

FAST INCREMENTAL METHOD FOR MATRIX COMPLETION: AN APPLICATION TO TRAJECTORY CORRECTION

Ricardo S. Cabral^{a,b}, João P. Costeira^a, Fernando De la Torre^b and Alexandre Bernardino^a

^aISR - Instituto Superior Técnico, Univ. Técnica de Lisboa (Lisboa, Portugal)

^bCarnegie Mellon University (Pittsburgh, USA)

ABSTRACT

We address the problem of incrementally recovering a matrix of tracked image points, based on partial observations of their trajectories. Besides partial observability, we assume the existence of gross, but sparse, noise on the known entries. This problem has obvious applications in real-time tracking and structure from motion, where observations are plagued by self-occlusion and outliers. Recently, 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, the sum of its singular values. Despite exhibiting several optimality properties, no available algorithms perform this minimization incrementally.

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 while maintaining performance comparable to the state-of-the-art.

Index Terms— Incremental Matrix Completion, Nuclear Norm, Outliers, Structure from Motion, Tracking

1. INTRODUCTION

Several applications in computer vision [1, 10, 2] rely on a tracking process to produce a set of point trajectories of an object moving along several image frames. However, this process is usually plagued by self-occlusion and drifting, resulting in gross outliers and partial trajectories [1]. We focus on the problem of recovering in an incremental fashion the full set of trajectories of a rigid object tracked along several images, while correcting existing outliers. We formulate a matrix completion algorithm which, for every addition of a new frame, minimizes the rank of the resulting measurement matrix, which we know beforehand to be close to 4. 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 while enforcing the obtained result to be as close to rank-4 as possible. SPOC is optimal, as long as the occlusion pattern follows a Young diagram, *i.e.*, the number of occluded coordinates is a monotonic function of the rows or columns and they are not alternated with known coordinates (*e.g.*, they occupy the upper triangular of the observation matrix). In practice, we do not have this pattern due to points disappearing and reappearing in different positions due to self-occlusion. Moreover, their method relies on a bounded estimate for the fifth singular value, which is contaminated by the existence of gross outliers.

In the last few years, work in the optimization community has spun several methods for matrix completion [4, 5, 6, 7, 8] when this matrix is known to be low rank. One popular approach is to minimize the nuclear norm (the sum of singular values) [4, 5, 6, 9]. It has been shown [5] that this function, the convex envelope of the rank operator, actually achieves the same minimizer under broad conditions. In contrast to SPOC, these approaches are able to deal with arbitrary patterns of occlusion and gross outliers, thus allowing a solution for the trajectory correction problem as a whole. Their use of singular value decompositions of the entire matrix at each iteration of the gradient, however, makes them impracticable for real time sequential estimation. The method we propose assumes, instead, an initial subset of frames is known. Then, it builds up on the convergence achieved for each frame to calculate the subsequent one. For each new 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 coping with outliers and arbitrary occlusion patterns in an iterative fashion, while keeping the number and size of Singular Value Decompositions in the nuclear norm minimization to a minimum, thus resulting in a faster algorithm.

Very recently, Balzano *et al.* [7] proposed an incremental matrix completion algorithm, by performing a gradient descent on the Grassmanian manifold of rank deficient matrices (GROUSE). In this paper, we show that this approach, albeit very fast, yields worse results than other methods (and our own) when data is subject to gross sparse outliers, due to the fact that it estimates the row and column subspaces of the complete matrix rather than the entries themselves.

Support for this research was provided by the Fundação para a Ciência e a Tecnologia (Portuguese Foundation for Science and Technology) through the Carnegie Mellon Portugal program under the project FCT/CMU/P11. Partially funded by FCT project Printart PTDC/EEA-CRO/098822/2008.

