

# Compressive Higher-order Sparse and Low-Rank Acquisition with a Hyperspectral Light Stage

Boris Ajdin<sup>\*</sup>, Manuel Finckh<sup>\*</sup>, Christian Fuchs<sup>\*</sup>,  
Johannes Hanika<sup>†</sup>, Hendrik P.A. Lensch<sup>\*</sup>

<sup>\*</sup> - Eberhard Karls Universität Tübingen

<sup>†</sup> - Weta Digital

Technical report WSI-2012-01

ISBN 0946-3852

Eberhard Karls Universität Tübingen

FB Informatik

Computergrafik

Sand 14

72076 Tübingen

# Compressive Higher-order Sparse and Low-Rank Acquisition with a Hyperspectral Light Stage

Boris Ajdin\*, Manuel Finckh\*, Christian Fuchs\*, Johannes Hanika†, Hendrik P.A. Lensch\*

\* - Eberhard Karls Universität Tübingen

† - Weta Digital



**Figure 1:** The LED-based hyperspectral light stage can emit a different controlled spectrum from each light source direction simultaneously enabling compressive bispectral measurements of reflectance fields. One slice showing different illumination wavelengths and directions is shown on the right for a set of text markers. The compressive sparse and low-rank reconstruction approach robustly estimates a low-rank approximation plus a sparse representation for the non-low-rank content for 14 wavelength bands and 196 directions from only 400 captured samples.

## Abstract

Compressive sparse and low-rank recovery (CSLR) is a novel method for compressed sensing deriving a low-rank and a sparse data terms from randomized projection measurements. While previous approaches either applied compressive measurements to phenomena assumed to be sparse or explicitly assume and measure low-rank approximations, CSLR is inherently robust if any such assumption might be violated. In this paper, we will derive CSLR using Fixed-Point Continuation algorithms, and extend this approach in order to exploit the correlation in high-order dimensions to further reduce the number of captured samples. Though generally applicable, we demonstrate the effectiveness of our approach on data sets captured with a novel hyperspectral light stage that can emit a distinct spectrum from each of the 196 light source directions enabling bispectral measurements of reflectance from arbitrary view-points. Bispectral reflectance fields and BTFs are faithfully reconstructed from a small number of compressed measurements.

**CR Categories:** I.4.1 [Image Processing and Computer Vision]: Digitization and Image Capture—Reflectance I.4.1 [Image Processing and Computer Vision]: Digitization and Image Capture—Sampling I.4.1 [Image Processing and Computer Vision]: Digitization and Image Capture—Scanning

**Keywords:** reflectance field, multispectral acquisition, multispectral rendering, compressed sensing, hyperspectral light stage, image based relighting, sparse and low rank decomposition

## 1 Introduction

In order to create images of a real-world scene lit using altered illumination one needs to capture its reflectance field. This operation, called image-based relighting, can be mathematically defined as follows [Peers et al. 2009; Sen and Darabi 2009]:

$$c = T \cdot I, \quad (1)$$

where  $T$  is the  $m \times n$  matrix which describes the light transport from  $n$  light sources of the scene to  $m$  camera pixels,  $I$  is the vector of light source radiances of length  $n$ , and  $c$  is a vector of camera pixels whose length is  $m$  (the outgoing radiances of the scene). Capturing reflectance fields is a complex problem, and in order to simplify it a common assumption is made that no fluorescent interactions are present in the scene — wavelengths of all reflected light rays have to match the wavelengths of incoming rays. Such an assumption is not valid for many natural scenes, as was observed by Hullin et al. [2010], and it is therefore important to consider capturing multi- or bispectral reflectance.

In this paper we will present an LED-based capturing device that allows for controlled hyperspectral illumination — rather than a single wavelength each light source can emit its own spectrum — and arbitrary selection of the camera view-point. Such devices introduce a high-dimensional sampling and reconstruction problem.

Previous approaches include compressed sensing and low-rank matrix approximations. Compressed sensing [Sen and Darabi 2009; Peers et al. 2009] is able to recover the signal using sparse sampling, but it requires the signal to be sparse or compressible in some basis. Furthermore, these approaches do not recognize correlations between different dimensions. They have been extended to address multi-dimensional correlations [Baron et al. 2009; Duarte and Baraniuk 2010], but the sparseness or compressibility requirement remains.

On the other hand, low-rank matrix approximations of the light transport operator [O’Toole and Kutulakos 2010; Wang et al. 2009] requires that the desired signal be low-rank. These approaches are based on symmetric setups, which are not always possible to build. Hullin et al. [2010] presents a two-stage capture system which first identifies the low-rank sub-space by performing dense sampling, after which only sparse measurements are needed for the remaining samples.

To circumvent problems present in previous approaches, we choose to reconstruct the light transport using both sparse *and* low-rank components from compressed/projected measurements. Sparse and low-rank decomposition from complete data has been well stud-

ied [Yuan and Yang 2009; Chandrasekaran et al. 2009a], and in [Waters et al. 2011] developed in parallel to our work a solution for the compressed sensing setting has been proposed. Our method differs from this work in that we do not require prior knowledge of the rank and sparsity of matrices we are recovering. Furthermore, we extend our approach to higher-order problems. This way, one can exploit correlation between different dimensions, e.g. wavelength, lighting and view directions and achieve significant savings in acquisition cost.

The main contributions of this paper are:

- a fully automated hyperspectral light stage which allows for mixed-spectra projected measurements;
- a novel algorithm for simultaneous sparse and low-rank matrix decomposition from compressed/projected measurements (CSLR);
- a higher-order extension of the proposed algorithm which is able to utilize the multi-dimensional signal correlations (HO-CSLR), resulting in even fewer images required for high-quality reconstruction;
- comparisons of the proposed method to either sparse or low-rank reconstructions, which shows clear improvement of the obtained results.

Note that both algorithms presented can be applied beyond the application proposed in this paper.

## 2 Previous Work

**Measuring Bispectral Reflectance.** A number of proposed devices providing illumination from a sphere or a hemisphere for capturing reflectance fields range from moving point light sources [Debevec et al. 2000; Müller et al. 2005b] to domes with a set of discretized light directions [Wenger et al. 2005; Müller et al. 2005a] to linear light sources [Gardner et al. 2003], to the use of curved reflectors and projectors to allow for illuminating with illumination basis functions over the hemisphere rather than individual directions [Dana and Wang 2004; Ghosh et al. 2010].

The simplified multispectral BTF measurement setup presented in [Rump and Klein 2010] attempts to deal with metamerism of reflectance measured with standard RGB devices. Sparse spectral data at individual points is captured together with dense RGB data. By using the additional spectral information they are able to correctly infer the color of the RGB sampled data. However, in order for this method to work the spectra of all points exhibiting metamerism have to be measured during scanning, which may be a non-trivial, expensive task in certain conditions.

Hullin et al. [2010] have applied narrow-band wavelength filters in front of the camera and in front of a white light source to measure view and illumination dependent bispectral reflection and reradiation matrices for homogeneous materials, with the primary goal of capturing the effect of fluorescence.

