The Generalized Trace-Norm and its Application to Structure-from-Motion Problems

Roland Angst   Christopher Zach   Marc Pollefeys

Computer Vision and Geometry Group, Eidgenössische Technische Hochschule Zürich, Switzerland

{rangst, christopher.zach, marc.pollefeys}@inf.ethz.ch

Abstract

In geometric computer vision, the structure from motion (SfM) problem can be formulated as a optimization problem with a rank constraint. It is well known that the trace norm of a matrix can act as a convex proxy for a low rank constraint. Hence, in recent work [7], the trace-norm relaxation has been applied to the SfM problem. However, SfM problems often exhibit a certain structure, for example a smooth camera path. Unfortunately, the trace norm relaxation can not make use of this additional structure. This observation motivates the main contribution of this paper. We present the so-called generalized trace norm which allows to encode prior knowledge about a specific problem into a convex regularization term which enforces a low rank solution while at the same time taking the problem structure into account.

While deriving the generalized trace norm and stating its different formulations, we draw interesting connections to other fields, most importantly to the field of compressive sensing. Even though the generalized trace norm is a very general concept with a wide area of potential applications we are ultimately interested in applying it to SfM problems. Therefore, we also present an efficient algorithm to optimize the resulting generalized trace norm regularized optimization problems. Results show that the generalized trace norm indeed achieves its goals in providing a problem-dependent regularization.

1. Introduction

Consider an optimization problem over a matrix-valued variable $\mathbf{X} \in \mathbb{R}^{m \times n}$ which is convex except for a constraint of the form $\text{rank}(\mathbf{X}) \leq r$. Facing this non-convex constraint there are two possibilities: Either using direct non-linear non-convex optimization methods or convexifying the non-linear rank constraint with a convex function. The result of direct non-linear optimization highly depends on the initial values and without any good initial guess there is

little hope to converge to a good solution. With a convex relaxation of the rank the original problem becomes convex and hence convex optimization is guaranteed to find the optimal solution. This solution however is the solution to a modified convex relaxed problem and might not equal the optimal solution to the non-relaxed problem. There are pros- and cons to either approach and we do not claim one to be superior to the other for general problems. However, for specific problems such as matrix completion problems the convex relaxation has been shown to work surprisingly well [15, 5], in theory as well as in practice. Recent results [9, 13] provide strong guarantees under which conditions the solutions of the convex relaxation and the original non-convex problem are equivalent. These observations together with the previously mentioned convex relaxation of the SfM problem based on the trace-norm [7] motivated the present paper and we will thus focus in the upcoming sections solely on convex relaxations of a non-convex rank-constraint problem. Related work will be presented interspersed throughout the paper since once the necessary background information is acquired and the notation is in place it is simpler to draw parallels and to highlight differences to related work.

2. Trace-Norm

The trace-norm of a matrix $\mathbf{X} \in \mathbb{R}^{m \times n}$ is defined as the sum of the singular values of the matrix $\mathbf{X}$

$$\|\mathbf{X}\|_* = \sum_{i=1}^{\min(m,n)} \sigma_i(\mathbf{X}),$$

where $\sigma_i(\mathbf{X})$ denotes the i^{th} singular value of $\mathbf{X}$. What the $L1$-norm is for vectors, the trace norm is for matrices. The trace norm is also known as the $\min(m, n)$-Ky-Fan norm or the nuclear norm [10]. We will stick with the name trace norm for reasons which will become clear shortly. The trace norm is proportional to the the tightest convex lower bound of the rank function over the set of matrices with spectral norm smaller than β, i.e. the convex envelope of the rank function $\text{rank}(\mathbf{X})$ over the set $\{\mathbf{X} \in \mathbb{R}^{m \times n} | \|\mathbf{X}\|_2 \leq \beta\}$

is $\frac{1}{\beta}\|\mathbf{X}\|_*$. The trace norm of a matrix $\mathbf{X} \in \mathbb{R}^{m \times n}$ can be computed with a Semi-Definite Programming (SDP) formulation (e.g. see [15])

$$\|\mathbf{X}\|_* = \frac{1}{2} \min_{\mathbf{X}_1, \mathbf{X}_2} \left(\text{trace}\,(\mathbf{X}_1) + \text{trace}\,(\mathbf{X}_2) \right) \qquad (1)$$

$$\text{s.t.} \begin{bmatrix} \mathbf{X}_1 & \mathbf{X} \\ \mathbf{X}^T & \mathbf{X}_2 \end{bmatrix} \succeq \mathbf{0}_{m+n \times m+n}, \qquad (2)$$

where $\succeq$ denotes a generalized inequality constraint on the cone of positive-semidefinite matrices $\mathbb{S}_+^{m+n}$. SDP problems are convex. As for example stated in [3] there are three other formulations of the trace norm based on a factorization of $\mathbf{X}$ into two matrices $\mathbf{U}$ and $\mathbf{V}$

$$\|\mathbf{X}\|_* = \min_{\mathbf{X}=\mathbf{U}\mathbf{V}^T} \|\mathbf{U}\|_F \|\mathbf{V}\|_F = \min_{\mathbf{X}=\mathbf{U}\mathbf{V}^T} \sum_i \|\mathbf{U}_{:,i}\|_2 \|\mathbf{V}_{:,i}\|_2$$

$$= \frac{1}{2} \min_{\mathbf{X}=\mathbf{U}\mathbf{V}^T} \left(\|\mathbf{U}\|_F^2 + \|\mathbf{V}\|_F^2 \right). \qquad (3)$$

3. Generalized Trace-Norm

The starting point for the generalized trace norm is the following: the objective of the SDP-formulation of the trace norm in Eq. (1) actually corresponds to an inner product between an identity matrix an the matrix appearing in Eq. (2)

$$\|\mathbf{X}\|_* = \frac{1}{2} \min_{\mathbf{X}_1, \mathbf{X}_2} \left\langle \mathbf{I}_{m+n}, \begin{bmatrix} \mathbf{X}_1 & \mathbf{X} \\ \mathbf{X}^T & \mathbf{X}_2 \end{bmatrix} \right\rangle \qquad (4)$$

$$\text{s.t.} \begin{bmatrix} \mathbf{X}_1 & \mathbf{X} \\ \mathbf{X}^T & \mathbf{X}_2 \end{bmatrix} \succeq \mathbf{0}_{m+n \times m+n}. \qquad (5)$$

