Fast incremental method for matrix completion: an application to trajectory correction

Ricardo S. Cabral
ISR/Instituto Superior Tcnico
rscabral@cmu.edu

Abstract

We address the problem of recovering a matrix of tracked features, based on partial observations of their trajectories. Besides partial observability, we assume the existence of gross, but sparse, noise on the known entries. Recently, a spur of work in the optimization community has spun optimal methods for matrix completion when this matrix is known to be low rank by minimizing the nuclear norm. Albeit exhibiting several optimality properties, there are no incremental versions available.

This problem has obvious applications in real-time tracking and structure from motion, where the observations are usually plagued by self-occlusion and outliers in the tracker process.

In this paper, we build upon the Nuclear Norm Robust PCA method and SPectrally Optimal Completion to propose a fast and incremental algorithm which is able to cope with outliers.

We present experiments showing the competitive speed of our method and show that we obtain this results while maintaining performance comparable to the optimal approaches.

1. Introduction

We focus on the problem of recovering trajectories of a rigid object along several pictures of it [**?**]. Inherent to this problem are the difficulties created by self-occlusion and outliers introduced by the feature tracking process [**?**]. We deal with both of this issues by formulating an optimization problem which seeks to minimize the rank of the resulting measurement matrix, which we know beforehand to live on the set of rank-4 matrices.

This problem has already been explored by Aguiar *et al*. [2, 3] in the same context. Their method, SPectrally Optimal Completion (SPOC), iteratively completes the matrix such that the fifth singular value of the matrix is minimized (thus enforcing the constraint that the resulting matrix should be as close to rank-4 as possible). SPOC yields an optimal solution, as long as the occlusion pattern follows a Young diagram *i.e*., the number of occluded coordinates are a monotonic function of the rows or columns and they are not alternated with known coordinates (*e.g*., they occupy the upper or lower triangular of the observation matrix). Besides this, it requires that no gross outliers are present.

More recently, a spur of work in the optimization community has spun optimal methods for matrix completion [4, 7, 5] when this matrix is known to be low rank by minimizing the nuclear norm [6, 4, 7, 5, 8]. In contrast to SPOC, these approaches are able to deal with arbitrary patterns of occlusion and gross outliers. Albeit showing interesting theoretical properties such as optimality and automatic rank discovery, these methods are plagued by the need to calculate the Singular Value Decompositions of the entire currently estimate of the data matrix on every iteration for each addition of a new frame. Therefore, these factors make them improper for use in situations when the estimation is to be done sequentially, *i.e*., by building up on the previous iterations to calculate subsequent ones.

The method we propose assumes, instead, an initial subset of frames is known. Then, for each frame, it alternates between the use of nuclear norm minimization for outlier removal on the known trajectories and SPOC for matrix completion. The combined use of these techniques allows the completion for an arbitrary pattern of occlusion in an iterative fashion, while keeping the number and size of Singular Value Decompositions in the nuclear norm minimization to a minimum, since they are only applied to smaller matrices, thus resulting in a faster algorithm.

2. Proposed approach

The problem of recovering the full set of trajectories can be formulated as estimating the incomplete entries of a matrix $\mathbf{W}$, while correcting a subset of known ones which have been contaminated by noise. For the case of a rigid object, it has been shown [9, 1] that the matrix piling the trajectories along the various frames can be obtained by the product of a matrix $\mathbf{M}$ comprising the camera motion matrices and the 3D shape matrix $\mathbf{S}$, as

$$\mathbf{W} = \mathbf{M}\mathbf{S}^\top. \tag{1}$$

As such, the resulting measurement matrix $\mathbf{W}$ has, in absence of noise, a rank less or equal than 4.