Several methods have been proposed which utilize compressed sensing in hyperspectral imaging ([Sun and Kelly 2009; Liu et al. 2010; Valiollahzadeh and Yin 2010]). The main differences to our work is that most of these methods exploit coherence in the spatio-spectral domain while our examples so far explicitly explore the angle-spectral domain. Furthermore, our approach acquires bispectral data and is therefore able to capture fluorescence, i.e. we can resolve reflected wavelength depending on incident wavelength rather than just capturing a multispectral image.

Our goal is to utilize the novel hyperspectral light stage design to obtain dense illumination spectrum measurements from all illumination directions. We extend the setup of Debevec et al. [2002] to include 14 different narrow band LEDs per illumination direction to cover the entire visible spectrum. Together with a motorized RGB camera and a turntable supporting the object in the center of the dome, we can capture bispectral and spatially varying reflectance data from arbitrary viewing directions, i.e. depending on incident and reflected radiance.

**Multiplexed Illumination.** Illumination multiplexing [Schechner et al. 2003] aims at improving the SNR over scanning approaches, where the contribution of each light source is captured in a single measurement, by always using a pseudo-random set of light sources in each capture, followed by an inversion of the measurement matrix. As the contribution of each light source is now observed in a set of measurements the SNR typically improves. The paper proposes the use of Hadamard codes, since they are optimal under certain conditions. A deeper analysis of illumination multiplexing in presence of noise and pixel saturation [Ratner and Schechner 2007] suggests that this problem is more complex than previously thought. If photon noise is negligible compared to camera signal-independent noise, then Hadamard codes indeed are the optimal solution. Otherwise using Hadamard codes can even result in a decreased measurement SNR [Wenger et al. 2005]. In order to fight the problem of potential pixel saturation when illuminating with multiple lights Ratner et al. [2007] propose novel codes keeping the measured reflectance below the saturation level.

For multispectral reflectance Park et al. [2007] derive a method which determines the optimal multiplexing sequence for a set of sources with known spectra and the given number of measurements to be taken. A simple empirical linear multispectral model is used to reconstruct the spectral data at each pixel position from the given set of input images with multiplexed multispectral information.

Unfortunately, the common problem of quantization within the sensor was not explicitly discussed in the past, namely the issue of multiplexing very dim light sources with very bright ones — in this case the reflection from dim light sources might be hard to reconstruct since their values might be too close to the noise floor of bright light reflection. Codes derived using some kind of optimization to maximize SNR (e.g. codes from [Park et al. 2007]) could in principle implicitly counter that problem.

**Compressive Sensing and Low-rank Approximations.** While multiplexed illumination only tries to improve the SNR keeping the same number of images as in the case of scanning, compressive sensing [Candès et al. 2006; Donoho 2006; Tsaig and Donoho 2006] is a method for recovery (most often using the  $\ell_1$ -norm) from sparsely measured signals, assuming the signal to be reconstructed is sparse or compressible in some basis (e.g. in Fourier space for sound or using wavelets for images). This work is a probabilistic relaxation of the Nyquist-Shannon sampling theorem — assuming the sparsity of the signal is high enough, the exact signal can be recovered with high probability.

Baron et al. [2009] extend the compressive sensing framework utilizing multiple sensing elements. They infer that multiple signals can be recovered from sparse measurements using *both* inter- and intra-signal sparsity (i.e. sparsity of each signal separately as well as the joint sparsity of the signal ensemble). Duarte and Baraniuk [Duarte and Baraniuk 2010] extend compressed sensing to multi-dimensional scenarios using the *Kronecker* product. When measurement matrices and the sparsifying basis for lower-dimensional  $d$ -sections are known, the Kronecker product is applied to obtain the combined measurement matrix and sparsifying basis. Such combined structures contain enough information to re-

cover the desired multi-dimensional signal. The use of the Kronecker product however hinders the exploitation of potential correlation across different dimensions.

A problem related to compressive sensing is matrix completion [Candès and Recht 2009; Candès and Plan 2009]. It tries to recover a full matrix from a sparse set of measured entries, assuming the resulting matrix is low-rank. More generally, matrix rank minimization [Goldfarb and Ma 2009; Ma et al. 2009] attempts to recover the matrix assuming the measured samples have been obtained after some linear mapping has been applied to the matrix being recovered.

Sparse and low-rank matrix decomposition [Yuan and Yang 2009; Chandrasekaran et al. 2009a; Chandrasekaran et al. 2009b] tries to recover sparse and low-rank components of a known matrix concurrently via convex optimization. The process implements a robust PCA where outliers which would disturb the low-rank assumption are effectively represented in a sparse matrix. Yuan and Yang [2009] use Fixed Point Continuation methods to minimize  $\ell_1$  and nuclear matrix norms for the sparse and the low-rank component respectively, while Chandrasekaran et al. [2009b] reduce the problem of finding such matrices to a semidefinite program. Conditions which have to be satisfied in order for the decomposition to be exact are studied in detail in [Chandrasekaran et al. 2009a].

Concurrently to our work Waters et al. [2011] derived a sparse and low-rank factorization algorithm from randomized projections, i.e. in a compressed sensing setting. There are two major differences between this work and ours: in their work an assumption about the sparsity and rank of the matrices being recovered needs to be made to drive the *K-term support* algorithm for the sparse component and the *svd* computation for the low-rank component, whereas we do not need to make such assumptions; furthermore, the factorization is computed on the entire image at once, which adds correlation between neighboring pixels, but can increase the rank of the low-rank component in certain cases — we compute the factorization per-pixel which guarantees the lowest possible rank of the obtained low-rank component, at the cost of increased computation time. In addition, we will derive a sparse and low-rank factorization algorithm for higher-order tensors that incorporate correlations across various problem dimension, drastically reducing the required input size.

**Compressive or Low-rank Light Transport.** Compressive sensing and the recovery of a low-rank representation have been thoroughly investigated in the context of reflectance acquisition. Compressive light transport sensing [Peers et al. 2009; Sen and Darabi 2009] has been proposed with different illumination patterns. While Peers et al. [2009] represent the patterns in a compression basis (Haar wavelets), Sen and Darabi [2009] use a more standard approach and project only Bernoulli patterns.

The typically encountered low-rankness of the light transport operator has been explored in the Kernel Nyström method [Wang et al. 2009] which recovers the light transport matrix by sampling a small number of rows and columns of the matrix filling the remaining entries, assuming the resulting matrix has small rank. This method incorporates nonlinear coherence of the light transport matrix rows/columns. A non-linear mapping of intensities, the so-called “kernel-trick”, is applied to reduce the effect of non-low-rank effects such as shadows and specular highlights. The method requires a symmetric, co-axial setup of cameras and projectors probing the light transport in its original and transposed form. A symmetric setup is also featured by O’Toole and Kutoulakos [2010] who implement Krylov subspace and Arnoldi methods for obtaining the first eigenvectors partially in the optical domain. While both approaches significantly reduce the number of required images for a

reasonable reconstruction they require the symmetric measurement setup which is complicated to build. Furthermore, some of their savings are due to restricting the illumination to be low-rank already, i.e. a projector is illuminating the scene through a diffuser plate. Our reconstruction method does not require a symmetric measurement setup and does not constrain the illumination.

