

Online Identification and Tracking of Subspaces from Highly Incomplete Information

Laura Balzano[#], Robert Nowak[#] and Benjamin Recht[†]

[#] Department of Electrical and Computer Engineering,
University of Wisconsin-Madison

[†] Department of Computer Sciences, University of Wisconsin-Madison

June 20, 2010

Abstract

This work presents GROUSE (Grassmanian Rank-One Update Subspace Estimation), an efficient online algorithm for tracking subspaces from highly incomplete observations. GROUSE requires only basic linear algebraic manipulations at each iteration, and each subspace update can be performed in linear time in the dimension of the subspace. The algorithm is derived by analyzing incremental gradient descent on the Grassmannian manifold of subspaces. With a slight modification, GROUSE can also be used as an online incremental algorithm for the matrix completion problem of imputing missing entries of a low-rank matrix. GROUSE performs exceptionally well in practice both in tracking subspaces and as an online algorithm for matrix completion.

1 Introduction

The evolution of high-dimensional dynamical systems can often be well summarized or approximated in low-dimensional subspaces. Certain patterns of computer network traffic including origin-destination flows can be well represented by a subspace model [13]. Environmental monitoring of soil and crop conditions [10], water contamination [15], and seismological activity [18] have all been demonstrated to be efficiently summarized by very low-dimensional subspace representations.

The subspace methods employed in the aforementioned references are based upon first collecting full-dimensional data from the systems and then using approximation techniques to identify an accurate low-dimensional representation. However, it is often difficult or even infeasible to acquire and process full-dimensional measurements. For example, collecting network traffic measurements at a very large number of points and fine time-scales is impractical. An alternative is to randomly subsample the full-dimensional data. If the complete full-dimensional data is well-approximated by a lower dimensional subspace, and hence is in some sense redundant, then it is conceivable that the subsampled data may provide sufficient information for the recovery of that subspace. This is the central intuition of our proposed on-line algorithm for identifying and tracking low-dimensional subspaces from highly incomplete (i.e., undersampled) data. Such an algorithm could enable rapid detection of traffic spikes or intrusions in computer networks [13] or could provide large efficiency gains in managing energy consumption in a large office building [8].

In this paper, we present GROUSE (Grassmanian Rank-One Update Subspace Estimation), a subspace identification and tracking algorithm that builds high quality estimates from very sparsely sampled vectors. GROUSE implements an incremental gradient procedure with computational complexity linear in dimensions of the problem, and is scalable to very high-dimensional applications. An additional feature of GROUSE is that it can be immediately adapted to an ‘online’ version of the matrix completion problem, where one aims to recover a low-rank matrix from a small, streaming random subsets of its entries. GROUSE is not only remarkably efficient for online matrix completion, but additionally enables incremental updates as columns are added or entries are incremented over time. These features are particularly attractive for maintaining databases of user preferences and collaborative filtering.

2 Problem Set-up

We aim to track an d -dimensional subspace of \mathbb{R}^n that evolves over time, denoted by $S[t]$. At every time t , we observe a vector $v_t \in S[t]$ at locations $\Omega_t \subset \{1, \dots, n\}$. Let \mathcal{P}_{Ω_t} denote the $|\Omega_t| \times n$ matrix that selects the coordinate axes of \mathbb{R}^n indexed by Ω_t . That is, we observe $\mathcal{P}_{\Omega_t} v_t$, $t = 1, 2, \dots$. We will measure the error of our subspace using the natural cost function: the squared Euclidean distance from our current subspace estimate $S_{\text{est}}[t]$ to the observed vector v_t *only on the coordinates revealed in the set in Ω_t* :

$$F(S_{\text{est}}[t], t) := \text{dist}(\mathcal{P}_{\Omega_t}(S_{\text{est}}[t]), \mathcal{P}_{\Omega_t}(v_t))^2 \quad (1)$$