Aguiar *et al*. [3] showed that for a matrix $\mathbf{W}$ with a single entry x missing

$$\mathbf{W} = \begin{bmatrix} x & \mathbf{v}^\top \\ \mathbf{u} & \mathbf{A} \end{bmatrix}, \tag{2}$$

the completion such that its fifth singular value $\hat{\sigma}_5$ is minimized is, under broad conditions, unique. To obtain the solution, they use the Cauchy Interlacing Theorem to place a tight bound on $\hat{\sigma}_5$, as

$$\hat{\sigma}_5 \approx \max\left\{\sigma_5\left(\begin{bmatrix} \mathbf{v}^\top \\ \mathbf{A} \end{bmatrix}\right), \sigma_5\left(\begin{bmatrix} \mathbf{u} & \mathbf{A} \end{bmatrix}\right)\right\}. \tag{3}$$

Then, they obtain the completion in closed form as the root of the characteristic polynomial of the matrix $\mathbf{W}\mathbf{W}^\top$

$$p(x) = \det(\mathbf{W}\mathbf{W}^\top - \hat{\sigma}_5^2\mathbf{I}) = 0, \tag{4}$$

which results in a quadratic equation

$$ax^2 + bx + c = 0 \tag{5}$$

where $b^2 = 4ac$, the coefficients a, b, c are given by

$$a = \det\begin{bmatrix} 0 & \mathbf{u}^\top \\ \mathbf{u} & \mathbf{B} \end{bmatrix} \tag{6}$$

$$b = 2\det\begin{bmatrix} 0 & \mathbf{v}^\top\mathbf{A}^\top \\ \mathbf{u} & \mathbf{B} \end{bmatrix} \tag{7}$$

$$c = \det\begin{bmatrix} 0 & \mathbf{v}^\top\mathbf{A}^\top \\ \mathbf{A}\mathbf{v} & \mathbf{B} \end{bmatrix}, \tag{8}$$

and $\mathbf{B} = \mathbf{u}\mathbf{u}^\top + \mathbf{A}\mathbf{A}^\top - \hat{\sigma}_5^2\mathbf{I}$.

Furthermore, to complete a matrix exhibiting a pattern of occlusion that is a Young diagram such as

$$\mathbf{W} = \begin{bmatrix} ? & ? & ? & ? & M_{1,5} \\ ? & ? & ? & M_{2,4} & M_{2,5} \\ ? & ? & M_{3,3} & M_{3,4} & M_{3,5} \\ ? & ? & M_{4,3} & M_{4,4} & M_{4,5} \\ M_{5,1} & M_{5,2} & M_{5,3} & M_{5,4} & M_{5,5} \end{bmatrix}, \tag{9}$$

they show that sequentially estimating the unknown values from the left to the right and from the bottom to the top yields the optimal reconstruction for $\mathbf{W}$. This approach, called SPectrally Optimal Completion (SPOC), has an overall complexity of $\mathcal{O}\big(N \times \max\{S, D\}\big)$, where S and D denote the complexities of performing the SVDs in (3) and the determinants in (6).

This approach is naturally extensible to an incremental version. In this case, we want to perform the spectrally optimal completion of the matrix

$$\mathbf{W}^{(i+1)} = \begin{bmatrix} \mathbf{?} & \hat{\mathbf{W}}^{(i+1)} & \mathbf{?} \\ & \mathbf{W}^{(i)} & \end{bmatrix} \tag{10}$$

that stacks a set of new measurements $\hat{\mathbf{W}}^{(i+1)}$ on top of the matrix $\mathbf{W}^{(i)}$ containing the reconstruction calculated on the previous iteration.

Despite its speed and easy extension to the incremental version, SPOC has two caveats: 1) It assumes the values given for the estimation are not corrupted by noise and 2) to be able to deal with arbitrary patterns, a permutation matrix has to be determined to turn the matrix into the young diagram format, something which is not always possible. In the remainder of this section, we show how to deal with these two shortcomings.

Finding Permutations Let us consider, without any loss of generality, we have already access to a complete set of trajectories $\mathbf{W}^i$ and are given a partial set of observations $\hat{\mathbf{W}}^{(i+1)}$, where only estimates for partial entries are available

As already discussed, in order to estimate the trajectories of the new frame using SPOC, we have to permute the whole observation matrix $\mathbf{W}^{(i+1)}$ such that it obeys a Young diagram. At this point, we should note that due to the nature of the problem, we know that the (x, y) coordinates for each point are either known or unknown together. Therefore, the permutation obtained is optimal in the sense that, for each new frame, it is always able to convert the matrix to the form of a Young diagram.

Also, since we know the coordinates x and y are intrinsically related, we make a single pass through all points and position them counting from the beginning or the end, if they are (respectively) known or unknown entries. The resulting algorithm (Alg. 1), has a linear complexity in the number of points, therefore not raising the overall order of complexity of the method.