### 3 Compressive Sparse and Low-rank Recovery

In this section we derive mathematical equations necessary for our decomposition approach. Based on simultaneous sparse and low-rank decomposition [Yuan and Yang 2009; Chandrasekaran et al. 2009b] for known matrices we extend this concept towards compressed sensing where only the projection of the matrix is sampled. In order to perform the reconstruction we adapt two algorithms to jointly work together: the Fixed Point Continuation (FPC) algorithms for  $\ell_1$ -regularized minimization [Hale et al. 2007] and the Unconstrained Nuclear Norm Minimization (UNNM) [Ma et al. 2009], used for reconstructing sparse or low-rank matrix components respectively. In Section 4 we will then extend our approach to deal with higher-order tensors.

#### 3.1 Simultaneous Sparse and Low-rank Decomposition

Let us describe simultaneous sparse and low-rank matrix decomposition as proposed by Yuan and Yang [2009]. The envisioned decomposition can be obtained with the following constrained minimization to which the augmented Lagrangian method with alternating directions is applied:

$$\begin{aligned} \min_{A,B} \quad & \gamma \|A\|_{\ell_1} + \|B\|_* \\ \text{s.t.} \quad & A + B = C, \end{aligned} \quad (2)$$

where matrix  $A$  is sparse and  $B$  is low-rank,  $C$  corresponds to the input matrix,  $\|\cdot\|_{\ell_1}$  is the  $\ell_1$  norm of the matrix defined by the sum of absolute values of all entries,  $\|\cdot\|_*$  is the nuclear norm defined by the sum of all singular values.  $\gamma > 0$  is a constant providing a trade-off between sparse and low-rank minimization. In their work, only decomposition of a known matrix  $C$  is discussed. Our problem is slightly different — instead of a full matrix we have only a *sparse* set of measurements

$$\begin{aligned} \min_{A,B} \quad & \gamma \|A\|_{\ell_1} + \|B\|_* \\ \text{s.t.} \quad & M(A + B) = C, \end{aligned} \quad (3)$$

where  $M$  is the measurement matrix fulfilling the restricted isometry constraint, e.g. a random 0 or 1 Bernoulli matrix,  $C$  is the matrix of projected measurements, and  $A, B$  are sparse and low-rank components of the light transport matrix  $T$  which is to be reconstructed.

Following Yuan and Yang [2009], we derive the augmented Lagrangian equation for Problem 3 with Lagrangian multiplier  $Z$ :

$$\begin{aligned} L(A, B, Z) = & \gamma \|A\|_{\ell_1} + \|B\|_* - \langle Z, M(A + B) - C \rangle \\ & + \frac{\beta}{2} \|M(A + B) - C\|_F^2 \end{aligned} \quad (4)$$

In the above equation,  $\langle \cdot, \cdot \rangle$  is the standard trace inner product and  $\|\cdot\|_F$  is the Frobenius norm. The parameter  $\beta > 0$  penalizes the deviation from the given constraint. Minimizing this equation optimizes Problem 3.



Once the augmented Lagrangian is set up we can derive an iterative procedure which solves for  $A$ ,  $B$  and  $Z$  in a serial fashion [Yuan and Yang 2009]. Sequential updates for  $A$  and  $B$  are determined using the partial derivative of the augmented Lagrangian, while the update for  $Z$  can be derived from the classical iterative solution to the augmented Lagrangian method [Nocedal and Wright 2006] incorporating the subgradients of the corresponding norms. For real matrices  $M$  the minimum is found by utilizing the following update steps:

$$\begin{aligned} \mathbf{0} &\in \gamma \partial(\|A^{k+1}\|_{\ell_1}) - M^T(Z^k - \beta(M(A^{k+1} + B^k) - C)) \\ \mathbf{0} &\in \partial(\|B^{k+1}\|_*) - M^T(Z^k - \beta(M(A^{k+1} + B^{k+1}) - C)) \quad (5) \\ Z^{k+1} &= Z^k - \beta(M(A^{k+1} + B^{k+1}) - C) \end{aligned}$$

It turns out that these optimality conditions for matrices  $A$  and  $B$  are known problems, and iterative schemes have been proposed for each problem individually. Namely, the optimal matrix  $A$  can be found using the FPC method for  $\ell_1$ -regularized minimization, while the optimal matrix  $B$  can be found using the FPC method for UNNM. The two minimization problems feature almost the same structure. Using the augmented Lagrangian approach proposed in Eq. 5 they are combined and  $A$  and  $B$  can be jointly optimized.

### 3.2 FPC for $\ell_1$ -regularized Minimization

Assuming  $B^k$  and  $Z^k$  are fixed we can easily derive the explicit update step for obtaining a sparse matrix  $A$  using the Fixed Point Continuation method for  $\ell_1$ -regularized minimization [Hale et al. 2007], which attempts to solve the following problem

$$\min_X \|X\|_{\ell_1} \quad \text{s.t. } \mathcal{M}(X) = N. \quad (6)$$

Here  $\mathcal{M}$  is a linear map and  $N$  is a matrix. In the presence of noise in  $N$  the constraint  $\mathcal{M}(X) = N$  must be relaxed. One way to achieve this is to use the Lagrangian form:

$$\min_X \mu \|X\|_{\ell_1} + \frac{1}{2} \|\mathcal{M}(X) - N\|_F^2. \quad (7)$$

Let the set of all optimal solutions of Formula 7 be  $\mathcal{X}$ . If  $X^* \in \mathcal{X}$  then  $X^*$  must satisfy the following optimality condition which has the same form as in Eq. 5.

$$\mathbf{0} \in \mu \partial\|X^*\|_{\ell_1} + \mathcal{M}^T(\mathcal{M}(X^*) - N). \quad (8)$$

For any  $\tau > 0$ , Equation 8 is equivalent to

$$\mathbf{0} \in \tau \mu \partial\|X^*\|_{\ell_1} + \tau \mathcal{M}^T(\mathcal{M}(X^*) - N). \quad (9)$$

Using substitution we get an equivalent optimization problem

$$\begin{aligned} \text{Let } Y &= X^* - \tau \mathcal{M}^T(\mathcal{M}(X^*) - N) \implies \\ \mathbf{0} &\in \tau \mu \partial\|X^*\|_{\ell_1} + X^* - Y \implies \quad (10) \\ X^* &= \min_X \tau \mu \|X\|_{\ell_1} + \frac{1}{2} \|X - Y\|_F^2. \end{aligned}$$

This optimization step has a closed form optimal solution, obtained using the  $\ell_1$ -shrinkage operator  $S_v^{\ell_1}$  (see [Hale et al. 2007])

$$X^* = S_v^{\ell_1}(Y) \quad (11)$$

defined for matrices component-wise as

$$S_v^{\ell_1}(X) = \text{sign}(X_{ij}) \max\{|X_{ij}| - v, 0\}. \quad (12)$$

Using substitution 10 and the closed form solution 11 an iterative algorithm for computing  $X^*$  can be written as

$$\begin{aligned} Y^k &= X^k - \tau \mathcal{M}^T(\mathcal{M}(X^k) - N) \\ X^{k+1} &= S_{\tau \mu}(Y^k). \end{aligned} \quad (13)$$