We can compute $F(S; t)$ explicitly for any subspace S . Let U be any matrix whose columns span S . Let U_{Ω_t} denote the submatrix of U consisting of the rows indexed by Ω_t . For a vector $v \in \mathbb{R}^n$, let $v_{\Omega_t} = \mathcal{P}_{\Omega_t}(v)$ denote a vector in $\mathbb{R}^{|\Omega_t|}$ whose entries are indexed by Ω_t . Then we have

$$F(S; t) = \min_a \|U_{\Omega_t} a - v_{\Omega_t}\|^2 \quad (2)$$

We will use this definition in Section 3 to derive our algorithm. If the matrix $U_{\Omega_t}^T U_{\Omega_t}$ has full rank, then we must have that $w = (U_{\Omega_t}^T U_{\Omega_t})^{-1} U_{\Omega_t}^T v_{\Omega_t}$ achieves the minimum in (2). Thus,

$$F(S; t) = v_{\Omega_t}^T (I - U_{\Omega_t} (U_{\Omega_t}^T U_{\Omega_t})^{-1} U_{\Omega_t}^T) v_{\Omega_t}. \quad (3)$$

In the special case where the subspace is time-invariant, that is $S[t] = S_0$ for some fixed subspace S_0 , then it is natural to consider the average cost function

$$\bar{F}(S) := \sum_{t=1}^T \text{dist}(\mathcal{P}_{\Omega_t}(S), \mathcal{P}_{\Omega_t}(v_t))^2. \quad (4)$$

The average cost function will allow us to estimate the steady-state behavior of our algorithm. Indeed, in the static case, our algorithm will be guaranteed to converge to a stationary point of $\bar{F}(S)$.

2.1 Relation to Matrix Completion

In the scenario where the subspace does not evolve over time and we only observe vectors on a finite time horizon, then the cost function (4) is identical to the matrix completion optimization

problem studied in [11, 6]. To see the equivalence, let $\Omega = \{(k, t) : k \in \Omega_t, 1 \leq t \leq T\}$, and let $V = [v_1, \dots, v_T]$. Then

$$\bar{F}(S) = \sum_{t=1}^T \min_a \|U_{\Omega_t} a - v_{\Omega_t}\|^2 = \min_{A \in \mathbb{R}^{d \times T}} \sum_{(i,j) \in \Omega} (UA - V)_{ij}^2 \quad (5)$$

That is, the global optimization problem can be written as $\min_{U,A} \sum_{(i,j) \in \Omega} (UA - V)_{ij}^2$, which is precisely the starting point for the algorithms and analyses in [11, 6]. The authors in [11] use a gradient descent algorithm to jointly minimize both U and A while [6] minimizes this cost function by first solving for A and then taking a gradient step with respect to U . In the present work, we consider optimizing this cost function one column at a time. We show that by using our online algorithm, where each measurement v_t corresponds to a random column of the matrix V , we achieve state-of-the-art performance on matrix completion problems.

3 Stochastic Gradient Descent on the Grassmannian

The set of all subspaces of \mathbb{R}^n of dimension d is denoted $\mathfrak{G}(n, d)$ and is called the *Grassmannian*. The Grassmannian is a compact Riemannian manifold, and its geodesics can be explicitly computed [7]. An element $S \in \mathfrak{G}(n, d)$ can be represented by any $n \times d$ matrix U whose columns form an orthonormal basis for S . Our algorithm derives from an application of incremental gradient descent on the Grassmannian manifold. We first compute a gradient of the cost function F , and then follow this gradient along a short geodesic curve in the Grassmannian.

We follow the program developed in [7]. To compute the gradient of F on the Grassmannian manifold, we first need to compute the partial derivatives of F with respect to the components of U . For a generic subspace, the matrix $U_{\Omega_t}^T U_{\Omega_t}$ has full rank provided that $|\Omega_t| > r$, and hence the cost function (2) is differentiable almost everywhere. Let Δ_{Ω_t} be the $n \times n$ diagonal matrix which has 1 in the j^{th} diagonal entry if $j \in \Omega_t$ and has 0 otherwise. We can rewrite

$$F(S; t) = \min_a \|\Delta_{\Omega_t}(Ua - v_t)\|^2 \quad (6)$$

from which it follows that the derivative of F with respect to the elements of U is

$$\frac{dF}{dU} = -2(\Delta_{\Omega_t}(v_t - Uw))w^T \quad (7)$$

$$= -2rw^T \quad (8)$$

where $r := \Delta_{\Omega_t}(v_t - Uw)$ denotes the (zero padded) residual vector and w is the least-squares solution in (2).