After applying the permutation, the obtained matrix $\mathbf{W}^{(i+1)}\mathbf{P}$ obeys a young diagram of the form

$$\mathbf{W}^{(i+1)}\mathbf{P} = \begin{bmatrix} \mathbf{?} & \hat{\mathbf{W}}^{(i+1)}\mathbf{P}_2 \\ & \mathbf{W}^{(i)}\mathbf{P} \end{bmatrix}, \tag{11}$$

which allows the completion to be obtained by feeding the matrix to SPOC.

Outlier removal The method as presented thus far is able to do the trajectory estimation sequentially, but it still lies on the assumption that the points known on the new frame do not contain errors. In this paragraph, we let go of this assumption to deal with the possible existence of gross, but sparse, outliers in the tracked coordinates.

We first note that there are several algorithms [4, 7, 5, 8] available for matrix completion, when cast as the optimization problem of minimizing the nuclear norm (the sum of singular values) of a matrix $\mathbf{A}$ that has a set of known entries, as

$$\begin{array}{ll} \text{minimize} & ||\mathbf{A}||_* \\ \text{subject to} & \mathbf{D}_{ij} = \mathbf{A}_{ij}, \quad \forall (i,j) \in \Omega. \end{array} \tag{12}$$

It has already been shown [?] that this minimization, which is the convex envelope of the optimization problem

$$\begin{array}{ll} \text{minimize} & \text{rank}(\mathbf{A}) \\ \text{subject to} & \mathbf{D}_{ij} = \mathbf{A}_{ij}, \quad \forall (i,j) \in \Omega. \end{array} \tag{13}$$

actually achieves the same minimizer, under broad conditions. Therefore, this method yields the optimal reconstruction of the low rank matrix $\mathbf{A}$.

The use of this technique would allow for solving the problem as a whole. Its use of singular value decompositions of the entire matrix at each iteration, however, makes this method impracticable for real time sequential estimation.

Final Algorithm To get the best of both worlds, we propose the use of Robust PCA [7] on the sub-matrix $\tilde{\mathbf{W}}^{(i+1)}$, that stacks $\hat{\mathbf{W}}^{(i+1)}\mathbf{P}_2$ in (11) with the respective columns of $\mathbf{W}^{(i)}\mathbf{P}$. This technique is intrinsically tied with the formulation of matrix completion as in (12), with its only difference being that there are no unknown entries, so (12) becomes

$$\begin{array}{ll} \text{minimize} & ||\mathbf{A}||_* + \lambda ||\mathbf{E}||_1 \\ \text{subject to} & \mathbf{D} = \mathbf{A} + \mathbf{E}, \end{array} \tag{14}$$

where $\mathbf{E}$ models a matrix of outliers, intended to be as sparse as possible. In this case as in (12), there is evidence [7] this problem exhibits the same minimizer as its non-convex counterpart

$$\begin{array}{ll} \text{minimize} & \text{rank}(\mathbf{A}) + \lambda ||\mathbf{E}||_0 \\ \text{subject to} & \mathbf{D} = \mathbf{A} + \mathbf{E}. \end{array} \tag{15}$$

Although we still are limited by the performance of singular value decompositions in the Robust PCA method, they do not put such a heavy burden on the global time since they are dealing with matrices of smaller dimension and with no occluded entries. The final algorithm, which joins both methods, is summarized in Alg. 2.

Algorithm 1 Finding permutations for SPOC

```
Initialize known count k = 0
Initialize unknown count u = 0
for all points i ∈ 1 : N do
    if point i is known then
        Increment k
        P_{i,k} = 1
    else
        P_{i,N−u} = 1
        Increment u
    end if
end for
```

Algorithm 2 Joint Matrix Completion using SPOC and RPCA

```
for all new frames i ∈ 1 : N do
    Calculate P s.t. W^(i+1) P has a Young pattern (Alg. 1)
    Perform RPCA (14) on W̃^(i+1) to remove possible outliers
    Estimate completion for frame (i+1) using SPOC with refined W̃^(i+1) and remaining data in W^(i)
end for
```

In the development of this method, we have assumed the existence of a set of initial known frames $\mathbf{W}^{(i)}$. It should be noted, however, that this initialization is not critical, as this sub-matrix can also be found from partial observations by using existing algorithms for problem (12). We should also note the fact that although doing sequential permutations might lead to suboptimal completions — since matrix as a whole may not be converted to a young diagram — this is attenuated by the fact that we use Robust PCA on each iteration. As the number of new frames grows, this iteration on the sub-matrices should correct all of the outliers in the original matrix, thus yielding asymptotically optimal results.