The update step for matrix  $A^{k+1}$  from Equation 5 can be rewritten to obtain the following form

$$\mathbf{0} \in \gamma \partial(\|A^{k+1}\|_{\ell_1}) + \beta M^T(MA^{k+1} + MB^k - C - \frac{1}{\beta} Z^k) \quad (14)$$

Putting  $N = -(MB^k - C) + \frac{1}{\beta} Z^k$  and multiplying the equation with  $\frac{\tau}{\beta}$  we obtain the form equivalent to Equation 9. Finally, using equations 13 and 14 we can derive the formula for computing the matrix update  $A^{k+1}$  for our problem if matrices  $B^k$  and  $Z^k$  are fixed:

$$\begin{aligned} Y_A^k &= A^k + \tau M^T(Z^k - \beta(M(A^k + B^k) - C)) \\ A^{k+1} &= S_{\frac{\tau \mu}{\beta}}(Y_A^k). \end{aligned} \quad (15)$$

### 3.3 FPC for UNNM

Once  $A$  has been updated, assuming  $A^{k+1}$  and  $Z^k$  are fixed, an explicit minimization step for  $B$  can be derived, which uses the FPC algorithm for unconstrained nuclear norm minimization [Ma et al. 2009]. FPC for UNNM attempts to solve the following minimization problem:

$$\min_X \|X\|_* \quad \text{s.t. } \mathcal{M}(X) = N. \quad (16)$$

Notice that Formulas 6 and 16 differ only in the choice of the norm of matrix  $X$ . Therefore the optimal solution  $X^*$  for the UNNM case can be found similarly to the previously described  $\ell_1$  case as:

$$X^* = \min_X \tau \mu \|X\|_* + \frac{1}{2} \|X - Y\|_F^2 \quad (17)$$

For the matrix norm minimization the closed form optimum is obtained using the *matrix shrinkage operator*  $S_v^*$

$$X^* = S_v^*(Y). \quad (18)$$

The operator applied to a matrix  $X$  whose SVD is given as  $X = U \text{Diag}(\sigma) V^T$  with singular values  $\sigma_i$  is defined as:

$$\begin{aligned} S_v(X) &= U \text{Diag}(\bar{\sigma}) V^T \\ \bar{\sigma}_i &= \begin{cases} \sigma_i - v, & \text{if } \sigma_i - v > 0; \\ 0, & \text{otherwise.} \end{cases} \end{aligned} \quad (19)$$

Applied to our problem,  $B^k$  is updated in each iteration, assuming fixed  $A^{k+1}$  and  $Z^k$ , as follows:

$$\begin{aligned} Y_B^k &= B^k + \tau M^T(Z^k - \beta(M(A^{k+1} + B^k) - C)) \\ B^{k+1} &= S_{\frac{\tau \mu}{\beta}}(Y_B^k) \end{aligned} \quad (20)$$

The complete algorithm for joint recovery of a sparse matrix  $A$  and a norm-minimized matrix  $B$  is presented in Figure 2.

```

input: measurements  $C$ , projection matrix  $M$ 
output:  $\ell_1$ -minimized  $A$ ,
        nuclear norm minimized  $B$  } s.t.  $M(A+B) = C$ 

initialize  $A^0, B^0, Z^0 \leftarrow 0$ 
do
     $Y_A = A^k + M^T(Z^k - \beta(M(A^k + B^k) - C))$ 
     $A^{k+1} = S_{\frac{\gamma}{\beta}}^{\ell_1}(Y_A)$ 

     $Y_B = B^k + M^T(Z^k - \beta(M(A^{k+1} + B^k) - C))$ 
     $B^{k+1} = S_{\frac{\gamma}{\beta}}^*(Y_B)$ 

     $Z^{k+1} = Z^k - \beta(M(A^{k+1} + B^{k+1}) - C)$ 

until convergence

```

**Figure 2:** CSLR: Algorithm for compressive sparse and low-rank recovery.

### 3.4 Choice of Parameters

In order to guarantee that the substitution in Equation 13 and 20 is a non-expansive operator, we choose parameter  $\tau = 1$  and normalize the linear map with the operator norm:  $M := \frac{M}{\sigma_{\max}(M)}$ .

Furthermore, in order to simplify the choice of the parameter  $\beta$  we assume the average value of  $C$  to be 1 and scale  $C$  if this condition is not met:  $C := \frac{N_C}{\|C\|_{\ell_1}} C$ , where  $N_C$  is the number of entries in  $C$ . After this normalization, in all our experiments we set parameters  $\beta = 0.25$  and  $\gamma = 0.1$  (see [Yuan and Yang 2009]). In principle,  $\beta$  can be increased during the optimization to yield slightly faster convergence.  $A$ ,  $B$ , and  $Z$  can be initialized arbitrarily.

As both  $M$  and  $C$  have been scaled the inverse scale has to be applied to the recovered  $A$  and  $B$  to ensure that they match.

## 4 Higher-order CSLR

The Compressive sparse and low-rank reconstruction method can be extended to higher-order tensors. Operating on higher-order tensors allows for exploiting the correlation between arbitrary dimensions. While Vasilescu and Terzopoulos [2004] applied HO-SVD to compress a given tensor we derive an algorithm for recovering sparse and low-rank tensors from random projections to drastically reduce the number of required samples. The derivation of the algorithm follows the insights made in the context of tensor completion [Liu et al. 2009] and convex multilinear estimation [Signoretto et al. 2011].

Assuming  $A$  and  $B$  are tensors several adaptation steps of the scheme proposed in Section 3 are necessary. In our framework the step of sparse  $\ell_1$  optimization of  $A$  hardly changes — the shrinkage  $S^{\ell_1}$  will simply be carried out per entry in the tensor.

Rank minimization in higher-order structures however first needs a proper definition of the norm to be minimized. In [Liu et al. 2009; Signoretto et al. 2011] it is shown that minimizing the rank of an  $n$ -order tensor is required:  $\text{rank}(X) = \sum_{i=1}^N \text{rank}(X_{\langle i \rangle})$ , where  $X_{\langle i \rangle}$  corresponds to unfolding the tensor into a matrix along dimension  $i$  and  $X^{\langle i \rangle}$  corresponds to the inverse operation of folding back to the tensor. According to the matrix case where the nuclear norm  $\|A\|_*$  is the greatest convex minorant of the rank function, minimizing the

```

initialize  $A^0, B^0, Z_1^0, \dots, Z_N^0 \leftarrow 0$ 
do
     $Y_A = A^k + M^T(Z^k - \beta(M(A^k + B^k) - C))$ 
     $A^{k+1} = S_{\frac{\gamma}{\beta}}^{\ell_1}(Y_A)$ 

    for  $i=1, \dots, N$ 
         $Y_i^B = B^k + M^T(Z_i^k - \beta(M(A^{k+1} + B^k) - C))$ 
         $B_i^{k+1} = (S_{\frac{\gamma}{\beta}}^*(Y_{i, \langle i \rangle}^B))^{\langle i \rangle}$ 
         $Z_i^{k+1} = Z_i^k - \beta(M(A^{k+1} + B_i^{k+1}) - C)$ 
    end for

     $B^{k+1} = \frac{1}{N} \sum_{i=1}^N B_i^{k+1}$ 

     $Z^{k+1} = Z^k - \beta(M(A^{k+1} + B^{k+1}) - C) = \frac{1}{N} \sum_{i=1}^N Z_i^{k+1}$ 

until convergence

```

**Figure 3:** HO-CSLR: Algorithm for higher-order compressive sparse and low-rank recovery.

rank of a tensor can be expressed by minimizing the nuclear norm in each tensor unfolding  $\sum_{i=1}^N \|X_{\langle i \rangle}\|_*$ .

