

Efficient Basis Decomposition for Scattered Reflectance Data

R. Peter Weistroffer

Kristen R. Walcott

Greg Humphreys

Jason Lawrence

University of Virginia

Abstract

Recent progress in acquisition technology has increased the availability and quality of measured appearance data. Although representations based on dimensionality reduction provide the greatest fidelity to measured data, they require assembling a high-resolution and regularly sampled matrix from sparse and non-uniformly scattered input. Constructing and processing this immense matrix becomes a significant computational bottleneck. We describe a technique for performing basis decomposition directly from scattered measurements. Our approach is flexible in how the basis is represented and can accommodate any number of linear constraints on the factorization. Because its time- and space-complexity is proportional to the number of input measurements and the size of the output, we are able to decompose multi-gigabyte datasets faster and at lower error rates than currently available techniques. We evaluate our approach by representing measured spatially-varying reflectance within a reduced linear basis defined over radial basis functions and a database of measured BRDFs.

Categories and Subject Descriptors (according to ACM CCS): I.4.1 [Digitization and Image Capture]: Reflectance

1. Introduction

Recent advances in digital cameras, lighting technology, and automated acquisition methods have increased the accessibility and quality of measured appearance data. These datasets are typically far too large to be used directly, so recent research has explored compact, accurate, and editable representations. Those based on basis function decomposition typically provide the greatest accuracy and have become practical alternatives to conventional analytic reflectance models in a variety of contexts including real-time rendering systems [MAA01, MMS*04], global illumination simulations [LRR04], and material design [LBAD*06].

One key limitation of these techniques is their requirement that the measured data be organized into a matrix. Only then can the matrix be factored, effectively projecting the input onto a reduced linear basis. However, the high dimensionality of light scattering functions makes it generally impractical to obtain the dense set of measurements necessary for matrix generation. Even the multi-gigabyte datasets returned by current acquisition systems [GLL*04, LKG*03, NDM05] contain only sparse and scattered measurements. To populate the matrix from such data, an expensive reconstruction effort is required, and the efficiency of the subsequent factorization is hindered by its massive size.

This paper introduces an efficient and general method for performing linear basis decomposition directly from a set of scattered and sparse measurements. Our key insight is that the estimated linear basis should itself be defined with respect to a fixed secondary basis. The choice of secondary basis is guided by prior work in accurate and efficient representations of measured reflectance (e.g., wavelets [MPBM03b], radial basis functions [ZERB05], a database of measured BRDFs [MPBM03a]). We introduce an iterative algorithm that solves a series of convex quadratic programming problems to perform factorization with respect to a secondary basis and scattered input. The complexity of the resulting algorithm is proportional to the number of input samples and the size of the output, a significant improvement over currently available techniques. Furthermore, our formulation can accommodate hard constraints on the decomposition such as non-negativity and energy conservation, which are particularly important in the context of material representation [CBCG02, LRR04, LBAD*06]. We explore the use of radial basis functions and a BRDF database [MPBM03a] in this framework and compare the efficiency of our approach to currently available methods using two datasets of measured spatially-varying reflectance.

2. Background and Prior Work

Although our work could be used to represent a variety of light scattering functions, this paper focuses on heterogeneous opaque reflectance, which is characterized by the 6-D Spatially-Varying Bidirectional Reflectance Distribution Function (SVBRDF) [NRH77]: $S(\vec{x}, \vec{\omega}_i, \vec{\omega}_o)$. The appearance of such materials may vary with 2-D spatial position along an object's surface \vec{x} , incoming light direction $\vec{\omega}_i$, and reflected direction $\vec{\omega}_o$ (we will often abbreviate the complete 4-D angular domain as $\vec{\theta} = (\vec{\omega}_i, \vec{\omega}_o)$).

Fitting Parametric Models: A common method for appearance representation is to fit analytic reflectance functions to data [War92, LFTG97, McA02, GTHD03, LKG*03, GCHS05, NDM05, MWAM05]. Although they provide a parameterized and compact representation that is well-defined over the entire domain, fitting these models to measured data can result in significant error [NDM05, LBAD*06]. Additionally, this approach relies on fragile non-linear optimization which is often unstable in practice. Furthermore, these methods are not appropriate for representing datasets with widely varying reflectance properties or those that exhibit global illumination or visibility effects typically described by the Bidirectional Texture Function (BTF) [DvNK99, MMS*04].

Scattered Data Interpolation: Several interpolation strategies have been used in the context of material representation. Most related to this paper is the work of Zickler et al. [ZERB05, ZERB06] which uses radial basis functions (RBFs) to recover a smooth approximation of an SVBRDF from a sparse set of images. One key limitation of their approach is that it assumes the angular component of surface reflectance varies smoothly across the surface, whereas our approach is designed to handle SVBRDFs that exhibit complex high-frequency spatial blending between component materials (Fig. 5 and Fig. 13).

Basis function decomposition: Numerous dimensionality-reduction algorithms have been applied to the task of material representation, including variants of principal component analysis (PCA) [KM99, FKIS02, VT04], homomorphic factorization [MAA01], independent component analysis (ICA) [TOS*03], k -means clustering [LM01], and variants of non-negative matrix factorization (NMF) [CBCG02, LRR04, PvBM*06, LBAD*06].