3. Experiments

In this section, we perform a few experiments with synthetic data to illustrate the effectiveness and speed of our method.

Example of completion To motivate the importance of this problem, we generate trajectories of a rigid object (a sphere) (Fig. 1) and depict what a typical scenario of tracking would give as a result (Fig. 2). In this case, we have originally $N = 49$ points tracked around $f = 30$ cameras, but only 70% are present in the measurement matrix. By using our method, we are able to retrieve the original set of trajectories (Fig 3), despite the significant percentage of occluded entries.

Comparison of methods We generate 40×20 matrices with points living in the interval $[-50, 50]$ and having approximately rank 4 and occlude, for each frame matrix

2×20, about 40% of the points. To compare with existing state-of-the-art, we measure the norm of the residuals against ground truth and the total runtime of the algorithm. As a representative of the state-of-the-art, we use Singular Value Thresholding (SVT) [4] and set its parameters τ and `tol` to achieve better accuracy in detriment of speed. We choose the values such that the accuracy results obtained are comparable with what is obtained by our method. The average residual plot (Fig. 4) shows that even for the case where the accuracy is worst (thus, speed is better) than what is obtained by our method (SPOC+RPCA), the latter still outperforms the former in runtime (Fig. 5) for several frames. Since we are not on the presence of outliers, SPOC is able to recover the matrix with less error than the remaining methods for a small set of.

To demonstrate the resilience of our method to outliers, we added to approximately 10% of the non occluded points sparse noise generated from a uniform distribution on the interval $[-15, 15]$ and measure the average residual (Fig. 6) and runtime (Fig. 6). Here, the performance of SPOC is clearly undermined, since we are not able to trust entirely on the known entries for completing the matrix. SVT and RPCA+SPOC achieve performance comparable to what was obtained for the case when no outliers are present, both in residual norm and runtime. In both cases, our method (RPCA+SPOC) clearly outperforms SVT in run time for the case when only a subset of frames are to be estimated. The monotonic tendency on RPCA+SPOC suggests that for as the number of frames to be estimated is significantly larger than the number of frames assumed known, it should take the same time than SVT to perform the completion. On this point, there is no difference between either

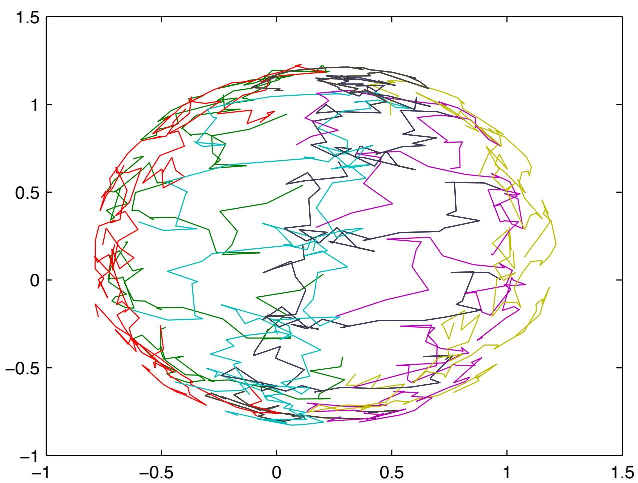

Figure 1. Ground truth trajectories.

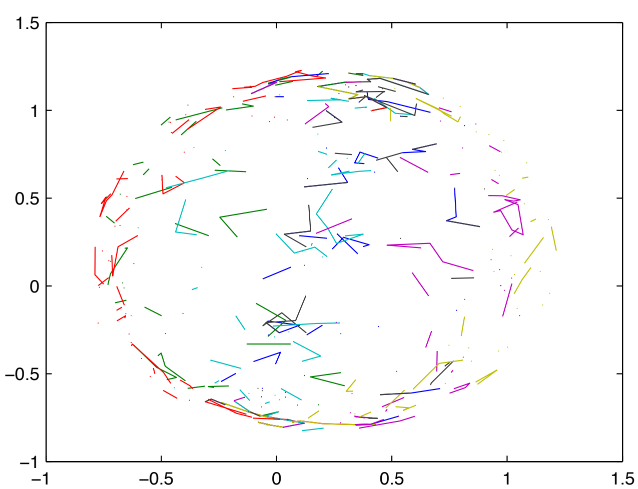

Figure 2. Partial data.