Consequently, our objective for the higher-order CSLR recovery is:

$$\min_{A, B} \gamma \|A\|_{\ell_1} + \frac{1}{N} \sum_{i=1}^N \|B_{\langle i \rangle}\|_* \quad (21)$$

s.t.  $M(A+B) = C$ ,

The augmented Lagrangian now requires a different optimizer for each order  $B_i$  and the corresponding multiplier  $Z_i$  (see [Signoretto et al. 2011]):

$$L(A, B_1, \dots, B_N, Z_1, \dots, Z_N) = \gamma \|A\|_{\ell_1} + \frac{1}{N} \sum_{i=1}^N (\|B_{i, \langle i \rangle}\|_* - \langle Z_i, M(A+B_i) - C \rangle) + \frac{\beta}{2} \|M(A+B) - C\|_F^2, \quad (22)$$

where the optimal tensor  $B$  is obtained as the average of tensors  $B_i$  [Liu et al. 2009]. In our case, as we do not favor minimization in any particular order, the unweighted average is sufficient:  $B = \frac{1}{N} \sum_{i=1}^N B_i$ . The formulation with the augmented Lagrangian allows one to minimize all sub problems

$$B_i = \min_B \|B_{\langle i \rangle}\|_* \quad \text{s.t.} \quad M(A+B_i) = C \quad (23)$$

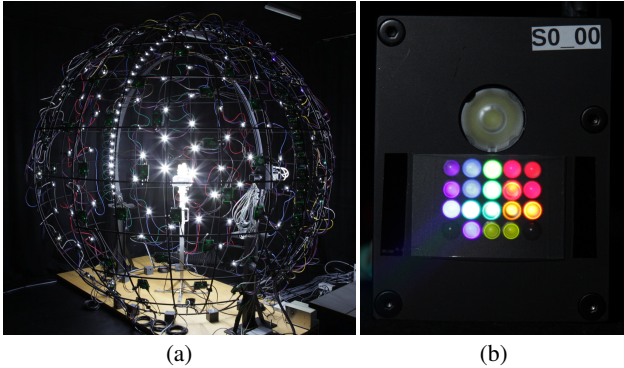
in parallel in each iteration. The update step for  $B_i$  follows Eq. 20 where the shrinkage is performed on the unfolded tensor:

$$B_i^{k+1} = S_{\frac{\gamma}{\beta}}^*(Y_{i, \langle i \rangle}^k)^{\langle i \rangle} \quad (24)$$

The complete algorithm for tensor-based CSLR is given in Figure 3. The choice of parameters is the same as described in Section 3. The measurement tensor  $M$  is normalized in a similar manner as in 3.4. However, since we cannot compute the largest singular value of the whole tensor, normalization is performed with the average of all largest singular values computed for each camera-view measurement matrix separately.

## 5 Hyperspectral Light Stage


Our custom-made light stage (see Figure 4) is a spherical device with a diameter of around 2.6m. Inside the sphere a motor-driven camera can move on and arc around the center more than  $180^\circ$ , from the equator of the sphere on one side to below the equator on the other side. In combination with the turntable in the center of the sphere which supports the scene, the target can be observed from almost any viewpoint without manual intervention.



**Figure 4:** Each of the 196 light sources on the light stage (a) are equipped with 14 LEDs covering the visible spectrum (b), 1 IR and one white LED. The weaker LEDs are doubled to increase their output.

Distributed around the sphere are 196 LED boards acting as light sources with a very dense distribution along the main arc (56 light sources) running just behind the camera’s path. Higher angular sampling rate is provided in the mirroring plane of a planar target surface, which is the most significant direction for most BRDFs. The position of each light source is calibrated.

Each light source (Figure 4) is equipped with 14 LEDs sampling the visible spectrum from 400 nm to 690 nm at approximately 20 nm increments, a near-infrared LED emitting at 950 nm, and a high-power (3 W) white LED covering most of the visible spectrum. A holographic diffuser foil with  $10^\circ$  diffusion cone is placed in front of the clear glass LEDs to form a rather uniform spot in the center of the dome. For calibration purposes, we measured the emitted spectrum of each individual LED using a photo spectrometer (Spectravista GER-1500). The spectral power distribution for one of our light sources is shown in Figure 5. The slightly irregular spacing of the narrow-band LEDs is due to the availability on the quickly developing LED market. Especially around 560nm no powerful LEDs were offered. From the plot, it is obvious that the LEDs vary



**Figure 5:** Spectra of all LEDs on a single board.

greatly with regards to their output power. We compensated for this by measuring the relative power of every individual LED mounted on the light stage using a Spectralon sphere, normalizing all measurements to 99.9% reflectivity.

The lighting configuration of the whole light stage can be changed at about 50 Hz via an I<sup>2</sup>C bus. The intensity of each LED is controlled with 8 bits but in all our measurements we either turn them on or off completely.

To minimize measurement noise, we use a DVC-4000C RGB camera featuring a 4 megapixel sensor and dual-stage electrical cooling, running the image sensor at about  $40^\circ\text{C}$  below ambient temperature which is kept at a constant  $18^\circ\text{C}$ . All our measurements are recorded as high-dynamic range images based on [Hasinoff et al. 2010] and [Granados et al. 2010]. Both the exposure time and the camera gain are varied to cover the dynamic range.

## 6 Experiments and Discussion

### 6.1 Mapping to Modes

When applying the CSLR and HO-CSLR method on data captured with our light stage one has to determine a mapping of the degrees of freedom in the setup to the modes of the tensor. The measurement setup actually provides different kinds of measurement modes. Similar to the list mentioned in [Duarte and Baraniuk 2010] we distinguish

- modes that can be sensed in parallel, in our case these are the camera pixels  $(x, y)$ , as well as the pixel color vector  $RGB$ ,
- modes that can only be acquired sequentially, for example the position of the camera on the arc as well as the turntable rotation angle  $(\theta_o, \phi_o)$ , and
- finally, modes that can be acquired by random projections, in our case the illumination direction  $(\theta_i, \phi_i)$  and the illuminating wavelength  $\lambda_i$ .


As a consequence, we perform an independent reconstruction for each color channel of each camera pixel. These reconstruction problems are completely independent and can be computed in parallel. Illumination patterns will operate and simultaneously modify the incident direction and the wavelength. The different camera perspectives will span a higher-order tensor to be reconstructed per point on the surface.