Here $\langle \mathbf{A}, \mathbf{B} \rangle = \text{trace}\,(\mathbf{A}^T \mathbf{B})$ denotes the Frobenius inner product. We define for any matrix $\mathbf{X} \in \mathbb{R}^{m \times n}$ the generalized trace norm as

$$\|\mathbf{X}\|_{*\mathbf{D}} = \frac{1}{2} \min_{\mathbf{X}_1, \mathbf{X}_2} \left\langle \mathbf{D}, \begin{bmatrix} \mathbf{X}_1 & \mathbf{X} \\ \mathbf{X}^T & \mathbf{X}_2 \end{bmatrix} \right\rangle \qquad (6)$$

$$= \frac{1}{2} \min_{\mathbf{X}_1, \mathbf{X}_2} \text{trace}\left(\mathbf{D}^T \begin{bmatrix} \mathbf{X}_1 & \mathbf{X} \\ \mathbf{X}^T & \mathbf{X}_2 \end{bmatrix} \right) \qquad (7)$$

$$\text{s.t.} \begin{bmatrix} \mathbf{X}_1 & \mathbf{X} \\ \mathbf{X}^T & \mathbf{X}_2 \end{bmatrix} \succeq \mathbf{0}_{m+n \times m+n}, \qquad (8)$$

where $\mathbf{D} \in \mathbb{S}_+^{m+n}$ is a positive definite matrix. Hence, compared to the standard trace norm, the generalized trace norm replaces the identity with a positive semi-definite matrix and thereby introduces a non-uniform weighting, similar to a weighted $L1$-vector norm.

It remains to be shown, that this definition fulfills the axioms of a norm over the vector space of $\mathbb{R}^{m \times n}$. If the matrix $\mathbf{D}$ is not block-diagonal, then the formulation in Eq. (6) is not a norm[1]. However, if the matrix $\mathbf{D} \in \mathbb{S}_{++}^{m+n}$ is positive-definite and is chosen block-diagonal with blocks $\mathbf{D}_r$ and

---

[1]This can be seen by a numerical experiment: choosing a random matrix $\mathbf{A}$, then for a norm it should hold $\|\mathbf{A}\| = \|-\mathbf{A}\|$, however this is not true for random positive definite $\mathbf{D}$.

$\mathbf{D}_c$ of sizes $m \times m$ and $n \times n$, respectively, then [2] shows that this definition indeed is a valid norm. In the following we thus restrict ourselves to such block-diagonal choices. Note that since $\mathbf{D} \in \mathbb{S}_+^{m+n}$ the diagonal blocks must also belong to the positive semi-definite cones $\mathbf{D}_r \in \mathbb{S}_+^m$ and $\mathbf{D}_c \in \mathbb{S}_+^n$. [2] also derives an equivalent formulation based on the standard trace norm

$$\|\mathbf{X}\|_{*\mathbf{D}} = \left\| \mathbf{C}_r \mathbf{X} \mathbf{C}_c^T \right\|_*, \qquad (9)$$

where $\mathbf{C}_r = \mathbf{\Lambda}_r^{\frac{1}{2}} \mathbf{V}_r^T$ and $\mathbf{C}_c = \mathbf{\Lambda}_c^{\frac{1}{2}} \mathbf{V}_c^T$ are based on the eigen-decompositions $\mathbf{D}_r = \mathbf{V}_r \mathbf{\Lambda}_r \mathbf{V}_r^T$ and $\mathbf{D}_c = \mathbf{V}_c \mathbf{\Lambda}_c \mathbf{V}_c^T$. This representation allows a first intuitive interpretation of the generalized trace norm. The row and column spaces of the matrix $\mathbf{X}$ are projected onto the eigenspaces of $\mathbf{D}_r$ and $\mathbf{D}_c$ and scaled according to the square roots of the eigenvalues of $\mathbf{D}_r$ and $\mathbf{D}_c$. By designing the norm on purpose in this way, the generalized trace-norm is *not* a unitarily invariant norm and hence does not reduce to a symmetric gauge function [10] applied to the singular values of $\mathbf{X}$. The generalized trace norm can rather be thought of as the analog for matrices of what the weighted $L1$-vector norm is for vectors.

3.1. Factorized Formulation

The generalized trace-norm also allows a factorized interpretation instead of the SDP interpretation given first. Indeed, the SDP formulation in Eq. (6) is equivalent to the factorized interpretation

$$\|\mathbf{X}\|_{*\mathbf{D}} = \min_{\mathbf{X}=\mathbf{U}\mathbf{V}^T} \frac{1}{2} \text{trace}\left(\begin{bmatrix} \mathbf{U}^T & \mathbf{V}^T \end{bmatrix} \mathbf{D} \begin{bmatrix} \mathbf{U} \\ \mathbf{V} \end{bmatrix} \right). \qquad (10)$$

If $\mathbf{D} = \mathbf{I}_{m+n}$, then Eq. (10) reduces to Eq. (3), i.e. a sum over the squared Frobenius norms of the two factors $\mathbf{U}$ and $\mathbf{V}$. The generalized trace norm thus replaces the standard inner product in the space of $m+n \times m+n$ matrices by the weighted inner product $\langle ., . \rangle_{\mathbf{D}}$. If the matrix $\mathbf{D}$ is chosen block-diagonal with blocks $\mathbf{D}_r \in \mathbb{S}_+^m$ and $\mathbf{D}_c \in \mathbb{S}_+^n$, then the generalized trace norm reduces to

$$\|\mathbf{X}\|_{*\mathbf{D}} = \min_{\mathbf{U}, \mathbf{V}} \frac{1}{2} \left(\langle \mathbf{U}, \mathbf{U} \rangle_{\mathbf{D}_r} + \langle \mathbf{V}, \mathbf{V} \rangle_{\mathbf{D}_c} \right) \qquad (11)$$

This representation shows that the generalized trace-norm with block-diagonal matrix $\mathbf{D}$ has an interpretation as working in two inner-product spaces with inner products $\langle \mathbf{a}, \mathbf{b} \rangle_{\mathbf{D}_r} = \mathbf{a}^T \mathbf{D}_r \mathbf{b}$ and $\langle \mathbf{c}, \mathbf{d} \rangle_{\mathbf{D}_c} = \mathbf{c}^T \mathbf{D}_c \mathbf{d}$ in the range and domain of the linear mapping induced by $\mathbf{X}$, respectively. From an algorithmic point of view, such a representation is useful for stochastic gradient based optimization techniques which represent the unknown matrix explicitly in a low-rank factorized form.