With the exception of homomorphic factorization [MAA01], which represents the input as a single product and cannot support linear combinations, these methods require tabulating the data into a matrix of the same resolution as the final representation. This matrix is typically several orders of magnitude larger than the combined size of the input and output [CBCG02, LBAD*06] and becomes a significant (and theoretically unnecessary) computational bottleneck. For this reason, existing systems often consider only a tiny subset of this matrix [LBAD*06] or limit the

resolution of the output [DvNK99, MMS*04]. This strategy is particularly inefficient for handling measurements of spatially-varying reflectance collected from curved surfaces [ZERB05, LKG*01, GCHS05] and is one reason prior work has exclusively used planar samples: to avoid significant interpolation of the data across a regularly sampled grid [DvNK99, MMS*04, LBAD*06]. We alleviate the need for this intermediate matrix through a more general and flexible definition of the basis.

3. Overview

We consider decompositions of the SVBRDF into a sum of products of 2- and 4-D functions

$$S(\vec{x}, \vec{\theta}) \approx \sum_{k=1}^K F_k(\vec{x}) G_k(\vec{\theta}), \quad (1)$$

where the $F(\vec{x})$ can be interpreted as spatially-varying coordinates (often called *blending weights*) within the K -dimensional linear basis spanned by the $G(\vec{\theta})$ (often called the *basis BRDFs*). For most real-world materials, S can be well approximated with a small K , enabling accurate and compact representations.

A standard approach is to perform the decomposition in Eq. 1 by computing a rank- K factorization of a matrix. Measurements of S are assembled into a matrix Z so that spatial variation is preserved across its rows and angular variation across its columns (Fig. 1b). Alternatively, we can represent Z as

$$Z \approx WH, \quad (2)$$

where the K columns of W and the K rows of H contain the discretely sampled blending weights and basis BRDFs, respectively. The particular algorithm used to factor Z will depend on the desired convergence properties, error guarantees, and any constraints placed on the elements of W and H . Note that we assume K is known and fixed; estimating this number automatically remains an open research topic [TJBB04].

This standard approach has two key drawbacks that this paper aims to address. First, because densely measuring S along a regular *angular* grid is often impractical, each row of Z must be interpolated in undersampled regions. This is a time-consuming and difficult operation, particularly for multi-dimensional functions and very sparse input like that collected in [LKG*03, ZERB05]. Second, the number of columns in Z establishes the resolution of the 4-D basis BRDFs (3-D for isotropic materials) and should be very large to enable accurate decompositions. Together, the size of Z becomes a significant computational bottleneck both in terms of the effort required to interpolate scattered data over such a large grid and the price of computing its subsequent factorization.

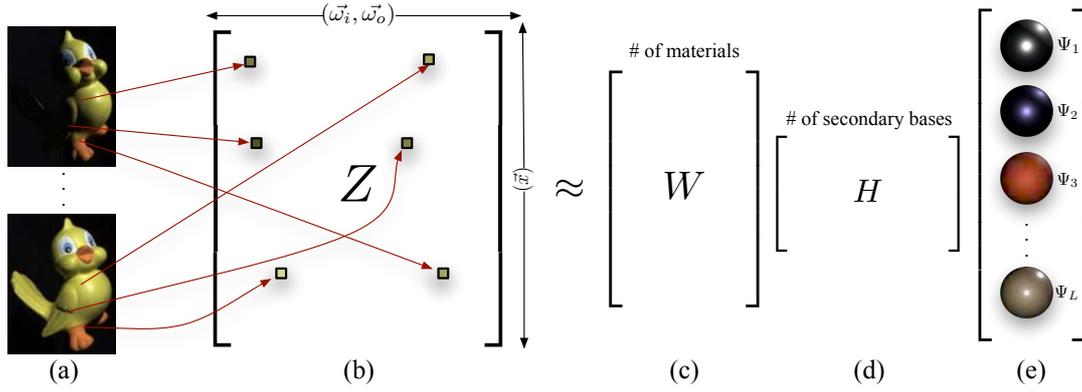


Figure 1: Our approach to SVBRDF factorization. (a) High dynamic range reflectance measurements from known geometry sparsely populate (b) a matrix Z whose rows vary with spatial position and columns with angular configuration. We factor Z into a product of (c) W , spatially distributed blending weights, and (d) H , another set of weights applied to (e) a secondary linear basis. Here we visualize the secondary basis as several lit spheres using the material database from Matusik et al [MPBM03a]. Note that we do not explicitly construct the full Z matrix, instead operating on only the scattered input data.

