, m$. It is clear that, without extra assumptions, PR will require $m \geq n$. In recent works, structural assumptions such as sparsity (see [3, 4, 5]) or low-rank (see [2]) have been incorporated into the PR problem in order to reduce the number of measurements m required for exact or accurate recovery. Low-rank has been used in two ways. One is to assume that a single signal re-arranged as a matrix (or a single image) is itself approximately or exactly low-rank. The goal is to recover this low-rank "signal" from its phaseless linear projections [6, 7]. The second is to assume that a time sequence of signals (or vectorized images) together form a matrix that is well modeled as being low-rank. Each signal/image is one column of this matrix. The measurements are phaseless linear projections of each signal or image (each column of the matrix) [2]. This problem has been referred to as "Low-Rank Phase Retrieval (LRPR)" in [2] where it was first studied.

Another way to interpret the LRPR problem is as follows [2]. We need to recover a time sequence of signals ℓ_t , t = 1, 2, ..., d, that are assumed to be generated from an unknown low dimensional subspace. Mathematically, $\ell_t = \mathbf{P}\mathbf{a}_t$, where \mathbf{P} is an $n \times r$ basis matrix (tall matrix with mutually orthonormal columns) with $r \ll n$, and \mathbf{a}_t is an $r \times 1$ coefficients' vector. For each $\ell_t, t = 1, 2, ..., d$, we have access to m phaseless measurements, $y_{i,t} = |\langle \mathbf{d}_{i,t}, \ell_t \rangle|, i =$ 1, 2, ...m, t = 1, 2, ...d. The goal is to either just recover the subspace, span(\mathbf{P}), or to recover both span(\mathbf{P}) and the coefficients and hence recover the signals ℓ_t 's (equivalently, recover the low-rank matrix $\mathbf{L} := [\ell_1, \ell_2, ..., \ell_d]$).

This work studies the Phaseless Subspace Tracking problem. This can be simply understood as the dynamic or timevarying subspace extension of LRPR. Thus, instead of the subspace span(P) being fixed, we assume that it can change with time, albeit slowly. Time varying subspaces is a better model (than just fixed subspace) for long data sequences, e.g., long image sequences or videos, because if one tries to use a single low-dimensional subspace to represent the entire data sequence, the required subspace dimension may end up being quite large. This can be problematic because it means that the resulting data matrix may not be sufficiently low-rank.

The most general model for time-varying subspaces is to allow the subspace to change at each time. However such a model involves too many unknowns. An r dimensional subspace in n-dimensional ambient space is fully specified by nrparameters. But the signal ℓ_t is an $n \times 1$ vector (has only n unknowns). Thus, allowing the subspace to change at each time will result in an increase in the number of unknowns per time instant from n to nr (rather than a decrease which is the purpose of incorporating structure into the PR problem). A less general model, but one that prevents an increase in the number of unknowns, is to assume that the data subspace is piecewise constant with time. This model has been extensively used in the robust subspace tracking literature [8, 9, 10]. We will use this model here as well. We explain it next.

Notation. $\|\cdot\|$ denotes the l_2 norm of a vector or the induced l_2 norm of a matrix. For other l_p norms, we use $\|\cdot\|_p$.

A matrix with mutually orthonormal columns is referred to as a "basis matrix". For basis matrix P we let P^{\perp} to be a basis for the subspace orthogonal to subspace spanned by \boldsymbol{P} . For basis matrices $\hat{\boldsymbol{P}}, \boldsymbol{P}$, the subspace error (SE) between their respective column spans is quantified by $\operatorname{SE}(\hat{\boldsymbol{P}}, \boldsymbol{P}) \coloneqq \|\hat{\boldsymbol{P}}^{\perp}\boldsymbol{P}\|$. This measures the sine of the principal angle between the subspaces.

For any two integers i_1, i_2 , the interval $[i_1 : i_2]$ denotes the set of integer values $\{i_1, i_1 + 1, \ldots, i_2\}$ and interval $[i_1 : i_2)$ denotes the set $\{i_1, i_1+1, \ldots, i_2-1\}$. Also, for integer values of t and α we let $[t; \alpha] = \{t - \alpha + 1, \ldots, t\}$.