2. PROPOSED APPROACH

The problem of recovering the full set of N point trajectories along F partial observations can be formulated as estimating incomplete entries of a rank deficient matrix $\mathbf{W} \in \mathbb{R}^{2F \times N}$, while correcting the subset of known ones that have been contaminated by noise. For a rigid object, it has been shown [10] that a matrix piling point trajectories along several frames can be obtained by the product of a matrix \mathbf{M} stacking the camera motion matrices and a 3D shape matrix \mathbf{S} , as

$$\mathbf{W} = \mathbf{M}\mathbf{S}^\top. \quad (1)$$

As such, the resulting measurement matrix \mathbf{W} has, in absence of noise, a rank that is less or equal than 4. Aguiar *et al.* [3] showed that for a matrix \mathbf{W} with a single entry x missing

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

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

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

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

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

a quadratic equation $ax^2 + bx + c = 0$ with $b^2 = 4ac$ and coefficients a, b, c given by

$$a = \det \begin{bmatrix} 0 & \mathbf{u}^\top \\ \mathbf{u} & \mathbf{B} \end{bmatrix} \quad (5)$$

$$b = 2 \det \begin{bmatrix} 0 & \mathbf{v}^\top \mathbf{A}^\top \\ \mathbf{u} & \mathbf{B} \end{bmatrix} \quad (6)$$

$$c = \det \begin{bmatrix} 0 & \mathbf{v}^\top \mathbf{A}^\top \\ \mathbf{A}\mathbf{v} & \mathbf{B} \end{bmatrix}, \quad (7)$$

where $\mathbf{B} = \mathbf{u}\mathbf{u}^\top + \mathbf{A}\mathbf{A}^\top - \hat{\sigma}_5^2 \mathbf{I}$. Furthermore, to complete a matrix exhibiting a pattern of occlusion that is a Young diagram, they show that sequentially estimating the unknown values from the right to the left 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 $\mathcal{O}(N_o \times \max\{S, D\})$, where $N_o < N$ is the number of occluded points and S, D denote the complexities of finding the bound in (3) and determinants in (5). The incremental nature of SPOC allows us to cast the problem of completing trajectories in this fashion. In this case, we want to perform the spectrally optimal completion of the matrix

$$\tilde{\mathbf{W}}^{(i+1)} = \begin{bmatrix} ? & \hat{\mathbf{W}}^{(i+1)} & ? \\ & \mathbf{W}^{(i)} & \end{bmatrix} \quad (8)$$

that stacks a set of new measurements $\hat{\mathbf{W}}^{(i+1)}$ with occluded entries $?$ on top of matrix $\mathbf{W}^{(i)}$, containing the reconstruction of the previous iteration. Despite its speed and immediate applicability to incremental estimation, SPOC has two caveats: 1) It assumes known values are not corrupted by noise and 2) to be able to deal with arbitrary patterns, a permutation matrix has to be determined to convert the matrix to a Young diagram format, in general 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 new set of observations $\hat{\mathbf{W}}^{(i+1)}$, for only a subset of the tracked points.

As already discussed, in order to estimate the trajectories of the new frame using SPOC, we have to permute the whole observation matrix $\tilde{\mathbf{W}}^{(i+1)}$ 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. Moreover, this intrinsic relation between coordinates allows to find the permutation in a deterministic fashion. We make a single pass through all points in frame $(i+1)$ and position them counting from the beginning or the end, according to whether they are (respectively) known or unknown entries. The resulting algorithm (Alg. 1) has linear complexity in the number of points N , therefore not raising the overall order of complexity of the method.

Algorithm 1 Finding permutations for SPOC at frame $(i+1)$

```

Initialize known count  $k = 0$ 
Initialize unknown count  $u = 0$ 
for all Points  $j \in 1 : N$  do
  if Point  $\tilde{W}_{(i+1),j}$  is known then
    Increment  $k$ 
     $P_{j,N-k} = 1$ 
  else
     $P_{j,u} = 1$ 
    Increment  $u$ 
  end if
end for

```

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

$$\tilde{\mathbf{W}}^{(i+1)}\mathbf{P} = \begin{bmatrix} ? & \hat{\mathbf{W}}^{(i+1)} \\ & \mathbf{W}^{(i)}\mathbf{P} \end{bmatrix}, \quad (9)$$

allowing its completion to be obtained with SPOC.