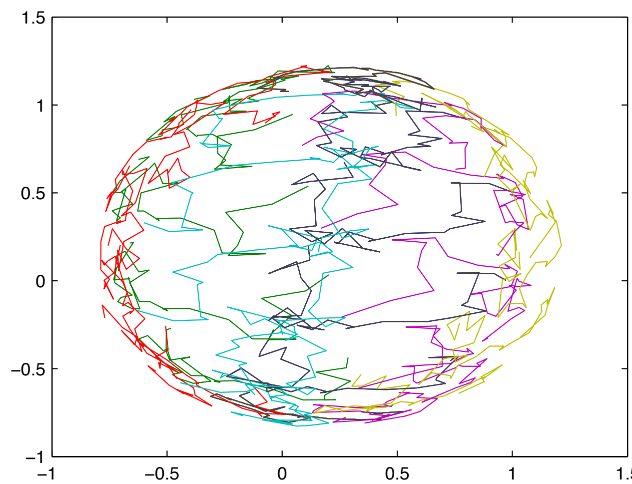

Figure 3. Estimated trajectories.

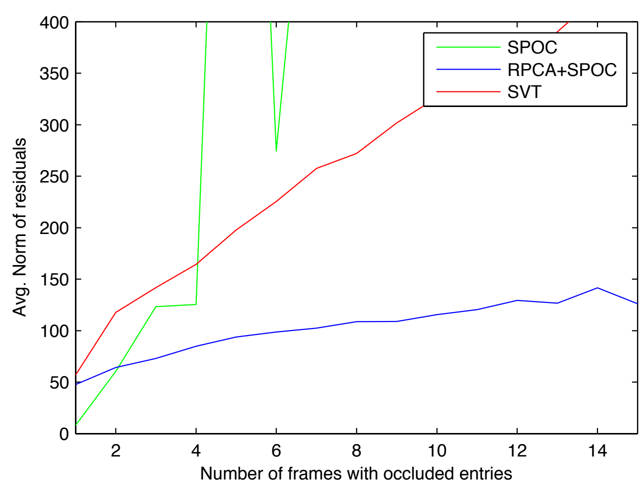

Figure 4. Variation of average norm of residuals with number of occluded frames without.

methods. It should be noted, however, that these numbers are in respect to estimating 75% of the matrix. As such, on an incremental framework, the considered time for each frame should be the one indicated for the estimation of one occluded frame (in the left region of the plot), where our method outperforms RPCA+SPOC by a significant margin, for comparable accuracy.

4. Conclusions

We have presented a method to recover point trajectories of a rigid object while being subject to gross outliers and self-occlusion. Results show that our method beats in terms of speed existing algorithms performing the same task, while achieving comparable estimation performance.

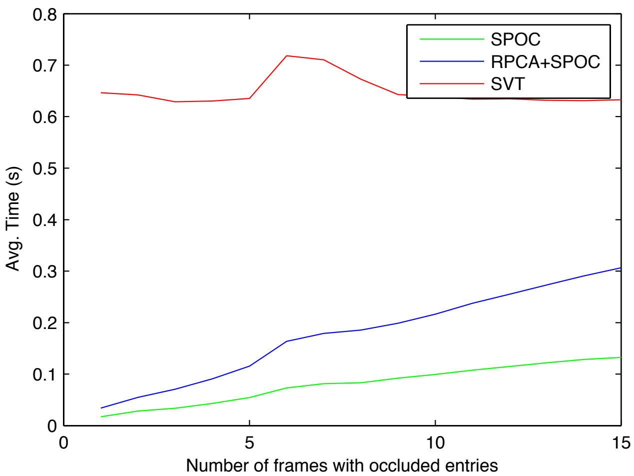

Figure 5. Variation of average runtime with number of occluded frames without outliers.