Using Equation (2.70) in [7], we can calculate the gradient on the Grassmannian from this partial derivative

$$\nabla F = (I - UU^T) \frac{dF}{dU} = -2(I - UU^T)rw^T = -2rw^T. \quad (9)$$

The final equality follows because the residual vector r is orthogonal to all of the columns of U . This can be verified from the definitions of r and w .

A gradient step along the geodesic with tangent vector $-\nabla F$ is given by Equation (2.65) in [7], and is a function of the singular values and vectors of ∇F . It is trivial to compute the singular value decomposition of ∇F , as it is rank one. The sole non-zero singular value is $\sigma = 2\|r\|\|w\|$ and the corresponding left and right singular vectors are $\frac{r}{\|r\|}$ and $\frac{w}{\|w\|}$ respectively. Let x_2, \dots, x_r be an orthonormal set orthogonal to r and y_2, \dots, y_r be an orthonormal set orthogonal to w . Then

$$-2rw^T = \begin{bmatrix} -\frac{r}{\|r\|} & x_2 & \dots & x_r \end{bmatrix} \text{diag}(\sigma, 0, \dots, 0) \begin{bmatrix} \frac{w}{\|w\|} & y_2 & \dots & y_r \end{bmatrix}^T \quad (10)$$

forms an SVD for the gradient. Now using (2.65) from [7], we find that for $\eta > 0$, a step of length η in the direction ∇F is given by

$$U(\eta) = U + \frac{(\cos(\sigma\eta) - 1)}{\|w\|^2} Uww^T + \sin(\sigma\eta) \frac{r}{\|r\|} \frac{w^T}{\|w\|} \quad (11)$$

$$= U + \left(\sin(\sigma\eta) \frac{r}{\|r\|} + (\cos(\sigma\eta) - 1) \frac{p}{\|p\|} \right) \frac{w^T}{\|w\|} \quad (12)$$

where $p := Uw$, the predicted value of the projection of the vector v onto S .

This geodesic update rule is remarkable for a number of reasons. First of all, it consists only of a rank-one modification of the current subspace basis U . Second, the term $\frac{\sin(\sigma\eta)}{\|r\|\|w\|} = \frac{\sin(\sigma\eta)}{\sigma}$ is on the order of η when $\sigma\eta$ is small. That is, for small values of σ and η this expression looks like a normal step along the gradient direction $-2rw^T$ given by (8). From the Taylor series of the cosine, we see that the second term is approximately equal to $\sigma^2\eta^2 \frac{pw^T}{\|p\|\|w\|}$. That is, this term serves as a second order correction to keep the iterates on the Grassmannian. Surprisingly, this simple additive term maintains the orthogonality, obviating the need for orthogonalizing the columns of U after a gradient step. Below, we will also discuss how this iterate relates to more familiar iterative algorithms from linear algebra which use full information.

The GROUSE algorithm simply follows geodesics along the gradients of F with a prescribed set of step-sizes η . The full computation is summarized in Algorithm 1. Our derivations have shown that computing a gradient step only requires the solution of the least squares problem (2), the computation of p and r , and then a rank one update to the previous subspace.

Each step of GROUSE can be performed efficiently with standard linear algebra packages. Computing the weights in Step 2 of Algorithm 1 requires solving a least squares problem in $|\Omega_t|$ equations and d unknowns. Such a system is solvable in at most $O(|\Omega_t|d^2)$ flops in the worst case. Predicting the component of v that lies in the current subspace requires a matrix vector multiply that can be computed in $O(nd)$ flops. Computing the residual then only requires $O(|\Omega_t|)$ flops, as we will always have zeros in the entries indexed by the complement of Ω_t . Computing the norms of r and p can be done in $O(n)$ flops. The final subspace update consists of adding a rank one matrix to an nr matrix and can be computed in $O(nd)$ flops. Totaling all of these computation times gives an overall complexity estimate of $O(nd + |\Omega_t|d^2)$ flops per subspace update.