Outlier removal The method as presented thus far is able to do trajectory estimation sequentially, but still relies on the assumption that points known at frame $(i+1)$ do not contain

errors. We let go of this assumption and deal with the possible existence of gross, but sparse, outliers in the tracked coordinates. To do so, we note that matrix completion algorithms based on nuclear norm minimization, albeit slow, provide this needed resilience [4, 5, 6, 9]. To get the best of both worlds, we propose the use of Robust PCA [5] on the known columns of $\tilde{\mathbf{W}}^{(i+1)}\mathbf{P}$. This technique is intrinsically tied with the formulation of nuclear norm matrix completion, with its only difference being that there are no unknown entries, as

$$\begin{aligned} & \text{minimize} && \|\mathbf{A}\|_* + \lambda\|\mathbf{E}\|_1 \\ & \text{subject to} && \mathbf{W} = \mathbf{A} + \mathbf{E}, \end{aligned} \quad (10)$$

where \mathbf{E} models a matrix of outliers, as sparse as possible. As with Matrix Completion, there is evidence [5] this problem exhibits the same minimizer as its non-convex counterpart

$$\begin{aligned} & \text{minimize} && \text{rank}(\mathbf{A}) + \lambda\|\mathbf{E}\|_0 \\ & \text{subject to} && \mathbf{W} = \mathbf{A} + \mathbf{E}. \end{aligned} \quad (11)$$

Although we are still 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 applied to smaller dimension matrices with no occluded entries. The final algorithm, which joins both methods, is summarized in Alg. 2.

Algorithm 2 Joint Matrix Completion using SPOC and RPCA

for all new frames $i \in 1 : F$ **do**
 Calculate \mathbf{P} s.t. $\tilde{\mathbf{W}}^{(i+1)}\mathbf{P}$ has a Young pattern (Alg. 1)
 Remove outliers from known columns of $\tilde{\mathbf{W}}^{(i+1)}\mathbf{P}$ using RPCA (10)
 Estimate $\mathbf{W}^{(i+1)}$ using SPOC on corrected $\tilde{\mathbf{W}}^{(i+1)}\mathbf{P}$
end for

Note that in this method, we assumed the existence of an initial set of 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 matrix completion algorithms. We should also note the fact that although doing sequential permutations might lead to suboptimal completions — since the 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, the RPCA step on the sub-matrices should correct all of the outliers in the original matrix.

3. EXPERIMENTS

To motivate the importance of this problem, we generate trajectories of a sphere (Fig. 1(a)) and depict a typical scenario of tracking with self-occlusion (Fig. 1(b)). In this case, we have $N = 64$ points tracked around $F = 30$ cameras, but only 70% of the entries are known and 20% have outliers.

Using our method, we are able to retrieve the full original set of trajectories (Fig. 1(c)).

Comparison with state-of-the-art We generate random matrices \mathbf{W} as in (1) with $N = 100$ points along $F = 20$ frames, where \mathbf{M} and \mathbf{S} are sampled from a uniform distribution living in the interval $[0, 4]$. We occlude from the first $2 \times N$ frame a variable number of entries ($N_o \in [1, 80]$). Then, we add to approximately 25% of the non occluded points sparse noise generated from a uniform distribution on the interval $[-3, 3]$. We measure the mean norm of the difference between completed entries against ground truth (normalized by the ground truth norm) and the total runtime of the algorithm. We compare our results against Singular Value Thresholding (SVT) [4], OptSpace [8] and GROUSE [7], using code provided by the methods authors. We use $\tau = 25\sqrt{N \times F}$ and $\delta = 2$ in SVT to increase speed while maintaining comparable accuracy. For OptSpace, we use $\text{tol} = 10^{-6}$. For each N_o , results are averaged among 50 tests, with approximately 10 being discarded in the OptSpace case due to numerical stability problems. Fig. 2 shows obtained error vs. time. Results show that incremental methods RPCA+SPOC and GROUSE are significantly faster than their non-incremental counterparts, with the former outperforming the latter in accuracy.