3.2. Relationship to Previous Work

According to our knowledge, concepts related to the generalized trace norm have appeared only sparsely in previous work. For example, the consistency of the trace norm regularization with an L2-loss has been investigated by Bach in [4]. This paper however addressed a different question than we do. Specifically, it addresses the question of how to balance a trace-norm regularization term and a L2-loss term between measurements and an unknown matrix when the number of measurements goes to infinity (the size of the matrix stays constant). Hence, in contrast to the situation we are interested in, there are many more measurements available than linear degrees of freedom. Bach also showed in this paper that a least squares fit solely based on the L2-loss term in order to derive weighting factors can be used to derive stronger consistency results. To the best of our knowledge, this is the only paper where a trace-norm regularization term of the form of Eq. (9) has appeared. However, in strong contrast to our contribution, this weighting scheme has not been used as a general means to include prior knowledge about the specific problem at hand but rather as a means to prove stronger consistency results in the setting when the number of measurements approaches infinity. The decomposition norms introduced in [3] are also related to the generalized trace norm. The authors define a regularization term w.r.t. any norm of the column vectors $\mathbf{U}_{:,i}$ and $\mathbf{V}_{j,:}$ of the factorization $\mathbf{X} = \mathbf{U}\mathbf{V}^T$. However the case of arbitrary norms induced by an inner product $\langle \mathbf{x}, \mathbf{y} \rangle_{\mathbf{D}} = \mathbf{x}^T \mathbf{D} \mathbf{y}$ has not been investigated, the authors have rather focused their attention to the case of choosing the L2 norm in one vector space and the L1-norm in the other. A recent paper about matrix completion [17] generalized the trace norm by introducing for each of the terms $\|\mathbf{U}_{i,:}\|_F^2$ and $\|\mathbf{V}_{j,:}\|_F^2$ in the factorized trace-norm formulation of Eq. (3) an individual scaling. The insight, that this actually corresponds to a diagonal choice for $\mathbf{D}$ has been missing. Another distinction of our work is that the way in which the generalized trace norm is derived and presented, allows to draw interesting connections to related work such as matrix completion and compressive sensing.

4. Trace-Norm Regularized M-Estimators

Let us assume we are given the measurements $\mathbf{z} \in \mathbb{R}^{m \times 1}$ and we are looking for a model $\mathbf{W} \in \mathbb{R}^{p \times q}$ with $m \ll pq$ which explains the measurements as accurately as possible. The accuracy is measured with a data cost function $\mathcal{L}(\mathbf{W}, \mathbf{z})$. The situation $m \ll pq$ means that there are far fewer measurements than actual degrees of freedom in our model. Regularized M-estimators add a regularization term $\|\mathbf{W}\|$ (usually a norm) in order to handle this situation

$$\mathbf{W}_* \in \arg\min_{\mathbf{W}} \mathcal{L}(\mathbf{W}, \mathbf{z}) + \mu \|\mathbf{W}\|. \qquad (12)$$

For concreteness, the data cost function is assumed to be a function of $\mathcal{A}(\mathbf{W}) - \mathbf{z}$ where $\mathcal{A} \in \mathbb{R}^{m \times pq}$ is a linear operator relating the matrix valued unknown $\mathbf{W}$ to the measurements $\mathbf{z}$. Recent work has analyzed the situation when the regularization is chosen to be the trace norm $\|\mathbf{W}\|_*$ [13, 15]. In this paper we are interested in the case where the regularizer corresponds to the generalized trace norm

$$\mathbf{W}_* \in \arg\min_{\mathbf{W}} \mathcal{L}(\mathcal{A}(\mathbf{W}) - \mathbf{z}) + \mu \|\mathbf{W}\|_{*\mathbf{D}}. \qquad (13)$$

The immediate question is: what is the influence of the block-diagonal matrix $\mathbf{D}$ and how should it be chosen? In order to answer these questions we appeal to a Bayesian interpretation of the generalized trace norm.

4.1. Bayesian Interpretation

The representation in Eq. (11) is related to a generalized version of the Bayesian probabilistic matrix factorization [16]. There, a Bayesian hierarchical model is introduced which represents the unknown matrix $\mathbf{W} = \mathbf{U}\mathbf{V}^T$ in a factorized form. Zero-mean spherical Gaussian priors are put on the unknown factors $\mathbf{U}$ and $\mathbf{V}$. It can be shown that the negative log-posterior distribution exactly corresponds to an M-estimator with a L2-loss function and a standard trace norm regularizer. Instead of putting a $\mathcal{N}(0,1)$ iid. prior on each element of the matrix factors $\mathbf{U}$ and $\mathbf{V}$, we allow for correlations between the entries. Thus the matrix $\mathbf{D}$ can be thought of as a precision matrix for a Gaussian prior on the space of $m + n \times m + n$ matrices. If the matrix $\mathbf{D}$ is chosen block-diagonal, then the row and column subspaces are uncorrelated and hence due to the Gaussianity independent.

This representation is very flexible, as it allows to build in prior knowledge about the low rank matrix. If we have prior knowledge about the row- and columns subspaces of the unknown matrix then this prior knowledge can be used to define the block-diagonal weight matrices $\mathbf{D}_r$ and $\mathbf{D}_c$. Since $\mathbf{D}_r^{-1}$ and $\mathbf{D}_c^{-1}$ can be thought of as covariance matrices the generalized trace norm allows to penalize variations along certain directions less than along other directions. As we will see in the experimental section, we can for example impose smooth subspaces while at the same time favoring low-rank matrices. In contrast to the standard trace norm which penalizes *all* the directions in the vector space of matrices $\{\mathbf{X} \in \mathbb{R}^{m \times n}\}$ uniformly, the generalized trace norm enables non-uniform regularization. It can be expected that with this non-isotropic penalization taking prior knowledge about the problem at hand into account the convex relaxation of the rank function will be stronger.

In summary, the generalized trace norm allows to combine more flexible subspace priors with a low-rank prior. The formulation still results in a convex problem which can be solved efficiently and globally optimal by any method of choice.