3.1 Step-sizes

In the static case where the subspace $S[t] = S_0$ for all t , we can guarantee that the algorithm converges to a stationary point of (4) as long as the stepsizes η_t satisfy

$$\lim_{k \rightarrow \infty} \eta_t = 0 \quad \text{and} \quad \sum_{k=1}^{\infty} \eta_t = \infty \quad (13)$$

That is, our algorithm computes a mixture of the current subspace estimate and a predicted orthogonal component. This mixture is determined both by the stepsize and the relative energy of v_t outside the current subspace.

4 Numerical Experiments

4.1 Subspace Identification and Tracking

Static Subspaces We first consider the problem of identifying a fixed subspace. In all of the following experiments, the full data dimension is $n = 700$, the rank of the underlying subspace is $d = 10$, and the sampling density is 0.17 unless otherwise noted. We generated a series of iid vectors v_t according to generative model:

$$v_t = U_{\text{true}}\alpha + \beta \tag{16}$$

where U_{true} is an $n \times d$ matrix whose d orthonormal columns span the subspace, α is a $d \times 1$ vector whose entries are realizations of iid $\mathcal{N}(0, 1)$ random variables, and β is an $n \times 1$ vector whose entries are iid $\mathcal{N}(0, \omega^2)$. Here $\mathcal{N}(0, \omega^2)$ denotes the Gaussian distribution with mean zero and variance $\omega^2 > 0$ which governs the SNR of our data.

We implemented GROUSE (see Algorithm 1 above) with a stepsize rule of $\eta_t := C/t$ for some constant $C > 0$. Figure 1(a) shows the steady state error of the tracked subspace with varying values of C and the noise variance. All the data points reflect the error performance at $t = 14000$. We see that GROUSE converges for C ranging over an order of magnitude, however with additive noise the smaller stepsizes yield smaller errors. When there is no noise, i.e., $\omega^2 = 0$, the error performance is near the level of machine precision and is flat for the whole range of converging stepsizes. In Figure 1(b) we show the number of input vectors after which the algorithm converges to an error of less than 10^{-6} . The results are consistent with Figure 1(a), demonstrating that smaller stepsizes in the suitable range take fewer vectors until convergence. We only ran the algorithm up to time $t = 14000$, so the data points for the smallest and the largest few stepsizes only reflect that the algorithm did not yet reach the desired error in the allotted time. Figures 1(c) and 1(d) repeat identical experiments with a constant stepsize policy, $\eta_t = C$. We again see a wide range of stepsizes for which GROUSE converges, though the region of stability is narrower in this case.

We note that the norm of the residual $\|r\|$ provides an excellent indicator for whether tracking is successful. As shown in Figure 2(a), the error to the true subspace is closely approximated by $\|r\|/\|v_t\|$. This confirms the theoretical analysis in [1] which proves that this residual norm is an accurate estimator of the true subspace error when the number of samples is appropriately large.

Subspace Change Detection As a first example of GROUSE’s ability to adapt to changes in the underlying subspace, we simulated a scenario where the underlying subspace abruptly changes at three points over the course of an experiment with 14000 observations. At each break, we selected a new subspace S uniformly at random and GROUSE was implemented with a constant stepsize. As is to be expected, the algorithm is able to re-estimate the new subspace in a time depending on the magnitude of the constant stepsize.

Rotating Subspace In this second synthetic experiment, the subspace evolves according to a random ordinary differential equation. Specifically, we sample a skew-symmetric matrix B with




Figure 1: (a) Performance sensitivity to noise for a diminishing stepsize policy $\eta_t = C/t$. Results are displayed for three values of the noise magnitude, ω . (b) The number of iterations required to achieve an error of 10^{-6} as a function of C . (c) and (d) are the same except that the stepsize policy is here $\eta_t = C$.

independent, normally distributed entries and set

$$\dot{U} = BU, \quad U[0] = U_0. \tag{17}$$

The resulting subspace at each iteration is thus $U[t] = \exp(\delta t B)$ where δ is a positive constant. The resulting subspace at each iteration is thus $U[t] = \exp(\delta t B)$ where δ is a positive constant. In Figure 3, we show the results of tracking the rotating subspace with $\delta = 10^{-5}$. To demonstrate the effectiveness of the tracking, we display the projection of four random vectors using both the true subspace (in blue) and our subspace estimate at that time instant (in red).