Problem Setting. For the identifiability reason explained above, we assume a piecewise constant (with time) model on subspace change. Denote the subspace change times by t_j for $j = 1, 2, \ldots, J$ and let $t_0 = 0$. We assume that $\ell_t = P_j a_t$ for all $t = t_j, t_j + 1, t_j + 2, \ldots, t_{j+1} - 1$. The goal is to recover the ℓ_t 's and their subspaces from m phaseless measurements at each time, i.e., from $y_{i,t} := |\langle d_{i,t}, \ell_t \rangle|, i = 1, 2, \ldots, m$ for each $t = 1, 2, \ldots, d$. Here, $d_{i,t}$ is measurement vector and we let $D_t = [d_{1,t}, d_{2,t}, \ldots, d_{m,t}]$ to be measurement wat it at time t. Similarly, we define phaseless measurement vector $y_t = [y_{1,t}, y_{2,t}, \ldots, y_{m,t}]'$ at time t. Then, y_t can be presented as $y_t = |D'_t \ell_t|$.

We assume that the subspace changes are gradual or "slow". This means the following:

- SE(P_{j-1}, P_j) ≤ Δ ≪ 1, or only a few subspace directions change at each t_j, or both (we need at least one of these two assumptions to hold);
- and the delay between subspace change times, q_j := min_j(t_{j+1} − t_j), is large enough. Of course this delay needs to be at least r to even compute the subspace even when ℓ_t's are directly available.

Under this model, the question is when can one solve this problem either using a smaller m per signal than what is needed for LRPR? Moreover, can we get a faster and more memory-efficient mini-batch solution instead of the LRPR algorithm which is a batch one (and hence slow and memory-intensive)?

The LRPR work [2] has demonstrated that just exploiting the low-rank assumption enables a reduction in the required m compared to simple PR done for each signal ℓ_t individually. However, the time complexity of LRPR is roughly rtimes higher than that of simple PR; while its memory requirement is equal to the size of the entire batch of data. This can be prohibitive for large-sized image sequences. This work aims to achieve a significant speed up while also significantly reducing the memory complexity of the resulting minibatch approach. Moreover, we also show empirically that our proposed algorithm, which we call "fast Phaseless Subspace Tracking (PhaST)", needs a smaller m or accurate recovery when compared with LRPR (and of course with simple PR, e.g., the truncated Wirtinger Flow (TWF) algorithm of [1]. Algorithm 1 An overview of the proposed PhaST algorithm; for simplicity we present a version of the algorithm that assumes t_j known; we explain the approach to detect t_j in the text. Our actual code used for all our experiments does not assume t_j known.