4.2. Change of Variable

If the weight matrices $\mathbf{D}_r$ and $\mathbf{D}_c$ are non-singular then the generalized trace norm is a true norm (otherwise if at least one of them is singular the generalized trace norm is only a semi-norm). In this case, the problem in Eq. (13) can be rewritten in the following way. Due to Eq. (9) we know that $\|\mathbf{W}\|_{*\mathbf{D}} = \|\mathbf{C}_r \mathbf{W} \mathbf{C}_c^T\|_*$. Let us introduce the new variable $\tilde{\mathbf{W}} = \mathbf{C}_r \mathbf{W} \mathbf{C}_c^T$ and reparameterize Eq. (13) w.r.t. this variable

$$\tilde{\mathbf{W}}_* \in \arg\min_{\tilde{\mathbf{W}}} \mathcal{L}(\mathcal{A}(\mathbf{C}_r^{-1}\tilde{\mathbf{W}}\mathbf{C}_c^{-T}) - \mathbf{z}) + \mu \left\|\tilde{\mathbf{W}}\right\|_*. \quad (14)$$

The variable we are actually interested in can be recovered by computing $\mathbf{W} = \mathbf{C}_r^{-1}\tilde{\mathbf{W}}_*\mathbf{C}_c^{-T}$. This problem formulation replaced the generalized trace norm with an ordinary trace-norm regularization. However, the data term now depends on the matrices $\mathbf{C}_r$ and $\mathbf{C}_c$. The formulation in Eq. (14) is equivalent to Eq. (13). However, the reformulation has algorithmic consequences as we will see shortly. This insight of the equivalence between the generalized trace-norm and the standard trace-norm with a reweighted data term is quite interesting in its own right. The matrix representation of the linear operator can absorb the weight matrices $\tilde{\mathcal{A}} = \mathcal{A}[\mathbf{C}_c^{-1} \otimes \mathbf{C}_r^{-1}]$ and hence we are left with the data cost $\mathcal{L}(\tilde{\mathcal{A}}(\tilde{\mathbf{W}}) - \mathbf{z})$. Recent results in compressive sensing [9, 13] are based on conditions which the linear operator appearing in the trace-norm regularized M-estimator framework has to meet in order to be able to recover the true low-rank solution with high probability. The previous derivation thus shows that the generalized trace norm can be used to modify this linear operator and hence we might expect the probability of successful low-rank matrix recovery to increase. As potential future work it might be interesting to follow this line of reasoning in order to establish quantitative statements for the generalized trace norm.

5. Application to Projective SfM

Even though the generalized trace norm is a very general concept with connections to compressive sensing, matrix completion, collaborative filtering, and general low-rank problems, we exemplify its power on a very specific computer vision problem, namely the so called Structure-from-Motion (SfM) problem. This problem deals about recovering the 3D structure from several projections of 3D points into multiple camera planes. Let the camera matrix at frame f be $\mathbf{P}_f \in \mathbb{R}^{3\times4}$ and the projective coordinates of the n^{th} feature point be $\mathbf{s}_n \in \mathbb{R}^{4\times1}$. Then the projection into the camera image plane is given by the projective equation $\lambda_{f,n}\mathbf{x}_{f,n} = \mathbf{P}_f\mathbf{s}_n$, where $\mathbf{x}_{f,n} = (u_{f,n}, v_{f,n}, 1)^T$ denotes the homogeneous coordinates of the given image observation. The goal of SfM is to recover the projective depths $\lambda_{f,n}$ from several image observations of the same

3D points at different points in time. Let F denote the total number of frames and N the total number of points. All the observations from all the points at all the frames are concisely described by introducing the stacked[2] camera matrix $\mathbf{P} = [\Downarrow_f \mathbf{P}_f] \in \mathbb{R}^{3F\times4}$, the structure matrix $\mathbf{X} = [\Rightarrow_n \mathbf{s}_n]^T \in \mathbb{R}^{N\times4}$, the combined observation matrix $\mathbf{x} = [\Downarrow_f\Rightarrow_n (u_{f,n}, v_{f,n}, 1)^T] \in \mathbb{R}^{3F\times N}$ and the combined projective depth matrix $\lambda = [\Downarrow_f\Rightarrow_n \lambda_{f,n}] \in \mathbb{R}^{F\times N}$. Then we get the matrix equation

$$[\lambda \otimes \mathbf{1}_{3\times1}] \odot \mathbf{x} = \mathbf{P}\mathbf{X}^T = \mathbf{W} \in \mathbb{R}^{3F\times N}, \quad (15)$$

where $\otimes$ denotes the Kronecker product and $\odot$ denotes the Hadamard product[3]. This problem is an instance of an affine rank minimization problem

$$\min_{\mathbf{W},\lambda} \text{rank}(\mathbf{W}) \text{ s.t. } \mathbf{W} = [\lambda \otimes \mathbf{1}_{3\times1}] \odot \mathbf{x}. \quad (16)$$

Surprisingly, the convex relaxation of this problem with the trace-norm has only recently been introduced [7]. We refer to [1, 14, 18, 8, 7] and references therein for other, mostly iterative approaches to solve this low-rank SfM problem. Here, we are interested in the behavior of the convex relaxation using the generalized trace norm. But firstly, we have to note that the formulation in Eq. (16) is ambiguous. Specifically there is a $F + N$-dimensional solution space since the camera matrices *and* the points are projective entities and hence only defined up to scale. This can easily be seen by recalling Lemma 5.1.2 in [10] which establishes the commutativity of the Hadamard product with diagonal matrices. Applied to our problem this looks like

$$[\mathbf{D}_P \otimes \mathbf{I}_3]\mathbf{P}\mathbf{X}^T\mathbf{D}_X = [[[\mathbf{D}_P \otimes \mathbf{I}_3][\lambda \otimes \mathbf{1}_{3\times1}]\mathbf{D}_X] \odot \mathbf{x}],$$

with arbitrary diagonal matrices $\mathbf{D}_P \in \mathbb{R}^{F\times F}$ and $\mathbf{D}_X \in \mathbb{R}^{N\times N}$. Applying the Kronecker product property on the right hand side we get

$$[\mathbf{D}_P \otimes \mathbf{I}_3]\mathbf{P}\mathbf{X}^T\mathbf{D}_X = [\mathbf{D}_P\lambda\mathbf{D}_X \otimes \mathbf{1}_{3\times1}] \odot \mathbf{x}, \quad (17)$$

Therefore λ is only unique up to a left- and right-multiplication with a diagonal matrix.

5.1. Avoiding Trivial Solutions

Without any further constraints, the optimal solution will lead to the trivial solution $\lambda = \mathbf{0}_{F\times N}$. In order to prevent this trivial solution, additional constraints on λ are required. In the past, the row- and column sum have been constrained while at the same time enforcing non-negative entries

$$\lambda \geq \mathbf{0}_{F\times N}, \lambda\mathbf{1}_{N\times1} = \mathbf{c} \in \mathbb{R}^{F\times1}, \lambda^T\mathbf{1}_{F\times1} = \mathbf{r} \in \mathbb{R}^{N\times1}. \quad (18)$$

---

[2]The arrows with sub-indices denote the way in which several matrices are stacked w.r.t. to each other.