To overcome these limitations, we employ a second level of indirection and define the primary basis functions G with respect to a secondary linear basis Ψ

$$S(\vec{x}, \vec{\theta}) \approx \sum_{k=1}^K \left(F_k(\vec{x}) \sum_{l=1}^L \lambda_{kl} \Psi_l(\vec{\theta}) \right). \quad (3)$$

This formulation offers several advantages:

- There are a number of bases (i.e., choice of Ψ) that offer accurate and compact representations of measured reflectance (e.g., wavelets [MPBM03a], spherical harmonics [KvDS96], and radial basis functions [ZERB05]). Defining G with respect to the appropriate basis can significantly reduce the number of free parameters in the decomposition, which in turn reduces computation time and memory requirements, and improves the stability of the optimization in the face of very sparse and noisy input.
- This strategy avoids an explicit interpolation stage.
- It generalizes prior work: standard factorization algorithms correspond to setting Ψ to the reconstruction kernel used for interpolation (e.g., hat, box, or Gaussian) and distributed according to the resolution of Z .
- The parameters F and λ can be estimated with a straightforward iterative algorithm based on solving a series of convex Quadratic Programming (QP) problems (Section 4.1). This also accommodates constraints on the decomposition, such as non-negativity and energy conservation, that are important for material representation [LRR04, LBAD*06].

Although the choice of Ψ greatly influences the quality of the resulting approximation, we first present our basic algorithm and examine the issues involved in selecting a suitable secondary basis in Sec. 5.

4. Basis Decomposition for Scattered Data

Although we do not explicitly construct the matrix Z , it is conceptually useful to think of it as an extremely sparse matrix with M rows and N_m scattered entries in the m^{th} row (Fig. 1b). The matrix W is $M \times K$ and corresponds to the blending weights we intend to estimate (i.e., discretely sampled representations of F_k in Eq. 3), and the matrix H is $K \times L$ and stores the weights $H_{kl} = \lambda_{kl}$ that are applied to Ψ in Eq. 3. The example shown in Fig. 1 uses a database of measured BRDFs [MPBM03a] as a secondary basis.

Let $\vec{\lambda}_k = (\lambda_{k1}, \dots, \lambda_{kL})$ denote the vector of weights that define G_k in Eq. 3. The value of G_k at an arbitrary location $\vec{\theta}$ is thus equal to the inner product of $\vec{\lambda}_k$ and a vector of the values of each Ψ_l evaluated at $\vec{\theta}$. Similarly, the values of G_k at a discrete set of N locations can be computed as the product of a matrix A and this weight vector

$$(G_k(x_1), \dots, G_k(x_N)) = \vec{\lambda}_k A, \quad (4)$$

where $A_{ij} = \Psi_i(x_j)$ and is $L \times N$, as desired. Next, consider our approximation of a single row in Z . In the case of the SVBRDF, z_m will contain N_m scattered measurements of a BRDF. We evaluate each Ψ_l at these scattered locations to form A_m as in Eq. 4, giving

$$z_m \approx w_m H A_m, \quad (5)$$

where $H_{kl} = \lambda_{kl}$ and w_m is the m^{th} row of W . Eq. 5 is the fundamental relationship that allows relating the parameters we intend to estimate (w_m and H) to the input. Fig. 4 provides a didactic illustration of how A_m is constructed for a simple 1-D example. We next describe our algorithm for computing the W and H that best approximate Z .

where $Q'_m = HA_m B_m B_m A_m^T H^T$ and $c'_m = -z_m B_m B_m A_m^T H^T$ (although not shown, E'_H can be derived similarly). In practice we fold these weights into the A matrices ($A' = AB$) to save additional memory and computation costs.

4.2. Error Metric

Following the recent work of Ngan et al. [NDM05] for fitting analytic BRDF models to measured data, we will seek decompositions that minimize a cosine-weighted L_2 distance

$$E = \sum [S(\vec{x}; \vec{\omega}_i, \vec{\omega}_o) \cos(\theta_i) - \tilde{S}(\vec{x}; \vec{\omega}_i, \vec{\omega}_o) \cos \theta_i]^2, \quad (15)$$

where \tilde{S} is the reconstructed SVBRDF and the sum is taken over the entire set of scattered measurements. Because our focus is on fitting very sparse and scattered angular measurements, we omit the solid angle term present in [NDM05]. Lastly, we disregard measurements with incoming or outgoing angles larger than 80 degrees. We can form a weighting matrix B by evaluating $\cos \theta_i$ at each scattered measurement which causes our algorithm to minimize Eq. 15, as desired.

4.3. Constraints

The ability to place constraints such as non-negativity and energy conservation on the resulting decomposition is important in many applications, including real-time rendering [MAA01], importance sampling for physically-based rendering [LRR04], and material design [LBAD*06]. Although this paper does not attempt to analyze this design space, focusing instead on the performance of our basic algorithm, it does support these types of constraints. The key difference between our algorithm and existing techniques is that constraints on the primary basis must be expressed in terms of the secondary basis. For example, if a collection of measured BRDFs are used as the secondary basis and each is guaranteed to be non-negative and conserve energy, we can restrict their weights (λ_{kl} in Eq. 3) to be non-negative and sum to a value less than or equal to 1, thereby guaranteeing the resulting primary basis is also physically plausible. We leave further investigation of handling constraints in our framework as future work. However, we do enforce non-negativity in the blending weights W in all of our experiments to encourage more intuitive separations.

4.4. Complexity

The space complexity of our algorithm is dominated by the A_m matrices and is $O(TL)$, where T is the total number of measurements ($T = \sum_m N_m$). Its time complexity is a function of the input size T , the size of the secondary basis L , and the number of terms in the factorization K . Updating each row of W requires sequentially solving M ($K \times K$) QP problems, and updating H requires solving a single ($KL \times KL$)

QP problem. The total run time depends on the number of iterations performed; we terminate the optimization once the error changes by less than 0.1% between steps.