Our algorithm can be applied sequentially, allowing for real-time implementations, where only a small subset of the data is available beforehand.

Further work should exploit the formulation of a spectrally optimal completion while subject to outliers in the obtained data points into a single optimization problem, so as to avoid the use of matrix completion algorithms on the same columns more than once.

References

[1] P. Aguiar and J. Moura. Rank 1 weighted factorization for 3d structure recovery: Algorithms and performance analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(9):1134–1149, Sept. 2003. 1

[2] P. Aguiar, J. Xavier, and M. Stosic. Spectrally optimal factorization of incomplete matrices. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pages 1 –8, 23-28 2008. 1

[3] P. M. Aguiar, M. Stosic, and J. Xavier. On singular values of partially prescribed matrices. *Linear Algebra and its Applications*, 429(8-9):2136 – 2145, 2008. 1, 2

[4] J.-F. Cai, E. J. Candes, and Z. Shen. A singular value thresholding algorithm for matrix completion. 2008. 1, 3, 4

[5] E. Candes and B. Recht. Exact low-rank matrix completion via convex optimization. In *Communication, Control, and Computing, 2008 46th Annual Allerton Conference on*, pages 806 –812, 23-26 2008. 1, 3

[6] S. Ji and J. Ye. An accelerated gradient method for trace norm minimization. In *ICML '09: Proceedings of the 26th Annual International Conference on Machine*

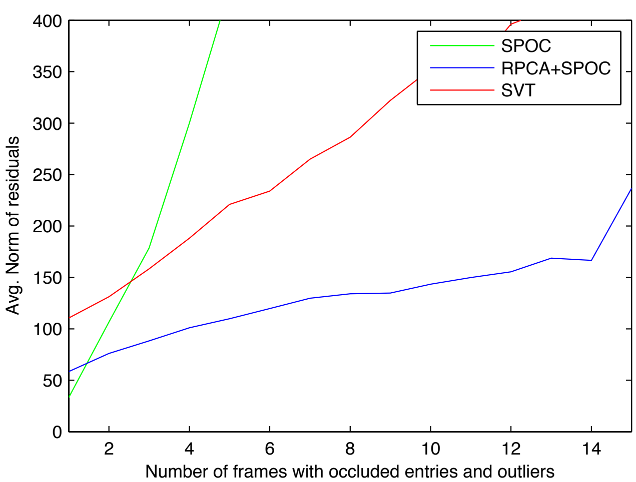

Figure 6. Variation of average norm of residuals with number of occluded frames with outliers.

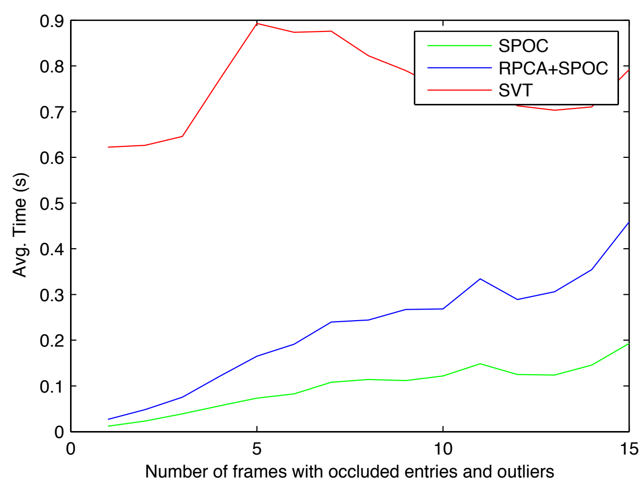

Figure 7. Variation of average runtime with number of occluded frames with outliers.

Learning, pages 457–464, New York, NY, USA, 2009. ACM. 1

[7] Z. Lin, A. Ganesh, J. Wright, L. Wu, M. Chen, and Y. Ma. Fast convex optimization algorithms for exact recovery of a corrupted low-rank matrix. preprint, July 2009. 1, 3

[8] K.-C. Toh and S. Yun. An accelerated proximal gradient algorithm for nuclear norm regularized least squares problems. preprint, 2009. 1, 3

[9] C. Tomasi and T. Kanade. Shape and motion from image streams under orthography: a factorization method. *Int. J. Comput. Vision*, 9(2):137–154, 1992. 1