[3]The Hadamard product is the element-wise multiplication between two matrices of the same size

Such a constraint set is known as the transportation polytope in the operations research community. The orientation of this transportation polytope depends upon the choice of $\mathbf{c}$ and $\mathbf{r}$. Note that the equality constraints are not independent, since the sum over the column sums has to equal the sum over the row sums and hence, the equality constraints fix $F + N - 1$ degrees of freedom. Remember that there is a $F + N$ dimensional variety of possible rank-4 solutions. The affine equality constraints of the transportation polytope slice through this $F + N$ dimensional variety and cut out a 1D variety of possible rank-4 factorizations. A proper choice of $\mathbf{c}$ and $\mathbf{r}$ is crucial for the success of the trace-norm relaxation. Usually the choice $\mathbf{c} = N\mathbf{1}_{F \times 1}$ and $\mathbf{r} = F\mathbf{1}_{N \times 1}$ is made. However, based on several experiments with synthetic data we concluded that such a choice can often result in a bad convex relaxation. Instead, we propose to replace the transportation polytope constraints Eq. (18) by the single scalar equality constraint $\mathbf{1}_{F \times 1}^T \lambda \mathbf{1}_{N \times 1} = FN \in \mathbb{R}$, which can be thought of as fixing the scale of the overall reconstruction. The intuition is that the algorithm should figure out the optimal scaling on its own. Our experiments showed that this single equality constraint does only lead to the trivial or to an almost trivial solution with a large fraction of the entries in λ set to zero for really challenging data. If such a situation is encountered additional constraints (e.g. $\lambda \geq 0.1$) can still be introduced in our formulation.

5.2. SDP Formulation for Inexact Measurements

The generalized trace norm minimization problem can be formulated as a standard SDP problem. Hence, any off-the-shelf SDP solver could be used to find the minima of the convex problem. Standard SDP solvers are usually based on second-order interior point methods which enjoy quadratic convergence near the minimum. However the cost per iteration of these methods is prohibitively high for large matrices and hence standard interior-point solvers can only be applied to small-sized problems. This is why only small-sized problems could have been solved in [7]. Due to the famous Netflix challenge, an increased interest in solvers for matrix completion problems has been observed (see [11, 12] and references therein). These solvers are based on first-order methods and do not require solving a costly linear system involving the Hessian matrix in each iteration. These solvers have solved matrix completion problems with matrix sizes up to 10000×10000 and 10^5 observed entries in about 2.5 hours [12]. The algorithm we are going to present makes use of the same underlying building blocks (especially the soft-thresholding operation) and hence also scales to thousands of image measurements. Unfortunately, the matrix completion solvers are tailored to the very specific linear equality constraints of the matrix completion problem and our problem formulation has more general linear equality constraints. Moreover, existing solvers for matrix

completion problems or robust PCA are either based on a pure L1-noise model or on a pure L2-noise model. A more suitable choice for our purposes is a robust cost function, like the Huber cost function

$$F(\mathbf{E}) = \|\mathbf{E}\|_\epsilon = \sum_{i,j} |\mathbf{E}_{i,j}|_\epsilon \text{ where } |e|_\epsilon = \begin{cases} \frac{|e|^2}{2\epsilon} & |e| < \epsilon \\ |e| - \frac{\epsilon}{2} & |e| \geq \epsilon \end{cases}.$$

Therefore, we propose to minimize the following cost function with a first-order primal-dual algorithm (see Alg. 1)

$$\min_{\mathbf{W},\lambda} \|\mathbf{W}\|_{*\mathbf{C}_r, \mathbf{C}_c} + \alpha \|\mathbf{E}\|_\epsilon \tag{19}$$

$$\text{s.t. } \mathbf{E} = \mathbf{W} - [\lambda \otimes \mathbf{1}_{d \times 1}] \odot \mathbf{x} \tag{20}$$

$$\mathbf{1}_{F \times 1}^T \lambda \mathbf{1}_{N \times 1} = FN. \tag{21}$$

The matrix-size-normalized trace-complexity measure presented in [17] suggests choosing α proportional to $\sqrt{dFN}$. Such a choice is also according to the consistency conditions derived in [4] and hence α is always chosen proportional to $\sqrt{dFN}$ in our formulation.

6. Primal-Dual Proximal Optimization

Proximal point optimization methods are based on applying proximal operators to sub-problems of the original problem at each iteration. Applied to the problem in Eq. (19) this requires the proximal operator of the generalized trace-norm. It is non trivial to derive this operator. However, we can side-step this issue by appealing to the equivalent formulation based on Eq. (14) and rather solve

$$\min_{\tilde{\mathbf{W}},\lambda} \|\tilde{\mathbf{W}}\|_* + \alpha \|\mathbf{E}\|_\epsilon \tag{22}$$

$$\text{s.t. } \mathbf{E} = \mathbf{C}_r^{-1} \tilde{\mathbf{W}} \mathbf{C}_c^{-T} - [\lambda \otimes \mathbf{1}_{d \times 1}] \odot \mathbf{x} \tag{23}$$

$$\mathbf{1}_{F \times 1}^T \lambda \mathbf{1}_{N \times 1} = FN. \tag{24}$$

This formulation only requires the proximal operator for the standard trace norm which is defined as the solution to the problem (see [12])

$$\mathbf{X}_* = \arg\min_X \tau \|\mathbf{X}\|_* + \frac{1}{2} \|\mathbf{X} - \mathbf{A}\|_F^2. \tag{25}$$

The well-known closed-form solution to this problem is the soft-thresholding given in Alg. 2. The projection operator of λ into the feasible region defined in Eq. (24) can be derived with Lagrange multipliers and is given in Alg. 3 (the Lagrange multipliers are denoted $\mathbf{y}$ in this algorithm). The proximity operator for the dual variables $\mathbf{q}$ is simply an element-wise clamping of the entries of $\mathbf{q}$ to a maximal absolute value of α. The complete algorithm is summarized in Alg. 1 and we refer the interested reader to [6] for more details about proximal splitting techniques. The equality constraints in Eq. (23) are