Tracking Chlorine Levels We also analyzed the performance of the GROUSE algorithm on simulated chlorine level monitoring in a pressurized water delivery system. The data were generated using EPANET ¹ software and were previously analyzed [15]. The input to EPANET is a network layout of water pipes and the output has many variables including the chemical levels, one of which is the chlorine level. The data we used is available from [15] ². This dataset has ambient dimension $n = 166$, and $T = 4610$ data vectors were collected, once every 5 minutes over 15 days. We tracked an $d = 6$ dimensional subspace and compare this with the best 6-dimensional SVD

¹<http://www.epa.gov/nrmrl/wswrd/dw/epanet.html>

²<http://www.cs.cmu.edu/afs/cs/project/spirit-1/www/>

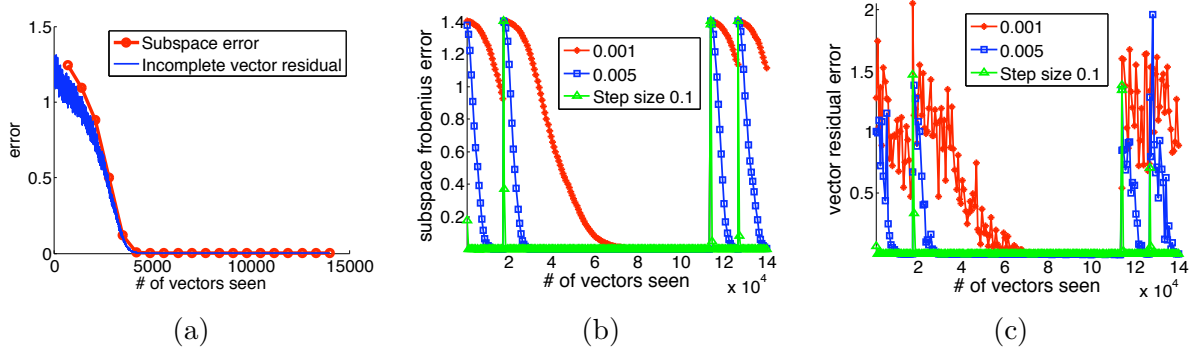


Figure 2: (a) Comparison of the distance to the true subspace and the norm of the residual in Step 4 of the GROUSE algorithm. The residual norm closely tracks the distance to the actual subspace. (b) Using constant stepsize to track sudden changes in subspace. We plot the transient behavior three constant stepsize policies. In (c), we again verify that the norm of the residual gives an accurate signature for anomaly detection and for tracking success.





Figure 3: Tracking a rotating subspace. Here we plot the norm of the projection of four random vectors over time. The blue curves denote the true values of these norms, and the red curves plot the GROUSE predictions. Note that except for very early transients, GROUSE fully tracks the subspace.

approximation of the entire complete dataset. The results are displayed in Figure 4. The table gives the results for various constant stepsizes and various fractions of sampled data. The smallest sampling fraction we used was 20%, and for that the best stepsize was $3e-2$; we also ran GROUSE on the full data, whose best stepsize was $5e-3$. As we can see, the performance error improves for the smaller stepsize of $5e-3$ as the sampling fraction increases; Also the performance error improves for the larger stepsize of $3e-2$ as the sampling fraction decreases. However for all intermediate sampling fractions there are intermediate stepsizes that perform near the best reconstruction error of about 0.12. The normalized error of the full data to the best 3-dimensional SVD approximation is 0.0704. Note that we only allow for one pass over the data set and yet attain, even with very sparse sampling, comparable accuracy to a full batch SVD which has access to all of the data.

Figure 4(a) and (b) show the original and the GROUSE reconstructions of five of the chlorine sensor outputs. We plot the last 500 of the 4310 samples, each reconstructed by the estimated subspace at that time instant.