Input: $\{y_t, D_t\}, K, \alpha, T_P = 2, T_a$. 1: Initialize: $j \leftarrow 0, k \leftarrow 1$. 2: $\hat{P}_0 \leftarrow \text{LRPR}$ algorithm with $\{y_t, D_t\}_{t=0:t_{\text{train}}}$ 3: for $t > t_{\text{train}}$ do 4: if $t = t_j + \alpha$ then $\hat{P}_{i1}^{new} \leftarrow \text{top } r \text{ eigenvectors of } \tilde{Y}_{P}(\hat{P}_{i-1}) \text{ with}$ 5:
$$\begin{split} \tilde{\boldsymbol{Y}}_{P}(\hat{\boldsymbol{P}}) & \text{as given in (3).} \\ \hat{\boldsymbol{P}}_{j,1}^{init} \leftarrow [\hat{\boldsymbol{P}}_{j-1}, \hat{\boldsymbol{P}}_{j,1}^{new}] \\ & \text{for } \tau \in [t; \alpha] \text{ do} \\ & \hat{\boldsymbol{a}}_{\tau}^{init} \leftarrow \text{top eigenvector of } \boldsymbol{Y}_{\boldsymbol{a}}(\hat{\boldsymbol{P}}_{j,1}^{init}) \text{ with} \end{split}$$
6: 7: 8. $Y_a(\hat{P})$ defined in (2). for $v \in [1:T_a]$ do 9: $\hat{\boldsymbol{C}}_{\tau} \stackrel{i}{\leftarrow} \text{Phase}(\boldsymbol{D}_{\tau} \hat{\boldsymbol{P}}_{j,1}^{init} \hat{\boldsymbol{a}}_{\tau}) \\ \hat{\boldsymbol{a}}_{\tau}^{init} \leftarrow \arg\min_{\tilde{\boldsymbol{a}}} \|\hat{\boldsymbol{C}}_{\tau} \boldsymbol{y}_{\tau} - \boldsymbol{D}_{\tau} \hat{\boldsymbol{P}}_{j,1}^{init} \tilde{\boldsymbol{a}}\|^2$ 10: 11: end for 12: end for $\hat{L}^{init} = \hat{P}_{j,1}^{init} \hat{A}_{t;\alpha}^{init}$ $\hat{P}_{j,1}^{(0)} \leftarrow \text{top } r \text{ singular vectors of } \hat{L}^{init}$ else if $t = t_j + k\alpha, k = 2, \dots$ then 13: 14: 15: 16: $\hat{oldsymbol{P}}_{j,k}^{(0)} \leftarrow \hat{oldsymbol{P}}_{j,k-1}.$ end if 17: 18: if $t = t_j + k\alpha$, $k = 1, \ldots$ then 19: if k = 1 set $T_p \leftarrow 3$ else if $k = 2, \ldots, K$ set 20: $T_p \leftarrow 2$ for $v = 1, ..., T_p$ do 21: for $\tau \in [t; \alpha]$ do 22: $\hat{a}_{\tau} \leftarrow$ top eigenvector of $Y_{a}(\hat{P}_{i,k}^{(v-1)})$ 23: end for 24: for $u = 1, ..., T_a$ do 25: for $\tau \in [t; \alpha]$ do 26: $egin{array}{lll} \hat{m{C}}_{ au} &\leftarrow ext{Phase} \left(m{D}_{ au}' \hat{m{P}}_{j,k}^{v-1} \hat{m{a}}_{ au}
ight) \ \hat{m{a}}_{ au} &\leftarrow ext{arg min}_{m{a}} \| \hat{m{C}}_{ au} m{y}_{ au} \end{array}$ 27: 28: $oldsymbol{D}_{ au} \hat{oldsymbol{P}}_{j,k}^{v-1} oldsymbol{a} \|^2$ end for 29: $\hat{\boldsymbol{P}}_{j,k}^{(v)} \leftarrow rg \min_{\boldsymbol{P}} \sum_{\tau \in [t;\alpha]} \sum_{i=1}^{m} \|\hat{\boldsymbol{C}}_{\tau} \boldsymbol{y}_{\tau} - \boldsymbol{D}_{\tau}' \boldsymbol{P} \hat{\boldsymbol{a}}_{\tau} \|^2$ 30: 31: end for 32: $\hat{P}_{j,k} \leftarrow \hat{P}_{j,k}^{(T_p)}$ if $t = t_{j+1}$ then 33: 34: $\hat{\boldsymbol{P}}_{j} \longleftarrow \hat{\boldsymbol{P}}_{j,k}^{(T_p)}, k \leftarrow 0, j \leftarrow j+1$ end if 35: 36: $k \leftarrow k + 1$ 37: 38: end if Output: $\hat{L}_{j,k} := \hat{P}_{j,k}^{(1)} \hat{A}_{j,k}$ with $\hat{A}_{j,k} := [\hat{a}_{\tau}]_{[t;\alpha]}$ 39: 40: end for

Our Contribution. This work takes the first steps towards a mini-batch and model-free phaseless subspace tracking algorithm that we call "fast Phaseless Subspace Tracking (PhaST)". We show via experiments that PhaST is significantly faster, and significantly more memory-efficient than LRPR, while still being more accurate than LRPR when m is small. Of course this holds only when the slow changing subspace assumption holds. Under this assumption, it also needs a smaller m than simple (single signal) PR methods. These methods need $m \ge n$, while both LRPR and PhaST can tolerate smaller m's because they exploit structural assumptions.

We studied a preliminary model-based version of the phaseless subspace tracking problem in [11]. That work assumed a very specific model on subspace change: it required that only one subspace direction change at each time while the rest of the subspace needed be fixed. Moreover, it needed a very accurate estimate of the previous subspace before it could detect or track changes; and it did not provide any experimental verification of a true tracking approach with multiple subspace changes. This work removes all of these requirements and demonstrates automatic detection and tracking of subspace changes.