---

Algorithm 1: First-Order Primal-Dual Algorithm

  Data: $\mathbf{L}, \alpha, \epsilon$
  input : Linear Equality Constraint Matrix: $\mathbf{L} \in \mathbb{R}^{m \times dFN+m}$;
      Weight parameter: $\alpha \in \mathbb{R}$; Robust Huber cost function
      parameter: $\epsilon \in \mathbb{R}$; Time step sizes $\tau\sigma < \frac{1}{\|\mathbf{L}\|_{op}^2}$
  output: Minimizer $\mathbf{W}$ and λ for problem 22

---

1 // Initialization:
2 $\mathbf{q}_0 = \mathbf{0}_{m \times 1}, \lambda_0 = \mathbf{1}_{m \times 1}, \mathbf{W}_0 = \mathbf{0}_{dF \times N}$;
3 **while** *not converged* **do**
4   // Primal Variable Update:
5     $\begin{pmatrix} \text{vec}\left(\tilde{\mathbf{W}}\right) \\ \text{vec}\left(\tilde{\lambda}\right) \end{pmatrix} = \begin{pmatrix} \text{vec}\left(\mathbf{W}_t\right) \\ \text{vec}\left(\lambda_t\right) \end{pmatrix} - \tau \mathbf{L}^T \mathbf{q}_t$;
6   // Primal Variable Update:
7     $\mathbf{W}_{t+1} = \texttt{ProxTrace}(\tilde{\mathbf{W}}, \tau)$ // Alg. 2 ;
8     $\lambda_{t+1} = \texttt{ProxFeasible}(\tilde{\lambda})$ // Alg. 3 ;
9   // Reflection step:
10     $\hat{\mathbf{W}} = 2\mathbf{W}_{t+1} - \mathbf{W}_t$;
11     $\hat{\lambda} = 2\lambda_{t+1} - \lambda_t$;
12   // Dual Variable Update:
13     $\tilde{\mathbf{q}} = \frac{1}{1 + \frac{\sigma\epsilon}{\alpha}}\left(\mathbf{q}_t + \sigma\mathbf{L}\begin{pmatrix}\text{vec}\left(\hat{\mathbf{W}}\right) \\ \text{vec}\left(\hat{\lambda}\right)\end{pmatrix}\right)$;
14     $\mathbf{q}_{t+1} = \texttt{ProxBox}(\tilde{\mathbf{q}}, \alpha)$;
15     $t = t + 1$;
16 **end**

Algorithm 2: Soft-Thresholding for Trace-Norm

  Data: $\mathbf{A}, \tau$
  input : Matrix to approximate: $\mathbf{A} \in \mathbb{R}^{dF \times N}$; soft-thresholding
      parameter: τ
  output: Soft-thresholded solution for Eq. (25): $\mathbf{W}_*$

---

1 $\mathbf{U}\mathbf{\Sigma}\mathbf{V}^T = \texttt{svd}(\mathbf{A})$;
2 $\tilde{\mathbf{\Sigma}} = \texttt{max}(\mathbf{\Sigma} - \tau, 0)$;
3 $\mathbf{X}_* = \mathbf{U}\tilde{\mathbf{\Sigma}}\mathbf{V}^T$;

Algorithm 3: Projection Operator into Feasible Region

---

  input : Vector to approximate: $\mathbf{a}$; Linear equality constraint
      matrix: $\mathbf{M}$; Right-hand side of equality constraints: $\mathbf{b}$
  output: Closest point to $\mathbf{a}$ in feasible region:
      $\mathbf{x}_* = \arg\min_{\mathbf{x}} \frac{1}{2}\|\mathbf{x} - \mathbf{a}\|_2^2$ s.t. $\mathbf{A}\mathbf{x} = \mathbf{b}$.

---

1 $\begin{pmatrix} \mathbf{x}_* \\ \mathbf{y}_* \end{pmatrix} = \begin{bmatrix} \mathbf{I} & -\mathbf{M}^T \\ \mathbf{M} & \mathbf{0} \end{bmatrix}^{\dagger} \begin{pmatrix} \mathbf{a} \\ \mathbf{b} \end{pmatrix}$;

---

linear in $\tilde{\mathbf{W}}$ and λ and hence can be written in the form $\text{vec}(\mathbf{E}) = \mathbf{e} = \mathbf{L}\left(\tilde{\mathbf{w}}^T, \text{vec}(\lambda)^T\right)^T$, where $\tilde{\mathbf{w}} = \text{vec}\left(\tilde{\mathbf{W}}\right)$[4]. For our problem, the matrix $\mathbf{L} \in \mathbb{R}^{dFN \times dFN+FN}$ looks like $\mathbf{L} = \left[\mathbf{C}_c^{-1} \otimes \mathbf{C}_r^{-1}, -\text{diag}(\mathbf{x})\left[\mathbf{I}_{FN} \otimes \mathbf{1}_{d \times 1}\right]\right]$.

Performance Improvements: The pseudo-inverse of the projection operator into the feasible region stays constant throughout all the iterations and could be pre-computed once and for all. However, even though the matrix is highly sparse, the matrix generally becomes dense

after taking its pseudo-inverse which poses a problem for large problem sets. Hence, instead of pre-computing the pseudo-inverse we rather pre-compute a sparse LU-decomposition which is guaranteed to stay sparse. Then in each iteration, a sparse LU-problem needs to be solved which can be done very efficiently. If the data is incomplete, i.e. if there are missing entries in our input data matrix $\mathbf{x}$ (maybe due to occlusions or tracking failure), there is no need to introduce a full F-by-N matrix for the unknown projective depths. We only need to keep track of those projective depths $\lambda_{f,n}$ which corresponds to visible points. Therefore we put all these entries in one single m-vector where m denotes the number of observations. For simplicity, this vector is still called λ in Alg. 1. In this case, the linear equality constraint matrix $\mathbf{L}$ of course only consists of those rows corresponding to observed points. There is one drawback of the formulation in Eq. (22) compared to Eq. (19). The equality constraints in Eq. (20) would lead to a highly sparse matrix $\mathbf{L}$, the equality constraint in Eq. (23) however lead to a dense first block in $\mathbf{L}$ due to the Kronecker-product $\mathbf{C}_c^{-1} \otimes \mathbf{C}_r^{-1}$. For large datasets this matrix can become huge. Luckily, there is no need to store the Kronecker-product: instead the Kronecker-product property $\text{vec}\left(\mathbf{C}_r^{-1}\tilde{\mathbf{W}}\mathbf{C}_c^{-T}\right) = \left[\mathbf{C}_c^{-1} \otimes \mathbf{C}_r^{-1}\right]\tilde{\mathbf{w}}$ is used in step 5 and 13 of Alg. 1.