Because the size of real-world reflectance datasets typically exceeds the capacity of main memory, we must perform some subsampling of the input. We follow the strategy in [LBAD*06] which proceeds in two passes. First, we perform a complete factorization on a subset of the rows of Z to form $\tilde{Z} \approx \tilde{W}H$ (subsampling the columns is also possible). In a second pass, we backproject each row in Z onto the basis H with a single application of the steps in Sec. 4.1.1. We investigate the sensitivity of our algorithm to how aggressively we subsample the input in Sec. 5 and Sec. 6.

5. Choice of Secondary Basis

Although our algorithm is mathematically independent of the choice of secondary basis, an appropriate set of functions will greatly improve its performance and stability. We have analyzed two bases: radial basis functions (RBFs) and a database of measured isotropic BRDFs [MPBM03a]. We selected these as opposing points in this design space. On the one hand, RBFs are designed to interpolate arbitrary smooth functions, but may require a large number of centers and input data to form accurate approximations. On the other hand, a material database is a more restrictive basis, but it is meant to represent real-world BRDFs and has a greater ability to yield plausible estimates of the reflectance far away from actual measurements.

5.1. Radial Basis Functions

Zickler et al. showed that a modest number of RBFs can approximate measured reflectance data well [ZERB05] (see [Pow92] for a thorough review of RBFs). Whereas their work focused on interpolating a single SVBRDF under the assumption that its angular component varies smoothly over the surface, we are instead interested in using RBFs to approximate each of the K basis BRDFs which may be distributed in arbitrarily complex spatial patterns. Similarly to their work, our approach may also be thought of as “sharing reflectance information,” but only across regions with similar optical properties as determined by the estimated blending weights.

We define each basis BRDF as

$$G_k(\vec{\theta}) = c_k + \sum_{l=1}^L \lambda_{kl} \Psi_l(\|\vec{\theta} - \vec{\theta}_l\|), \quad (16)$$

where c_k is a constant term and Ψ_l are functions of radial distance. The number of functions L required for accurate approximations is heavily influenced by the parameterization of the BRDF domain, the analytic form of Ψ , and the center locations $\vec{\theta}_l$.

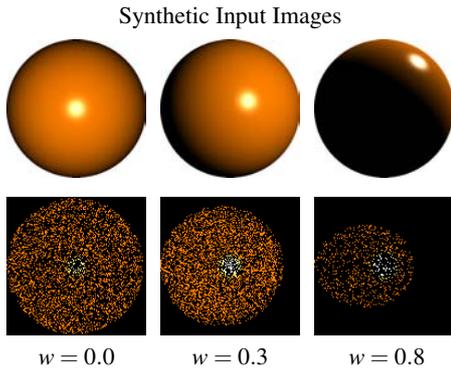


Figure 3: Reflectance samples captured in (top) several images of a sphere under varying point illumination (bottom) reparameterized as in [ZERB05]. The sphere’s reflectance is modeled as a Torrance-Sparrow BRDF [TS66] ($\sigma = 0.007$ and $k_s = 0.15$) under distant point illumination from ten randomly chosen directions. The bottom row shows the distribution of the reflectance samples contained in these images at three slices of constant w . Note that the specular highlight is aligned along the w -axis and the shape of the horizon causes the irregular clipping near grazing angles.

Parameterization: We use the parameterization described in [ZERB05] for representing isotropic BRDFs

$$(u, v, w) = \left(\sin \theta_h \cos 2\phi_d, \sin \theta_h \sin 2\phi_d, \frac{2\theta_d}{\pi} \right), \quad (17)$$

where θ_h , θ_d , and ϕ_d denote the elevation and azimuthal angles in the half/difference frame introduced by [Rus98] (ϕ_h may be ignored for isotropic BRDFs). Fig. 3 visualizes this parameterization for a synthetic BRDF.

Choice of Radial Function: We have experimented with a number of radial functions (e.g., linear, cubic, exponential) and found that a simple exponential $\Psi(r) = \exp(-\frac{r^2}{2\sigma^2})$ works well for a wide range of materials (recall r is equal to $\|(u, v, w) - (u_l, v_l, w_l)\|$ in our case). As indicated by the blue dots in the bottom row of Fig. 4, we place centers along nine uniformly spaced slices of w (3 of 9 are shown) to form a radial pattern in the (u, v) plane which is more dense near $(u, v) = (0, 0)$ where specular highlights occur. Additionally, we increase their overall density as $w \rightarrow 1$ to better capture the brightening and sharpening of the highlight near grazing angles. In all, this pattern totals 381 centers plus a constant and was used for all the RBF results shown in this paper.