### 6.2 Illumination Patterns


The choice for the projection matrix  $M$  is actually not straight forward. In a future paper we would like to thoroughly analyze the effect of multiplexing with drastically different contributions in the context of noise and HDR imaging. Results shown in this paper neither made use of the white nor the IR LED, reducing the set of different LEDs per board to 14. While we can emit arbitrary spectra from arbitrary directions we restricted ourselves to either fully turn on or turn off an LED, resulting in random combinations of subsets of the  $196 \times 14$  LEDs.

In our experiments, a Bernoulli matrix with 50% of the LEDs turned on at random for each picture led to input images with very little variation. The influence of the individual directions and the wavelengths can hardly be discerned as approximately half of the LEDs are on at every board. The images are pretty close to a full white image.

Much better results we achieved by constructing the random patterns as a tensor product: by first randomly selecting 50% of the



**Figure 6:** Comparison of a small subset of 100 images from a full dataset (laid out as  $10 \times 10$ ) generated by different projection patterns  $M$ . Clearly, the tensor product design features the largest differences between captured images.



**Figure 7:** Example scenes under white uniform illumination.

boards/directions and then randomly selecting 50% of the LEDs on each of the selected boards. This way, in each image half of the boards are completely off. We obtain much clearer information about the directions. Even compared to a 25% Bernoulli pattern the input images of our tensor product scheme feature much higher variance between images and are thus more informative. See Figure 6 for a comparison.

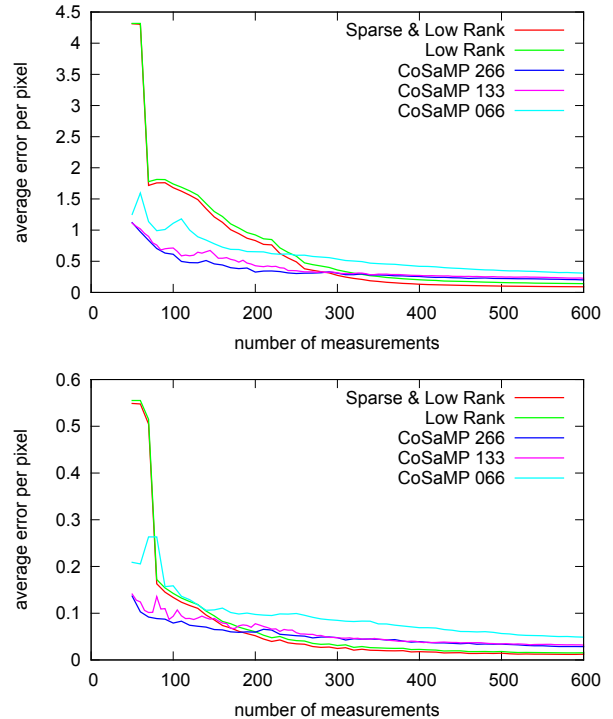
### 6.3 Single-View CSLR

Let's first investigate compressive sparse and low-rank recovery for a single camera position, i.e. fixed  $(\theta_o, \phi_o)$ . We seek matrices  $A$  and  $B$  representing the 2D arrangement of 196 illumination directions vs. 14 wavelengths ( $A, B \in \mathbb{R}^{196 \times 14}$ ). We will capture a set of  $N < 196 \cdot 14$  images, each with a different illumination as outlined in Section 6.2. The corresponding projection patterns  $M \in \mathbb{R}^{N \times (196 \cdot 14)}$  produces one measurement vector  $C \in \mathbb{R}^N$  per pixel (reconstruction is performed on Bayer interpolated data to reduce computation time and to avoid reconstructing from interpolated data). To compute the product  $M(A+B)$  matrices  $A$  and  $B$  are vectorized to  $\mathbb{R}^{196 \cdot 14}$ .

CSLR is applied to capture the reflectance fields of fluorescent markers (Fig. 1), flowers (Fig. 7(a), Fig. 8), a greeting card (Fig. 7(b)) and a tree model casting a complex shadow (Fig. 7(c)). For each particular scene we restricted the set of LEDs to those which actually illuminate the scene.

### 6.4 Capturing Fluorescence

Data collected with the hyperspectral light stage measures bispectral reflectance providing calibrated  $RGB$  color in dependence of the incident wavelength. The bispectral reflectance data captures the effect of fluorescence as visualized in Figure 1. Though illuminated with a narrow band LED each marker might re-emit a different color. In the accompanying video a comparison of the recovered markers to ground truth is shown. They match very well, wrongly reconstructed colors are however noticeable around 570 nm, where the illumination is very weak, and for the second left



**Figure 10:** Dependence of the reconstructed RMSE on the number of provided samples for the marker and the tree data set. At around 200 – 300 images CSLR outperforms CoSaMP yielding an RMSE which is about 2.4 times lower. Pure low-rank optimization leaving out the sparse components increases the error again by 25%.

marker which appears blue even under red illumination. Fluorescence is also present in Figure 8 for a set of flowers.


Compared to Hullin et al. [2010] who measured the bispectral reflectance function for a single homogeneous material, the presented setup can capture spatially varying bispectral reflectance fields and BTFs.

### 6.5 Comparison to Wavelet-based Compressive Sensing


For the single view CSLR the optical quality of the reconstruction of the tree data set is compared to ground truth and to the results obtained from a more traditional compressive sensing scheme in Figure 9. The highly efficient Compressive Sampling Matching Pursuit (CoSaMP) [Needell and Tropp 2009] exploits sparsity in the wavelet domain. The recovered reflectance field is represented using CDF-2,2 wavelets and the reconstruction is performed on the same input data as our CSLR scheme, optimizing each pixel independently. Given 400 input images the CSLR result is very close to ground truth though very dim regions might be estimated to bright. On the other hand, the CoSaMP reconstruction is still very noisy. Even with varying numbers of estimated wavelet coefficients CoSaMP reconstructs neighboring pixels differently indicating either that the number of coefficients or the number of input images is insufficient.

The dependence on the number of images of the two algorithms is further analyzed in Figure 10. CoSaMP produces results closer to ground truth for very few input images. As soon as sufficiently many constraints are provided (around 300 images) the performance of CSLR is about three times better than CoSaMP.






**Figure 8:** Slices of the bispectral reflectance field of flowers. The reflectance field reconstructed from 400 images recovers the dependence of the reflectance/re-emittance on the incident wavelength. Due to fluorescence wavelength shifts are observed. The flowers can be relit from arbitrary directions. The high contrast around the silhouettes are due to subtle motion during acquisition.



**Figure 9:** Visual comparison of ground truth, CoSaMP and CSLR on 9 different illumination directions and 14 different wavelengths. The CoSaMP reconstruction shows per-pixel noise and has difficulties to recover the correct color for near UV 400 nm independent of how many wavelet coefficients are allowed. CSLR is much closer to ground truth though it partially loses contrast in the shadow region and overestimates some of the very dim entries.



**Figure 11:** In highlight regions on the card (towards the bottom), the recovered sparse signal typically grows stronger as highlights partially violate the low-rank assumption. The apparent structure is not noise but actually due to the micro geometry of the card.