Additional Insights: First order proximal methods are known to converge quickly at the beginning. When using L1-error terms however, these methods often oscillate around the true optimum and only slowly converge towards this optimum due to the non-differentiability of the L1-norm at 0. The robust Huber cost function replaces the L1-error around the origin with a L2-penalty which is differentiable. Hence, the Huber cost function is not only a more appropriate model for outliers and inliers contaminated by Gaussian noise it also leads to less oscillations and hence faster convergence.

7. Experiments

Synthetic Experiments The robustness of factorization approaches, and actually of any SfM approach, is known to largely depend on the motion. With a larger variation in camera poses, the robustness of SfM method increases at the price of a more difficult feature point matching stage. In our synthetic experiments, we have generated realistic projections of feature points based on parameters of a real camera. The camera internal calibration matrix is generated with zero skew, square pixels, principal point in the center of the image, a focal length of 10mm, and a resolution of 1920-by-1080. In order to generate smooth camera motions, the external calibration matrices are based on spline interpolated camera rotations and translations. The mean of the camera centers is chosen at $(8, 0, 0)$ an varies smoothly along a spline curve with 10 uniformly sampled control

---

[4]$\text{vec}(\mathbf{A})$ denotes the vectorization of a matrix into a vector by stacking all its columns below each other

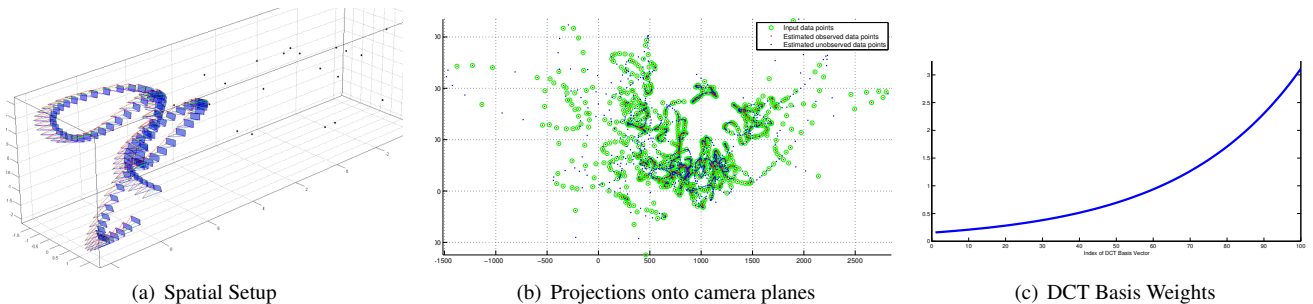| (a) Spatial Setup | (b) Projections onto camera planes | (c) DCT Basis Weights |

Figure 1. Synthetic experiment: In Fig. 1(a) the spatial configuration of the cameras and points is shown whereas Fig. 1(b) shows the observed input data (green) and the reprojection of the estimated points (in red (blue) projections corresponding to observed (missing) entries) for all the frames. The smooth trajectories are clearly visible. Our approach achieves an error with a sub-pixel accuracy of 0.2 pixels. Fig. 1(c) shows the weighting of the DCT basis components. The matrix $\mathbf{C}_r$ is chosen such that the columns of $\mathbf{W}$ (which correspond to 3D point trajectories) are transformed into a three-fold Direct Cosine Transform (DCT) basis (one DCT basis for each (x,y,z) component). Motivated by the fact that natural data follows a power law decay rate in a DCT basis (e.g. also used for jpg-compression) we have chosen an exponential weighting scheme for the DCT basis vectors. As shown in the image, smooth directions in the DCT are penalized less heavily than non-smooth ones.

points in $[6, 10] \times [-2, 2] \times [-2, 2]$ throughout the motion. The cameras are oriented such that the principal axis always points to the origin and rotations around the principal axis are limited to ± 25 degrees again smoothed with a spline interpolation. 100 camera poses have been sampled from the resulting spline interpolation curves. 20 points with a uniform distribution in a cube of size $[-5, 5] \times [2, 2] \times [2, 2]$ are generated in order to simulate realistic variations in depth. Fig. 1(a) shows the resulting spatial setup. Fifty percent of the entries in the resulting data matrix are marked uniformly at random as missing. No noise is added in order to facilitate the search for optimal parameters of the Huber cost function ϵ (ϵ is fixed to zero) and of the data fidelity trade-off α. The standard projective factorization approach with ordinary trace norm regularization (which has also been presented in [7]) failed to recover the underlying structure for whatever choice of parameter value α. The smallest reprojection error we could achieve was 27.5 pixels. Note that the parameter α can always be chosen such that the resulting SDP formulation will indeed provide a rank-4 solution. Our experiments however show that such a low-rank solution is not equal to the true underlying low rank matrix: in order to get a rank-4 solution the trace norm regularization term has to be large enough compared to the data term which on the other hand implies a stronger regularization of the unknown matrix *uniformly* in all directions. This uniform penalization property is exactly the drawback of the standard trace norm. Either we choose a small α and get a highly regularized rank-4 matrix or we choose a larger α which leads to a full rank solution to the SDP problem. In either way, the solution of the convex relaxation is not the one we are looking for. Adding weight matrices $\mathbf{C}_r$ and $\mathbf{C}_c$ introduces *non-uniform* regularization, i.e. not all directions are penalized in the same way. We have

chosen the matrix $\mathbf{C}_r$ such that smooth directions are penalized less heavily than non-smooth ones (see Fig. 1(c)). The matrix $\mathbf{C}_c$ down-weights the direction $\mathbf{1}_{1 \times N}$ by a factor of roughly 4 since we already know that this direction should be present in our solution (this direction corresponds to the properly scaled homogeneous coordinates of the points, i.e. the last row of $\mathbf{X}^T$). With such a choice the generalized trace norm formulation successfully recovered the structure up to a sub-pixel reprojection error of 0.2 pixels (see Fig. 1(b) for a visual depiction of the measured input points and the recovered estimates). Note that for a fair comparison, all the reprojection errors we state are indeed based on the best rank-4 factorization in the Frobenius-norm sense of the SDP solution and *not* directly on the potentially full-rank SDP solution. This is important since especially for the standard trace norm regularized solutions, there might be quite a difference between the SDP solution and its best rank-4 approximation. Running the very same experiment with Gaussian noise with a standard deviation of 0.5 (2) pixels changed the average reprojection error of our method with the choice $\epsilon = 0.5$ ($\epsilon = 2$) for the Huber cost threshold to 0.6 (2.5) pixels. The algorithm is implemented in Matlab (single-threaded) on a Core i7 740QM and needs $0.023s$ per iteration, the algorithm converged after 1100 iterations for a 600×50 matrix with 50% missing entries. This is orders of magnitude faster than interior point solvers.