Verification: We verified these design decisions using two synthetic datasets meant to simulate the type of data acquired in practice (we present results for measured data in Sec. 6). We recovered a BRDF represented as RBFs from the measurements in Fig. 3 by solving the corresponding QP problem in Eq. 12 (this corresponds to estimating H for a single-term factorization and fixed unit intensity blending weights W). We handle color by estimating a separate set

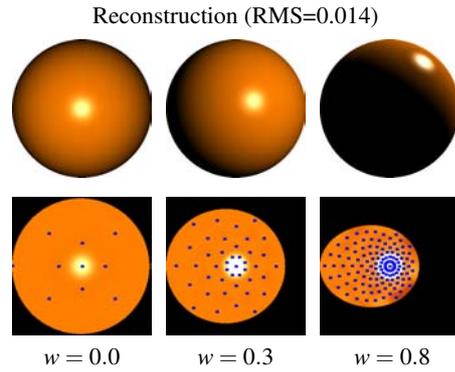


Figure 4: A BRDF represented as a weighted combination of RBFs estimated from the samples shown in Fig. 3. For comparisons, we show corresponding (top) spheres under point illumination and (bottom) several slices within the Zickler parameterization. Blue points mark the location of centers.

of RBF weights for each channel in the RGB colorspace. Fig. 4 shows the resulting approximation as spheres rendered under the same lighting conditions and for the same (u, v) slices shown in Fig. 3. Aside from some slight ringing near extreme grazing angles, the specular highlight and diffuse color are accurately reproduced. We also report the RMS error, taken as the square root of the normalized weighted metric, in Eq. 15.

We performed additional verification using a synthetic SVBRDF dataset consisting of the Torrance-Sparrow BRDF from Figures 3 and 4 mixed with an Oren-Nayar BRDF [ON94] ($\sigma = 0.2$) to form a checkerboard pattern mapped to a sphere. Figure 5 shows two of the ten input images and a 2-term factorization computed with our technique. We subsampled the input to consider only 4,096 of the possible 200K valid pixels, or 2%, to perform the factorization and then recovered a complete set of blending weights during a backprojection phase. These are visualized as grayscale images in Fig. 5 and clearly show the checkerboard pattern and the basis BRDFs are visualized as lit spheres. The decomposition converged after 6 iterations or ~ 3.5 minutes. Figure 6 compares the visual quality of this factorization to two of the ten input images and at one light position not included in the training set.

5.2. Material Database

We have also investigated using the MERL-MIT database of isotropic BRDFs collected by Matusik et al. [MPBM03a] as a secondary basis. We manually selected 55 of the highest quality examples (from a total of 100) that span the range of materials present in the collection and expanded each into three separate functions, one for each color channel, for a total of 165 different functions (i.e., $L = 165$ in Eq. 3).

Using a *data-driven* secondary basis has a number of advan-

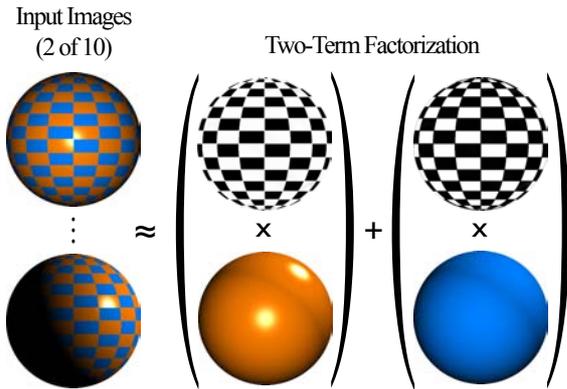


Figure 5: Illustration of a two-term factorization in which the basis BRDFs are represented using RBFs. Left: two of the ten synthetic input images, rendered with pbrt [PH04]. Right: each term consists of a set of blending weights, visualized here as grayscale images, and a basis BRDF, visualized as lit spheres.

tages. First, the size of the basis is relatively small so the factorization converges quickly. Second, the quality of the match is directly related to the quality and completeness of the database. Third, each function has wide support over the entire domain, allowing robust and reliable approximations even for very sparse input.

We used the synthetic dataset shown in Fig. 5 to verify the accuracy of this approach. Fig. 7 charts the RMS error in a two-term factorization as a function of the size of the subset of Z used to compute the decomposition (recall that we process only a subset of the input and recover complete blending weights during a backprojection phase as described in Sec. 4.4). Compared to RBFs, the error in the factorization is lower at very small input sizes, but flattens out once additional samples fail to change the choice of the best-fitting material weights. In contrast, the flexibility afforded by the larger number of RBFs allows a steadier reduction in error, but its accuracy relies on a larger number of input samples. Fig. 6 shows images rendered using the factorizations generated from a randomly chosen subset of 4,096 locations.

Figure 8 shows blending weights and basis BRDFs estimated from a measured dataset of spatially-varying reflectance (to save space only 3 of 5 terms are shown, the additional terms capture the color in the eyes and more subtle variation in the reflectance along the body). Although Sec. 6 provides further analysis of this dataset, we present it here to help illustrate the use of a material database in our framework. Note that each basis BRDF is itself represented as a weighted combination of the elements in the MERL-MIT BRDF database. The weights used to recover the basis BRDF shown in the middle column of Fig. 8 are visualized in Fig. 9.