4.2 Matrix Completion Problems

As described in Section 2.1, matrix completion can be thought of as a subspace identification problem where we aim to identify the column space of the unknown low-rank matrix. Once the column space is identified, we can project the incomplete columns onto that subspace in order to



$ \Omega_t /n$	η_t	reconstruction error
0.2	5e-3	0.2530
	3e-2	0.1244
0.4	5e-3	0.1797
	1e-2	0.1233
	3e-2	0.1432
0.7	5e-3	0.1289
	7e-3	0.1221
	3e-2	0.1684
1	5e-3	0.1253
	3e-2	0.2217

Figure 4: (a) Actual sensor readings. (b) Predicted sensor readings from tracked subspace. In these figures we are displaying the values at sensors 4, 17, 148, 158, and 159. The table lists errors in tracking the chlorine data set for varying sampling densities and stepsizes. The error between the data and the best SVD approximation is 0.0704.

complete the matrix. We have examined GROUSE in this context with excellent results. Our approach is to do descent on the column vectors in random order, and allow the algorithm to pass over those same incomplete columns a few times.

Our simulation set-up aimed to complete 700×700 dimensional matrices of rank 10 sampled with density 0.17. We generated the low-rank matrix by generating two factors Y_L and Y_R with i.i.d. Gaussian entries, and added normally distributed noise with variance ω^2 . The robustness to step-size and time to recovery are shown in Figure 5.

In Figure 6 we show a comparison of five matrix completion algorithms and GROUSE (labeled as Stochastic Gradient). Namely, we compare to the performance of OPT-SPACE [11], FPCA [14], SVT [4], SDPLR [3, 16], and NNLS [17]. We downloaded each of these MATLAB codes from the original developer’s websites when possible. We use the same random matrix model as in Figure 5. GROUSE is faster than all other algorithms, and achieves higher quality reconstructions on many instances. We subsequently compared against NNLS, the fastest batch method, on very large problems. Both GROUSE and NNLS achieved excellent reconstruction, but GROUSE was twice as fast.

5 Discussion and Future Work

The inherent simplicity and empirical power of the GROUSE algorithm merit further theoretical investigations to determine exactly when it provides consistent estimates. It is possible that an adaptation of the analysis of [11] to this online setting will yield such consistency bounds. The main difficulty lies in characterizing the trajectory of the GROUSE algorithm from a random starting subspace. Though we can guarantee that there is a basin of attraction around the global minimum, it is not yet clear how to characterize when GROUSE will end up in this basin.

We would also like to investigate how to adapt step-size in the GROUSE algorithm automatically to varying data. This is the only parameter required to run the GROUSE algorithm, and we have seen that there are substantial performance gains when this parameter is optimized. Using techniques from Least-Squares Estimation, it may be possible to automatically tune the step size based on our current error residuals.

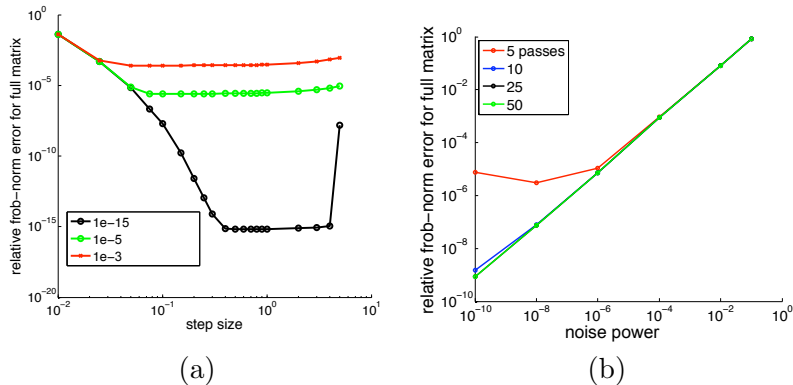



Figure 5: (a) Performance sensitivity to noise parameter ω and stepsize for matrix completion. Here we use a diminishing stepsize $\eta = C/t$. (b) Here we plot the time to get to a desired Frobenius norm error on the hidden matrix. We see that GROUSE converges to the noise floor after at most 10 passes over the columns of the matrix.

6 Acknowledgments

This work was partially supported by the AFOSR grant FA9550-09-1-0140. R. Nowak would also like to thank Trinity College and the Isaac Newton Institute at the University of Cambridge for support while this work was being completed.