## 6.6 Sparse vs. Low-Rank

The separation into sparse and low-rank components is essential for the performance of the CSLR method. Only performing  $\ell_1$  minimization in the pixel basis does not produce any reasonable results as the intended reflectance field is not sparse at all. On the other hand, pure nuclear norm minimization does not perform that badly. There are subtle but important differences to the simultaneous recovery of both, where  $A$  will typically catch outliers. Deviations from a low-rank model will occur for example in highlight regions (Figure 11). As can be seen in Figure 10 omitting  $A$  in the estimation would increase the error compared to the full CSLR scheme by about 25%.

## 6.7 Multi-View HO-CSLR

If we add different viewing directions to the problem we can either apply the CSLR scheme to each individual view in turn or we can try to exploit the correlation between the views with the HO-CSLR. In order to align the surface points in the captured views a homography is estimated from each to the central view. Even though the alignment might not be perfect for pixels with different depth it definitely increases the correlation.

If only  $\theta_o$  is varied in the viewing directions  $A$  and  $B$  are turned into



**Figure 12:** Visualization of the multi-view tensor. The tensor  $A + B$  includes 5 different view directions shown as different columns. One clearly sees the highlight moving in the different views while the remaining entries are highly correlated. It was reconstructed from  $5 \times 200$  input images.

3-mode tensors  $A, B \in \mathbb{R}^{196 \times 14 \times N_\theta}$ , where  $N_\theta$  is the number different viewing directions. We would create a 4-mode tensor accordingly when  $\theta_o$  and  $\phi_o$  are sampled. For each view we will capture  $N$  independent measurements ( $C_{RGB} \in \mathbb{R}^{N \times N_\theta}$ ), each acquired with a different set of illumination patterns, i.e.  $M \in \mathbb{R}^{N \times (196 \cdot 14) \times N_\theta}$ . Unlike the single-view CSLR approach, here we have to perform the reconstruction on interpolated RGB data since different camera views need to be aligned, for which Bayer pattern data is not sufficient. In order to compute  $M(A + B)$  tensors  $A$  and  $B$  are matrixized into  $\mathbb{R}^{(196 \cdot 14) \times N_\theta}$ ; for each view one matrix-vector product is computed.

Multi-view reconstruction results are demonstrated on the card scene (Figure 12). We combined only 5 different views in one reconstruction. They were however sufficient to demonstrate the additional savings by SLR minimization of the tensor. Compared to reconstructing each view individually similar visual quality is achieved with 50% of the input images. The reconstruction time of the HO-CSLR is however currently about 10 times slower.

## 6.8 Savings in Acquisition Time

Besides the previously reported savings in the number of images when comparing the CSLR or HO-CSLR to ground truth data sets where each image is illuminated with just a single LED, also the savings in exposure time are significant.

The chosen projection patterns typically turn on approx. 600 LEDs on average resulting in a significantly brighter illumination compared to a single LED. In the HDR sequence the longest exposure time could therefore be reduced from 4 s with 12 dB gain to only 0.12 s with 12 dB gain. The total time per HDR image changed from 15 s to 0.5 s accordingly. Combined with the savings in the number of images, acquisition time was reduced from about 4 hours to 17 minutes.

## 6.9 Limitations

One common problem in all the presented results is the relatively poor SNR for 574 and 590 nm where the intensity of the LEDs is very low. In all input images their small signal is mixed with brighter ones. A similar effect is happening in shadow regions which move from light source to light source (see Figure 9). As it is highly unlikely that in any illumination pattern the region is in full shadow without being directly illuminated by another LED the intensity might be overestimated. More research is necessary to predict the expected loss in the recovered signal.

A more severe error occurs whenever the captured frames deviate from the linear light transport assumption. We identified two cases: One is due to object motion. The other appears on highly specular surfaces where a very bright highlight is affecting a single sensor pixel. Even if correctly exposed in the HDR sequence during demosaicing its energy might spread non-linearly to neighbor pixels. This highlight will occur whenever the corresponding illumination direction is included in the pattern, which typically is in at least 25% of the samples.

## 7 Conclusion

Compressive sensing typically exploits sparseness in the model for more efficient acquisition. Matrix and tensor completion or norm minimization on the other hand try to find a good low-rank approximation for a given data set. The proposed compressive sparse and low-rank recovery algorithms (CSLR) optimizes both simultaneously on compressed input data. They take random projections and simultaneously optimize for low-rank and sparse decomposition of the underlying model. Joint recovery of low-rank and sparse data yields significantly better SNR at fewer input samples than wavelet-based compressive sensing. Extending the same formalism to higher-order tensors (HO-CSLR) even the correlation along different modes is successively explored to obtain further savings in the number of input samples.

The recovery algorithms implement relatively simple iterative schemes follow an augmented Lagrangian approach and are based on the shrinkage operators for the  $\ell_1$  and nuclear norm respectively. At this point, we stress that the approach can be easily adapted to other higher-order compressive sensing scenarios.

The proposed application scenario in this paper is the acquisition of spatially varying, bispectral reflectance data captured in a LED-based hyperspectral light stage. Reflectance fields are known to be locally low-rank and therefore constitute a perfect match to the proposed CSLR and HO-CSLR recovery. Most of the energy is indeed found to be low-rank while strong view and illumination

dependent effects such as highlights are to some extent recovered in the sparse part of the reconstruction.

The novel application of higher-order compressive sensing opens a new set of problems to be addressed in the future. More research is necessary to determine optimal higher-order illumination codes, in particular in the context of HDR imaging. Another area of improvement is of course the computation time. Here, parallelization on GPUs and the use of Monte Carlo SVD and alike are promising directions.

## References

- BARON, D., WAKIN, M. B., DUARTE, M. F., SARVOTHAM, S., AND BARANIUK, R. G. 2009. Distributed compressed sensing. Tech. rep.
- CANDÈS, E. J., AND PLAN, Y. 2009. Matrix completion with noise. *CoRR abs/0903.3131*.
- CANDÈS, E. J., AND RECHT, B. 2009. Exact matrix completion via convex optimization. *Foundations of Computational Mathematics* 9, 6, 717–772.
- CANDÈS, E. J., ROMBERG, J., AND TAO, T. 2006. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on Information Theory* 52, 2 (February), 489–509.
- CHANDRASEKARAN, V., SANGHAVI, S., PARRILO, P. A., AND WILLISKY, A. S. 2009. Rank-sparsity incoherence for matrix decomposition. *Arxiv 02139*, 1–24.
- CHANDRASEKARAN, V., SANGHAVI, S., PARRILO, P. A., AND WILLISKY, A. S. 2009. Sparse and low-rank matrix decompositions. In *Proceedings of the 47th annual Allerton conference on Communication, control, and computing*, Allerton’09, 962–967.
- DANA, K. J., AND WANG, J. 2004. Device for convenient measurement of spatially varying bidirectional reflectance. *Journal of the Optical Society of America A* 21, 1, 1–12.
- DEBEVEC, P., HAWKINS, T., TCHOU, C., DUIKER, H.-P., SAROKIN, W., AND SAGAR, M. 2000. Acquiring the reflectance field of a human face. *SIGGRAPH ’00*, 145–156.
- DEBEVEC, P., WENGER, A., TCHOU, C., GARDNER, A., WAESE, J., AND HAWKINS, T. 2002. A lighting reproduction approach to live-action compositing. In *Proc. ACM SIGGRAPH*, 547–556.
- DONOHO, D. L. 2006. Compressed sensing. *Information Theory, IEEE Transactions on* 52, 4, 1289–1306.
- DUARTE, M. F., AND BARANIUK, R. G. 2010. Kronecker compressive sensing. *IEEE Transactions on Image Processing* (Nov.).
- GARDNER, A., TCHOU, C., HAWKINS, T., AND DEBEVEC, P. 2003. Linear light source reflectometry. In *ACM SIGGRAPH 2003, SIGGRAPH ’03*, 749–758.
- GHOSH, A., HEIDRICH, W., ACHUTHA, S., AND O TOOLE, M. 2010. A basis illumination approach to brdf measurement. *International Journal of Computer Vision* 90, 183–197.
- GOLDFARB, D., AND MA, S. 2009. Convergence of fixed point continuation algorithms for matrix rank minimization. *CoRR abs/0906.3499*.