2. FAST PHASELESS SUBSPACE TRACKING (PHAST)

Our proposed algorithm, which we term "PhaST", consists of two parts. The first part is to detect if there has been any change in the subspace. The second part is to update (obtain better and better estimates of) the changed subspace on-thefly (in a mini-batch fashion).

To understand the main idea of the update steps, first assume that t_j is known. We split the update step into recovering α -consecutive-column-sub-matrices of L at a time. Let $L_{t;\alpha} := [\ell_{t-\alpha+1}, \ell_{t-\alpha+2}, \dots, \ell_t]$ to be one α column interval in $[t_j : t_{j+1})$. Then $L_{t;\alpha} := PA_{t;\alpha}$ in this interval with $P = P_j$. For each such sub-matrix, we have access to m measurements per column, i.e., we have $y_{i,\tau} :=$ $|d_{i,\tau}'P_ja_{\tau}|, i \in [1:m], \tau \in [t;\alpha]$; and we also have access to a subspace estimate \hat{P} available from the previous interval. Recall that $[t;\alpha] := [(t-\alpha+1):t]$. Suppose $SE(\hat{P}, P) \leq b$ with b small. We would like to solve for \hat{P} , \tilde{a} that minimize

$$\sum_{\tau \in [t;\alpha]} \sum_{i \in [1:m]} |y_{i,\tau} - |\boldsymbol{d}_{i,\tau}' \tilde{\boldsymbol{P}} \tilde{\boldsymbol{a}}_{\tau}||^2,$$
(1)

subject to the constraint that $SE(\hat{P}, P) \leq b$. Alternatively when b is unknown, we could instead try to solve the Lagrangian (unconstrained) version of this problem. In either case, if we tried to develop an alternating minimization (altmin) solution that alternates over P, a_{τ} 's, and the measurements' phase matrices $C_{\tau} := \text{diag}(\text{phase}(D_{\tau}'Pa_{\tau}))$, we would not get a closed-form update rule for P. One trick that is often used in such situations is the following: an indirect way to impose a constraint of the form $distance(\mathbf{x}_0, \mathbf{x}) \leq b$ (small) when minimizing a cost function over \mathbf{x} is to start the alt-min algorithm with initialization \mathbf{x}_0 and then to run only one (or a few) iterations of alt-min for the unconstrained problem over \mathbf{x} . In our setting, this means we initialize the alt-min algorithm with \hat{P} , and run only one or two iterations of altmin to minimize (1) over P; of course for each value of P, we run many iterations of alt-min over a_{τ} 's and C_{τ} 's. This also needs init for a_{τ} 's. Given a \hat{P} , we initialize a_{τ} using the approach introduced for LRPR [2]: compute a_{τ} as the top eigenvector of

$$\tilde{Y}_{\boldsymbol{a}}(\hat{\boldsymbol{P}}) := \hat{\boldsymbol{P}}'\left(\frac{1}{m}\sum_{i=1}^{m}y_{i,t}^{2}\boldsymbol{d}_{i,t}\boldsymbol{d}_{i,t}'\right)\hat{\boldsymbol{P}}.$$
(2)

When b is not small, but the subspaces are still close in other metrics (for example, if a large part of the subspace does not change at all, c.f. the model of [11]), we can still leverage knowledge of \hat{P} as follows. We initialize the altmin solution not with \hat{P} , but with a spectral initialization step that uses \hat{P} to first find the "newly added subspace" (basis for $\hat{P}^{\perp'}P$), followed by using a simple trick to delete the removed directions. Let P_{add} denote a basis for $\hat{P}^{\perp'}P$. Using a simple modification of the idea developed in our earlier work [11], our spectral initialization step computes \hat{P}_{add} as the top r eigenvectors of the matrix

$$\tilde{\boldsymbol{Y}}_{P}(\hat{\boldsymbol{P}}) := (\boldsymbol{I} - \hat{\boldsymbol{P}}\hat{\boldsymbol{P}}') \left(\frac{1}{m\alpha} \sum_{\tau \in [t;\alpha]} \sum_{i \in [1:m]} y_{i,\tau} \boldsymbol{d}_{i,\tau} \boldsymbol{d}_{i,\tau}'\right) (\boldsymbol{I} - \hat{\boldsymbol{P}}\hat{\boldsymbol{P}}').$$
(3)