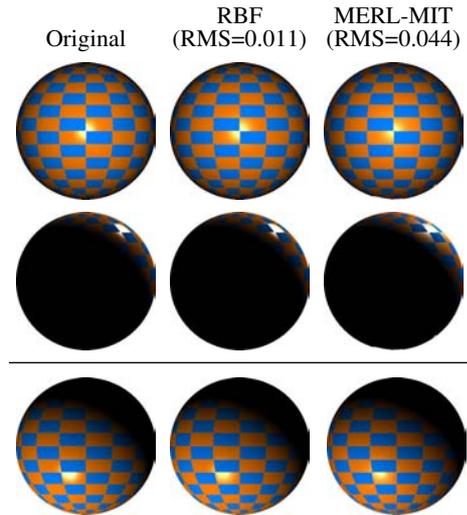


Figure 6: Comparison of factorizations generated using two different secondary bases for a synthetic SVBRDF dataset. Left: ground truth images (only the top two were included in the training set). Middle: reconstructions at the same light positions using RBFs as the secondary basis. Right: reconstructions using a BRDF database as the secondary basis.

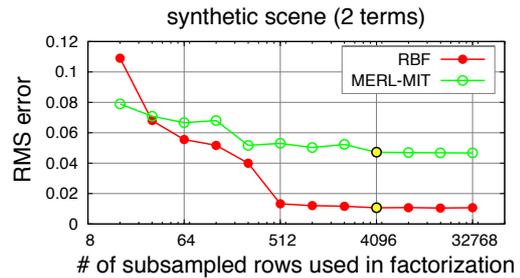


Figure 7: Reconstruction error as a function of input size for two secondary bases.

6. Comparison to Prior Work

We compare our algorithm to Alternating Constrained Least Squares (ACLS) introduced in [LBAD*06] which is an iterative procedure for performing the decomposition in Eq. 2 for a regular matrix Z that allows placement of arbitrary linear constraints on W and H . Similarly to our approach, it works by solving a series of small convex QP problems.

For these comparisons, we use the *bird* and *angels* spatially varying BRDF datasets from Lensch et al [LKG*01] which contain 1.9 million and 1.6 million BRDF samples, respectively. These samples are obtained from 25 and 27 images of known geometry. We empirically determined their respective number of primary basis functions to be $K = 5$ and $K = 11$. Unlike planar samples, images of curved surfaces contain very sparse angular measurements at any sin-

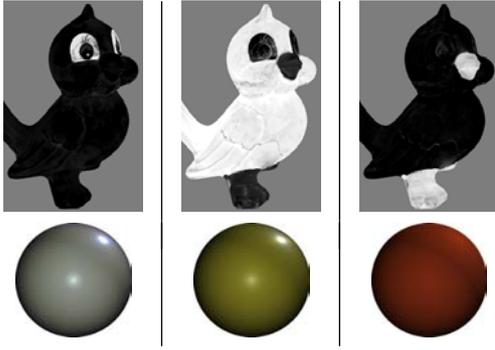


Figure 8: Three of five terms in a factorization of the *bird* dataset using a BRDF database as the secondary basis. Top: Blending weights, constrained to be non-negative, are shown as grayscale images. Bottom: corresponding basis BRDFs are visualized as lit spheres.

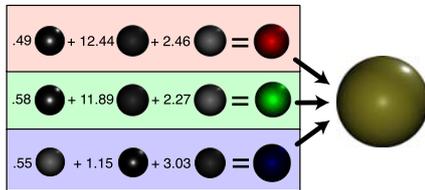


Figure 9: The secondary basis BRDFs from Matusik et al. [MPBM03a] are combined to form the primary linear basis functions (only those few secondary basis functions with non-zero weights from our factorization of the *bird* dataset are shown here).

gle location, but cover a wide range of normals in a single image. For these datasets, each row of Z contains approximately eight reflectance measurements on average.

Although ACLS allows low confidence values for undersampled regions of Z , significant interpolation is still necessary for this extremely sparse data. We reconstruct each row of Z at the same angular resolution reported in [LBAD*06] of $100 \times 30 \times 15$ ($\theta_h \times \theta_d \times \phi_d$) using the push-pull algorithm [GGSC96] and a Gaussian reconstruction kernel (note that Z therefore has 135,000 columns). The width of the reconstruction kernel was carefully selected to prevent over-smoothing while providing complete coverage of the domain, although some smoothing was unavoidable.

We compare our decomposition algorithm, using both RBFs and the MERL-MIT BRDF database, to ACLS in terms of their respective running times, memory consumption, and visual and numerical error using a Dell Precision 390 with a 1.8GHz Core 2 Duo processor and 3GB of memory. We compare their performance as a function of the number of subsampled rows used to compute the factorization and report the RMS error for the weighted metric in Eq. 15 across the entire dataset.

	Pre	UpdateW	UpdateH
ACLS	$MP\bar{N}$	MK^3	PK^3
OURS	$ML\bar{N}$	$MKL\bar{N} + MK^3$	$ML^2(\bar{N}^2 + K^2) + K^3L^3$

Table 1: Asymptotic running time of the precomputation, update W step, and update H step for ACLS and our algorithm. (M =subsamped rows; P =columns in regular sampled intermediate matrix; \bar{N} =average number of angular measurements per spatial position; K =number of primary basis functions; L =number of secondary basis functions)

Table 1 lists the asymptotic running time of each major stage in these algorithms. The graph in Fig. 10 shows the memory consumption of each algorithm and Fig. 11 plots measured running times for each dataset, along with a breakdown of the time spent in each of the major phases: precomputation, factorization, and backprojection. Because ACLS must store and process a matrix with $P = 135,000$ columns, it exceeds the memory capacity of our machine at $\sim 1K$ rows and we do not report statistics beyond this limit. The running time and memory consumption of ACLS is clearly dominated by pre-computing and processing the enormous $M \times P$ intermediate matrix. On the other hand, our algorithm is bound by the solving a $(KL \times KL)$ QP problem during the update H step (recall that we cannot estimate the columns in H independently as can be done in ACLS). This bottleneck becomes apparent when both L and K are relatively large as is the case when fitting RBFs to the *angel* dataset (see Fig. 11).