- GRANADOS, M., AJDIN, B., WAND, M., THEOBALT, C., SEIDEL, H.-P., AND LENSCH, H. P. A. 2010. Optimal hdr reconstruction with linear digital cameras. In *The Twenty-Third IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2010, San Francisco, CA, USA, 13-18 June 2010*, IEEE Computer Society, 215–222.
- HALE, E. T., YIN, W., AND ZHANG, Y. 2007. A fixed-point continuation method for  $l_1$ -regularized minimization with applications to compressed sensing. Tech. rep., Rice University.
- HASINOFF, S. W., DURAND, F., AND FREEMAN, W. T. 2010. Noise-optimal capture for high dynamic range photography. In *The Twenty-Third IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2010, San Francisco, CA, USA, 13-18 June 2010*, IEEE Computer Society, 553–560.
- HULLIN, M. B., HANIKA, J., AJDIN, B., SEIDEL, H.-P., KAUTZ, J., AND LENSCH, H. P. A. 2010. Acquisition and analysis of bispectral bidirectional reflectance and reradiation distribution functions. *ACM Trans. Graph. (Proc. SIGGRAPH 2010)* 29, 4, 97:1–97:7.
- LIU, J., MUSIALSKI, P., WONKA, P., AND YE, J. 2009. Tensor completion for estimating missing values in visual data. In *ICCV'09*, 2114–2121.
- LIU, H., LI, Y., WU, C., WANG, K., AND WANG, Y. 2010. Compression of hyperspectral imagery based on compressive sensing and interband prediction. In *Optics + Photonics: Optical Engineering + Applications*, The International Society for Optical Engineering., vol. 7810, 781016–781016–9.
- MA, S., GOLDFARB, D., AND CHEN, L. 2009. Fixed point and bregman iterative methods for matrix rank minimization. *CoRR abs/0905.1643*.
- MÜLLER, G., BENDELS, G. H., AND KLEIN, R. 2005. Rapid synchronous acquisition of geometry and btf for cultural heritage artefacts. In *The 6th International Symposium on Virtual Reality, Archaeology and Cultural Heritage (VAST)*, 13–20.
- MÜLLER, G., MESETH, J., SATTLER, M., SARLETTE, R., AND KLEIN, R. 2005. Acquisition, synthesis and rendering of bidirectional texture functions. *Computer Graphics Forum* 24, 1 (Mar.), 83–109.
- NEDELL, D., AND TROPP, J. 2009. Cosamp: Iterative signal recovery from incomplete and inaccurate samples. *Applied and Computational Harmonic Analysis* 26, 3, 301–321.
- NOCEDAL, J., AND WRIGHT, S. 2006. *Numerical optimization*, 2. ed. ed. Springer series in operations research and financial engineering. Springer, New York, NY.
- O'TOOLE, M., AND KUTULAKOS, K. N. 2010. Optical computing for fast light transport analysis. In *ACM SIGGRAPH Asia 2010, SIGGRAPH ASIA '10*, 164:1–164:12.
- PARK, J., LEE, M., GROSSBERG, M. D., AND NAYAR, S. K. 2007. Multispectral Imaging Using Multiplexed Illumination. In *IEEE International Conference on Computer Vision (ICCV)*.
- PEERS, P., MAHAJAN, D. K., LAMOND, B., GHOSH, A., MATUSIK, W., RAMAMOORTHY, R., AND DEBEVEC, P. 2009. Compressive light transport sensing. *ACM Trans. Graph.* 28 (February), 3:1–3:18.
- RATNER, N., AND SCHECHNER, Y. Y. 2007. Illumination multiplexing within fundamental limits. In *2007 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2007), 18-23 June 2007, Minneapolis, Minnesota, USA*, IEEE Computer Society.
- RUMP, M., AND KLEIN, R. 2010. Spectralization: Reconstructing spectra from sparse data. In *SR '10 Rendering Techniques*, J. Lawrence and M. Stamminger, Eds., 1347–1354.
- SCHECHNER, Y. Y., NAYAR, S. K., AND BELHUMEUR, P. N. 2003. A theory of multiplexed illumination. In *9th IEEE International Conference on Computer Vision (ICCV 2003), 14-17 October 2003, Nice, France*, IEEE Computer Society, 808–815.
- SEN, P., AND DARABI, S. 2009. Compressive Dual Photography. *Computer Graphics Forum* 28, 2, 609 – 618.
- SIGNORETTO, M., DE LATHAUWER, L., AND SUYKENS, J. A. K. 2011. Nuclear norms for tensors and their use for convex multilinear estimation. *Linear Algebra and Its Applications (preprint)*.
- SUN, T., AND KELLY, K. 2009. Compressive sensing hyperspectral imager. Optical Society of America, CTuA5.
- TSAIG, Y., AND DONOHO, D. L. 2006. Compressed sensing. *IEEE Trans. Inform. Theory* 52, 1289–1306.
- VALIOLLAHZADEH, S. M., AND YIN, W. 2010. Hyperspectral data reconstruction combining spatial and spectral sparsity. Tech. rep., Rice University.
- VASILESCU, M. A. O., AND TERZOPOULOS, D. 2004. Tensortextures: multilinear image-based rendering. *ACM Trans. Graph. (Proc. SIGGRAPH)* 23 (August), 336–342.
- WANG, J., DONG, Y., TONG, X., LIN, Z., AND GUO, B. 2009. Kernel Nyström method for light transport. *ACM Trans. Graph.* 28 (July), 29:1–29:10.
- WATERS, A. E., SANKARANARAYANAN, A. C., AND BARANIUK, R. G. 2011. Sparcs: Recovering low-rank and sparse matrices from compressive measurements. In *Neural Information Processing Systems (NIPS)*.
- WENGER, A., GARDNER, A., TCHOU, C., UNGER, J., HAWKINS, T., AND DEBEVEC, P. 2005. Performance relighting and reflectance transformation with time-multiplexed illumination. *ACM Trans. Graph.* 24 (July), 756–764.
- YUAN, X., AND YANG, J., 2009. Sparse and low-rank matrix decomposition via alternating direction methods.