We then let $P^{init,temp} := [\hat{P}, \hat{P}_{add}]$. Notice here that because of the initialization approach, $\hat{P}^{init,temp}$ is a $n \times 2r$ matrix. Using this we initialize the coefficient vectors, a_{τ} , (which are now 2r length vectors) as explained above but with \hat{P} is replaced by $\hat{P}^{init,temp}$ followed by obtaining improved estimates via minimizing (1) by alt-min over a_{τ} and C_{τ} (phase matrix) with \hat{P} held constant at $\hat{P}^{init,temp}$. We then set $\hat{L}^{init,temp}_{t;\alpha} := \hat{P}^{init,temp} \hat{A}$. By compute its top rleft singular vectors, we get a good r-dimensional initialization of the subspace. We denote this by \hat{P}^{init} .

The complete algorithm is summarized in Algorithm 1. It splits the matrix L into α -consecutive column intervals, followed by recursively getting better and better estimates of the subspace $P := P_j$. When a subspace change is detected, the subspace change is assumed to be large (*b* is not small enough) and hence at this time the above initialization step is needed. After at least one estimate of the new subspace is computed, it assumes *b* is small enough and P can be initialized just with the previous final subspace estimate.

3. EXPERIMENTS

This section provides two distinct set of experiments. To generate the data we adopt a similar approach as in [11]



Fig. 1: Subspace error of P_j vs. time for n = 500, m = 400, r = 10, $r_{add} = 10$, and $\theta = 15$. Here, J = 3, $\alpha = 300$, and $t_0 = 0$, $t_1 = 3300$, $t_3 = 6300$. LRPR- α shows subspace error when it is fed with one mini-batch at each time. LRPR-whole shows error when data of all time sequences is fed to LRPR and the obtained error is repeated for the time to simplify the representation in this case. Average time taken per column and the memory complexity is reported in the legend. Memory needed by PhaST and LRPR- α is $n\alpha = 300n$. TWF needs smaller memory of size n and is faster.

for a general case with more than one direction is changing. Let $P_{j-1,chg}$ denote directions from span (P_{j-1}) that changes at t_j and let $P_{j,chd}$ denote its changed version, then span (P_{j-1}) can be written as span $([P_{j-1,fix}, P_{j-1,chg}])$ and span (P_j) as span $([P_{j-1,fix}, P_{j,chd}])$, where $P_{j-1,fix}$ is an $n \times (r - r_{add})$ matrix corresponding to the fixed part of the subspace at t_j . This model allows only a part of the subspace to change. In our experiments here, we generate $P_{j,chd}$ as $P_{j,chd} = P_{j,add} \sin \theta_j - P_{j-1,chg} \cos \theta_j$ with a single angle for all changed columns (only for simplicity). Here $P_{j,add}$ contains a set of $r_{add} \leq r$ directions from span (P_{j-1}^{\perp}) . Beginning with $t_0 = 0$, the subspace changes at $t_1 = 3300$ and $t_2 = 6300$. All experiments are averaged over 50 Monte-Carlo repeats.

To illustrate the performance of proposed method, we consider two different scenarios. The first scenario models situations where the change is very small and aims at evaluating robustness of PhaST for cases with small values of θ . The second scenario is designed to show performance of the algorithm with larger values of θ . Both experiments are done for the setting with n = 500, m = 400, r = 10. PhaST uses K = 3, $\alpha = 300$, $T_P = 2$, $T_a = 10$. Detection is done with a similar approach as in [11]. The algorithm assumes that the subspace remains unchanged at least for three α time sequences. After three α time intervals of the last detected change, detection phase of the algorithm starts. Threshold constant for detection is set to 1.75 for both experiments. Figure 1 illustrates simulation results for the first scenario



Fig. 2: Subspace error of P_j vs. time for n = 500, m = 400, r = 10, $r_{add} = 4$, and $\theta = 75$. Here, J = 3, $\alpha = 300$, and $t_0 = 0$, $t_1 = 3300$, $t_3 = 6300$. Average time taken per column is reported

with $\theta = 15$ and $r_{add} = 10$. As figure suggests, PhaST is robust to small changes. Figure 2 belongs to the second scenario with $\theta = 75$, $r_{add} = 3$, and shows the ability of PhaST to detect larger changes.