The graph in Fig. 12 reports error in the resulting factorization as a function of input size for the different algorithms and Fig. 13 shows side-by-side comparisons to accompany these measurements. Because Euclidean distance is a crude approximation to perceptual differences the visual comparisons reveal more about the relative performance of these methods. These results clearly indicate that the secondary basis provides a direct trade-off between the flexibility (and accuracy) of the representation and the amount of data (and

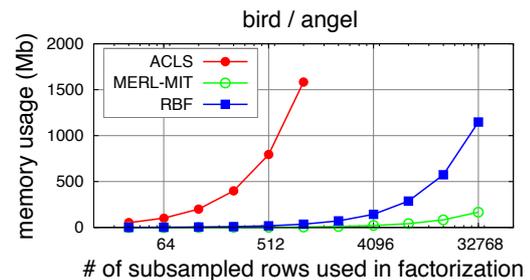


Figure 10: Memory usage of ACLS and our approach for two different secondary bases. This trend is dominated by the size of the interpolated data matrices for ACLS and by the A matrices in our algorithm (Eq. 4). Although these datasets use a different number of primary basis functions K , they exhibit nearly identical memory consumption.

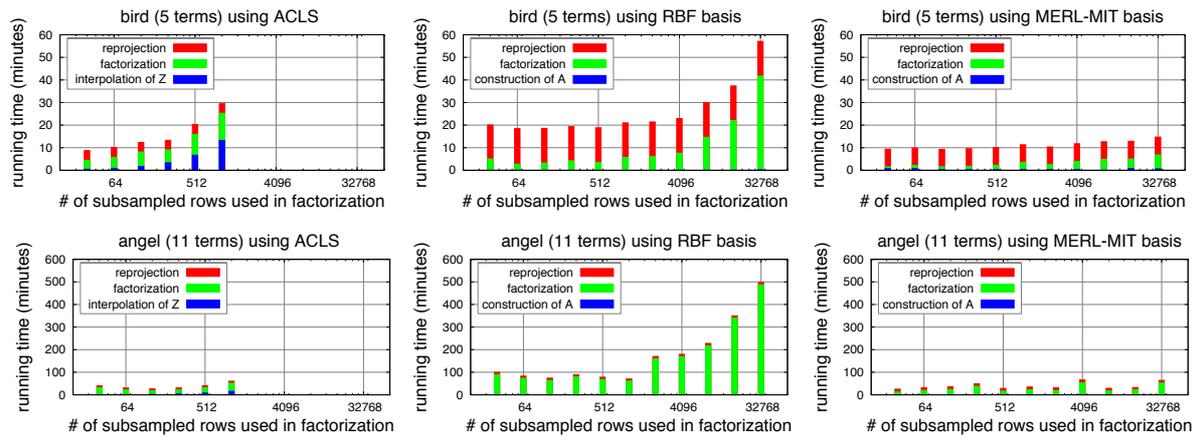


Figure 11: Running times for factoring the *bird* dataset into 5 terms as a function of the percentage of the input considered for different choices of secondary basis functions.

corresponding running times) that must be accommodated to make the optimization robust. RBFs are able to achieve an overall better match to measured data (in terms of numerical error) for increasingly large input sets, but at the cost of longer running times and greater memory consumption. In fact, ACLS can be regarded as the extreme point along this axis: it has enormous flexibility in its representation but requires a considerable amount of data to make this estimation robust and justify its expense. Not surprisingly, for this type of extremely sparse input it is not an appropriate representation and the significant interpolation required can attenuate or altogether miss important features in the data. This is most noticeable in the poorly recovered specular highlights for the *bird* model in Fig. 13. At the opposite extreme, a smaller and more restrictive secondary basis like the BRDF database performs very well for small input sizes and with less space and time requirements but clearly reaches a limit beyond which additional input samples will not affect the optimization.

7. Discussion and Future Work

We have presented a technique for efficiently performing basis decomposition from scattered appearance measurements. The key benefit of our approach is that it allows selecting a representation whose flexibility is appropriate for the nature and density of the available data. This is accomplished by representing each primary basis function as a weighted combination of a predetermined set of secondary basis functions. We introduced an iterative procedure that computes the optimal decomposition using quadratic programming methods that can accommodate additional linear constraints. Our results compare favorably with currently available methods such as ACLS for the case of decomposing and representing SVBRDFs, and we have analyzed the use of RBFs and a BRDF database as secondary bases. Here, we discuss some

existing limitations of our approach and propose further directions of study and applications of our work.