References

- [1] L. Balzano, B. Recht, and R. Nowak. High-dimensional matched subspace detection when data are missing. Submitted for publication. Preprint available at <http://arxiv.org/abs/1002.0852>, 2010.
- [2] M. Brand. Fast low-rank modifications of the thin singular value decomposition. *Linear Algebra and its Applications*, 415(1):20–30, 2006.
- [3] S. Burer and R. D. C. Monteiro. Local minima and convergence in low-rank semidefinite programming. *Mathematical Programming*, 103(3):427–444, 2005.
- [4] J.-F. Cai, E. J. Candès, and Z. Shen. A singular value thresholding algorithm for matrix completion. *SIAM Journal on Optimization*, 20(4):1956–1982, 2008.
- [5] P. Comon and G. Golub. Tracking a few extreme singular values and vectors in signal processing. *Proceedings of the IEEE*, 78(8), August 1990.
- [6] W. Dai and O. Milenkovic. SET: An algorithm for consistent matrix completion. In *Proceedings of ICASSP*, 2010.
- [7] A. Edelman, T. A. Arias, and S. T. Smith. The geometry of algorithms with orthogonality constraints. *SIAM Journal on Matrix Analysis and Applications*, 20(2):303–353, 1998.
- [8] N. Gershenfeld, S. Samouhos, and B. Nordman. Intelligent infrastructure for energy efficiency. *Science*, 327(5969):1086, 2010.
- [9] M. Gu and S. Eisenstat. A stable and efficient algorithm for the rank-one modification of the symmetric eigenproblem. *SIAM journal on Matrix Analysis and Applications*, 15:1266–1276, 1994.
- [10] J. Gupchup, R. Burns, A. Terzis, and A. Szalay. Model-based event detection in wireless sensor networks. In *Proceedings of the Workshop on Data Sharing and Interoperability (DSI)*, 2007.



Problem parameters				GROUSE			NNLS	
n_r	n_c	r	dens	rel.err.	passes	time	rel. err.	time (s)
5000	20000	5	0.006	1.10e-4	2	14.8	4.12e-04	66.9
5000	20000	10	0.012	1.5e-3	2	21.1	1.79e-4	81.2
6000	18000	5	0.006	1.44e-5	3	21.3	3.17e-4	66.8
6000	18000	10	0.011	8.24e-5	3	31.1	2.58e-4	68.8
7500	15000	5	0.005	5.71e-4	4	36.0	3.09e-4	62.6
7500	15000	10	0.013	1.41e-5	4	38.3	1.67e-4	68.0

Figure 6: The figure compares five matrix completion algorithms against GROUSE. The table gives a comparison of GROUSE and NNLS on large random matrix completion problems.

- [11] R. H. Keshavan, A. Montanari, and S. Oh. Matrix completion from a few entries. 2009. Preprint available at <http://arxiv.org/abs/0901.3150>.
- [12] H. J. Kushner and G. G. Yin. *Stochastic Approximation and Recursive Algorithms and Applications*. Springer-Verlag, New York, second edition, 2003.
- [13] A. Lakhina, M. Crovella, and C. Diot. Diagnosing network-wide traffic anomalies. In *Proceedings of SIGCOMM*, 2004.
- [14] S. Ma, D. Goldfarb, and L. Chen. Fixed point and Bregman iterative methods for matrix rank minimization. Preprint available at http://www.optimization-online.org/DB_HTML/2008/11/2151.html, 2008.
- [15] S. Papadimitriou, J. Sun, and C. Faloutsos. Streaming pattern discovery in multiple time-series. In *Proceedings of VLDB Conference*, 2005.
- [16] B. Recht, M. Fazel, and P. Parrilo. Guaranteed minimum rank solutions of matrix equations via nuclear norm minimization. *SIAM Review*, 2007. To appear. Preprint Available at <http://pages.cs.wisc.edu/~brecht/publications.html>.
- [17] K.-C. Toh and S. Yun. An accelerated proximal gradient algorithm for nuclear norm regularized least squares problems. Preprint available at http://www.optimization-online.org/DB_HTML/2009/03/2268.html, 2009.
- [18] G. S. Wagner and T. J. Owens. Signal detection using multi-channel seismic data. *Bulletin of the Seismological Society of America*, 86(1A):221–231, February 1996.