To demonstrate the superiority of proposed algorithm, we compare it with two other existing state of the art PR methods, namely, LRPR and TWF. Comparison with LRPR can be done in two different ways. First is to feed LRPR with data of all time sequences. Hence $q = \sum_j q_j$ for LRPR. Although more measurements are available in this case, LRPR is unable to recover the subspace.

Another way of evaluating LRPR is to feed it with data of one mini-batch at each time. Since LRPR does not take advantage of the previous subspace estimates, it shows almost the same performance for different such batches. On the other hand TWF is also not taking advantage of the previous knowledge and hence its performance is almost the same for every column. Therefore since these approaches do not work for all time sequences as it can be seen in Fig. 1, we just report the average error of first 5 mini-batches for 50 independent trials in Fig. 2. Number of loop iterations for LRPR and TWF are 20. Number of power method iterations used in TWF initialization is 50 and parameter settings are the same as suggested in [1]. Simulation results confirm that PhaST needs less number of measurements to track the subspace than what is needed by TWF and LRPR. Besides it makes shorter delay than LRPR.

4. CONCLUSION

This paper proposed a novel algorithm for model-free, minibatch and fast Phaseless Subspace Tracking and provided first simulation results for complete PST algorithm. Experiments show that PhaST is memory efficient and needs less sample complexity in comparison with existing PR methods. Providing convergence analysis and real world applications will be part of the future work.

5. REFERENCES

- Y. Chen and E. Candes, "Solving random quadratic systems of equations is nearly as easy as solving linear systems," in *Adv. Neural Info. Proc. Sys. (NIPS)*, 2015, pp. 739–747.
- [2] N. Vaswani, S. Nayer, and Y. C. Eldar, "Low rank phase retrieval," *IEEE Trans. Sig. Proc.*, August 2017.
- [3] K. Jaganathan, S. Oymak, and B. Hassibi, "Recovery of sparse 1-d signals from the magnitudes of their fourier transform," in *IEEE Intl. Symp. on Information Theory* (*ISIT*). IEEE, 2012, pp. 1473–1477.
- [4] Y. Shechtman, A. Beck, and Y. C. Eldar, "Gespar: Efficient phase retrieval of sparse signals," *IEEE Trans. Sig. Proc.*, vol. 62, no. 4, pp. 928–938, 2014.
- [5] A. Szameit, Y. Shechtman, E. Osherovich, E. Bullkich, P. Sidorenko, H. Dana, S. Steiner, E. B. Kley, S. Gazit, T. Cohen-Hyams, S. Shoham, M. Zibulevsky, I. Yavneh, Y. C. Eldar, O. Cohen, and M. Segev, "Sparsity-based single-shot subwavelength coherent diffractive imaging," *Nature Materials*, vol. 11, pp. 455–9, Apr. 2012.
- [6] S. Tu, R. Boczar, M. Soltanolkotabi, and B. Recht, "Low-rank solutions of linear matrix equations via procrustes flow," *arXiv preprint arXiv:1507.03566*, 2015.
- [7] Q. Zheng and J. Lafferty, "A convergent gradient descent algorithm for rank minimization and semidefinite programming from random linear measurements," in *Adv. Neural Info. Proc. Sys. (NIPS)*, 2015.
- [8] C. Qiu, N. Vaswani, B. Lois, and L. Hogben, "Recursive robust pca or recursive sparse recovery in large but structured noise," *IEEE Trans. Info. Th.*, pp. 5007–5039, August 2014.
- [9] P. Narayanamurthy and N. Vaswani, "Provable dynamic robust pca or robust subspace tracking," *IEEE Trans. Info. Theory*, to appear, 2018.
- [10] P. Narayanamurthy and N. Vaswani, "Nearly optimal robust subspace tracking," in *ICML 2018*, 2018, pp. 3701– 3709.
- [11] S. Nayer and N. Vaswani, "Phaseless subspace tracking," Accepted for publication in IEEE Global Conference on Signal and Information Processing (Global-SIP), 2018.