Fig. 2. Average time vs. average error for single frame completion. Time standard deviation is negligible at shown scale.

To assess outlier correction and incremental completion performance, we generate $N = 100$, $F = 10$ matrices \mathbf{W} and corrupt 25% of its entries with outliers (generated as before). Then, we incrementally concatenate this matrix with another $F = 10$ measurements, varying the number of frames having half its entries randomly occluded. Figure 3 shows the normalized residual norm separated by known and unknown entries of the entire matrix, averaged over 40 tests. Results corroborate the incremental stability of SPOC+RPCA. The monotonic tendency in the error for known entries shows the effectiveness of random permutations, necessary for conversion to young formats, in the correction of randomly positioned outliers in the entire matrix by applying RPCA to just a subset of the matrix. Furthermore, the known entries results show that when compared to baseline (original SPOC,




Fig. 1. Tracking a ball with self-occlusion and outliers. (a) Ground truth; (b) Partial data; (c) Estimated trajectories.

which leaves known entries intact), nuclear norm methods such as SVT are able to correct outliers, whereas subspace estimation methods like OptSpace or GROUSE increase the corruption.

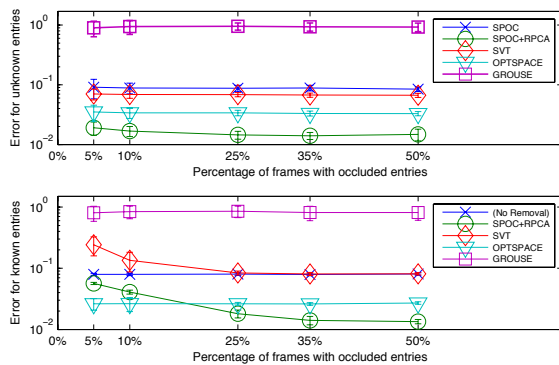


Fig. 3. Average normalized residual norm for unknown (top) and known (bottom) entries of matrix \mathbf{W} .

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 outperforms existing algorithms performing the same task, both in accuracy and speed. 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 Robust PCA algorithms on the same columns more than once.

5. REFERENCES

- [1] B. D. Lucas and T. Kanade, “An iterative image registration technique with an application to stereo vision,” in *IJCAI81*, 1981, pp. 674–679.
- [2] P. Aguiar, J. Xavier, and M. Stolic, “Globally optimal solution to exploit rigidity when recovering structure from motion under occlusion,” in *Image Processing, 2008. ICIP 2008. 15th IEEE International Conference on*, 2008, pp. 197–200.
- [3] Pedro M.Q. Aguiar, Marko Stolic, and João Xavier, “On singular values of partially prescribed matrices,” *Linear Algebra and its Applications*, vol. 429, no. 8-9, pp. 2136–2145, 2008.
- [4] Jian-Feng Cai, Emmanuel J. Candes, and Zuowei Shen, “A singular value thresholding algorithm for matrix completion,” 2008.
- [5] 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.
- [6] E.J. Candes and B. Recht, “Exact low-rank matrix completion via convex optimization,” in *Communication, Control, and Computing, 2008 46th Annual Allerton Conference on*, 23-26 2008, pp. 806–812.
- [7] Laura Balzano, Robert Nowak, and Benjamin Recht, “Online identification and tracking of subspaces from highly incomplete information,” *CoRR*, vol. abs/1006.4046, 2010.
- [8] Raghunandan H. Keshavan, Andrea Montanari, and Seungwon Oh, “Matrix completion from a few entries,” *IEEE Trans. Inf. Theor.*, vol. 56, pp. 2980–2998, June 2010.
- [9] K.-C. Toh and S. Yun, “An accelerated proximal gradient algorithm for nuclear norm regularized least squares problems,” preprint, 2009.
- [10] Carlo Tomasi and Takeo Kanade, “Shape and motion from image streams under orthography: a factorization method,” *Int. J. Comput. Vision*, vol. 9, no. 2, pp. 137–154, 1992.