Real World Data As in [7], the Dinosaur Oxford data sequence has been used for verification purposes on a real-world data sequence. Only those points which were visible in at least 10 frames have been considered. This resulted in a sequence with $F = 17$ frames and $N = 47$ points with roughly 30 percent of missing entries. Exactly the same weight matrices as for the synthetic data experiment have been chosen. The robust Huber parameter has been set to
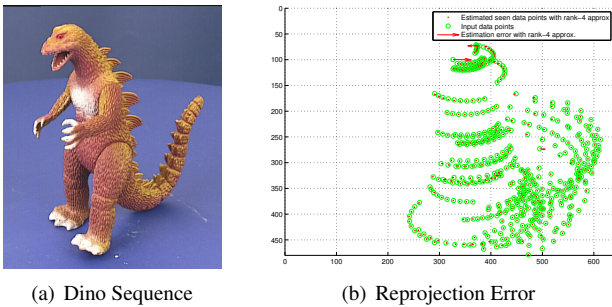
(a) Dino Sequence   (b) Reprojection Error

Figure 2. For the dinosaur sequence, our approach leads to an average reprojection error of 0.95 pixels which is a reduction by a factor of two compared to the standard trace-norm regularized formulation.

a 2 pixel threshold. Our method achieves an average reprojection error of 0.95 pixels whereas the standard trace norm regularization achieves a reprojection error of 1.84 pixels. Hence, we get a reduction by a factor of two (see also Fig. 2). Note that the full dinosaur sequence has a very strong band diagonal pattern of known entries. Such a pattern proved to be very challenging for the SDP relaxations we considered in this paper. We leave it as an open question for future work to address the issue of strong band-diagonal patterns of known entries.

8. Conclusion and Future Work

This paper presented a generalized trace norm which allows to encode prior knowledge about a specific problem at hand into a convex regularization term which enforces a low rank solution while at the same tame taking the specific problem structure into account. We have exemplified the generalized trace norm on the projective SfM factorization problem where it was shown that the standard trace norm regularization failed whereas the generalized trace norm succeeds in finding the correct rank-4 solution. As expected the generalized trace norm with its built-in prior knowledge can indeed handle more missing entries. From an algorithmic point of view, an efficient first-order algorithm for generalized trace-norm regularized problems has been presented which can handle considerably larger problems than previous work. In future work, we would like to perform more extensive quantitative evaluations based on varying noise levels, fraction of missing entries, and outliers. Recently, new results for M-estimators of the form in Eq. (12) with trace-norm regularization have been derived [13, 9] which tell that if the operator $\mathcal{A}$ fulfills certain conditions then the trace-norm regularized convex problem will provide the correct low-rank solution with high probability. We plan to investigate whether similar results can be stated for the operator appearing in the projective SfM formulation with a generalized trace norm regularization.

References

[1] H. Aanæs, R. Fisker, K. Åström, and J. Carstensen. Robust factorization. *IEEE Trans. PAMI*, pages 1215–1225, 2002.

[2] R. Angst, C. Zach, and M. Pollefeys. Supplemental material. www.inf.ethz.ch/personal/rangst/publications, 2011.

[3] F. Bach, J. Mairal, and J. Ponce. Convex sparse matrix factorizations. Technical report, HAL-00345747, 2008.

[4] F. R. Bach. Consistency of trace norm minimization. *J. Mach. Learn. Res.*, 9:1019–1048, June 2008.

[5] E. J. Candes. The power of convex relaxation. In *21st Annual ACM-SIAM*, SODA'10, pages 1321–1321, 2010.

[6] A. Chambolle and T. Pock. A First-Order Primal-Dual Algorithm for Convex Problems with Applications to Imaging. *Journal of Math. Imaging and Vision*, pages 1–26, 2010.

[7] Y. Dai, H. Li, and M. He. Element-wise factorization for n-view projective reconstruction. In *Proc. of ECCV*, pages 396–409, 2010.

[8] A. Eriksson and A. van den Hengel. Efficient computation of robust low-rank matrix approximations in the presence of missing data using the l1 norm. In *IEEE Conf. on CVPR*, pages 771–778, June 2010.

[9] D. Gross. Recovering low-rank matrices from few coefficients in any basis. *arXiv*, Oct 2009.

[10] R. A. Horn and C. R. Johnson. *Topics in matrix analysis*. Cambridge University Press, Cambridge, 1991.

[11] Z. Lin, M. Chen, L. Wu, and Y. Ma. The Augmented Lagrange Multiplier Method for Exact Recovery of Corrupted Low-Rank Matrices. *ArXiv e-prints*, Sept. 2010.

[12] R. Mazumder, T. Hastie, and R. Tibshirani. Spectral regularization algorithms for learning large incomplete matrices. *J. Mach. Learn. Res.*, 11:2287–2322, August 2010.

[13] S. Negahban, P. Ravikumar, M. J. Wainwright, and B. Yu. A unified framework for high-dimensional analysis of m-estimators with decomposable regularizers. *CoRR*, 2010.

[14] J. Oliensis and R. Hartley. Iterative extensions of the sturm/triggs algorithm: Convergence and nonconvergence. *IEEE Trans. PAMI*, 29:2217–2233, Dec. 2007.

[15] B. Recht, M. Fazel, and P. A. Parrilo. Guaranteed Minimum-Rank solutions of linear matrix equations via nuclear norm minimization. *ArXiv e-prints*, June 2007.

[16] R. Salakhutdinov and A. Mnih. Bayesian probabilistic matrix factorization using markov chain monte carlo. In *ICML*, pages 880–887, New York, NY, USA, 2008. ACM.

[17] R. Salakhutdinov and N. Srebro. Collaborative filtering in a non-uniform world: Learning with the weighted trace norm. *CoRR*, abs/1002.2780, 2010.

[18] A. Zaharescu and R. P. Horaud. Robust factorization methods using a gaussian/uniform mixture model. *IJCV*, 81(3):240–258, March 2009.