Broader Analysis: Further study of our proposed decomposition technique is warranted to more completely evaluate its applicability for representing more general scattered visual data. We would like to evaluate additional secondary bases (e.g., wavelets) for a variety of high-dimensional measured data beyond the SVBRDFs that we have described. In the future, we would like to assess our technique's effectiveness in representing BSSRDFs [PvBM*06], BTFs [DvNK99], or time-varying reflectance data [GTR*06]. Other such applications could include representing motion capture data, 3D surface data, and datasets outside computer graphics that involve scattered measurements in higher-dimensions.

Probabilistic Framework: There is a strong connection between our decomposition algorithm and the use of Expectation-Maximization (EM) [DLR77] to estimate the parameters of a probabilistic data model. Developing generative models for more general classes of material appearance along with methods for inferring their parameters from measured data is an interesting direction of future research.

Data-Driven Bases: We would like to further explore how data-driven reflectance models might be used to guide basis estimation of higher-dimensional appearance functions. For this approach to be applicable to a wider range of data, there must be increased availability of high-quality (low noise), densely sampled, and comprehensive datasets.

Acknowledgments

The authors wish to thank Dave Koller for his help with early drafts of the paper, Hendrik Lensch for generously sharing his reflectance data, and Wojciech Matusik for making the isotropic BRDF database available. We also acknowledge the helpful comments from the reviewers.

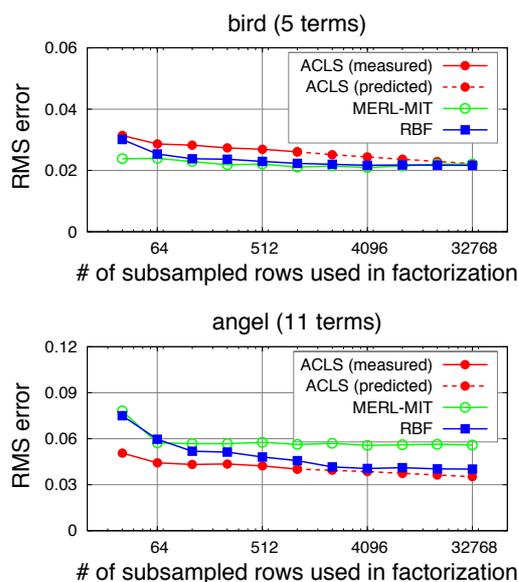


Figure 12: Numerical error of factorizations as a function of the input size and choice of secondary BRDF basis.

References

- [CBCG02] CHEN W.-C., BOUGUET J.-Y., CHU M. H., GRZESZCZUK R.: Light field mapping: Efficient representation and hardware rendering of surface light fields. *ACM Transactions on Graphics* 21, 3 (2002), 447–456.
- [DLR77] DEMPSTER A., LAIRD N., RUBIN D.: Maximum likelihood from incomplete data via the EM algorithm. *J. Royal Statistical Soc., Ser. B* 39, 1 (1977), 1–38.
- [DR82] DUFF I. S., REID J. K.: *MA27: A set of Fortran subroutines for solving sparse symmetric sets of linear equations*. Tech. Rep. AERE R10533, AERE Harwell Laboratory, London, England, 1982.
- [DvNK99] DANA K., VAN GINNEKEN B., NAYAR S., KOENDERINK J.: Reflectance and texture of real-world surfaces. *ACM Transactions on Graphics* 18, 1 (1999), 1–34.
- [FKIS02] FURUKAWA R., KAWASAKI H., IKEUCHI K., SAKAUCHI M.: Appearance based object modeling using texture database: acquisition, compression and rendering. In *Eurographics Workshop on Rendering* (2002), pp. 257–266.
- [GCHS05] GOLDMAN D. B., CURLESS B., HERTZMANN A., SEITZ S. M.: Shape and spatially-varying BRDFs from photometric stereo. In *IEEE International Conference on Computer Vision* (2005).
- [GGSC96] GORTLER S., GRZESZCZUK R., SZELISKI R., COHEN M.: The lumigraph. In *Proceedings of ACM SIGGRAPH 1996* (1996), pp. 43–54.
- [GLL*04] GOESELE M., LENSCH H. P. A., LANG J., FUCHS C., SEIDEL H.-P.: DISCO: Acquisition of translucent objects. *ACM Transactions on Graphics* 23, 3 (2004), 835–844.
- [GTHD03] GARDNER A., TCHOU C., HAWKINS T., DEBEVEC P.: Linear light source reflectometry. *ACM Transactions on Graphics* 22, 3 (2003), 749–758.
- [GTR*06] GU J., TU C., RAMAMOORTHY R., BELHUMEUR P., MATUSIK W., NAYAR S.: Time-varying surface appearance: Acquisition, modeling, and rendering. *ACM Transactions on Graphics (SIGGRAPH 2006)* 25, 3 (2006).
- [GW03] GERTZ E. M., WRIGHT S. J.: Object-oriented software for quadratic programming. *ACM Transactions on Mathematical Software* 29, 1 (2003), 58–81.
- [KM99] KAUTZ J., MCCOOL M.: Interactive rendering with arbitrary BRDFs using separable approximations. In *Eurographics Workshop on Rendering* (1999), pp. 247–260.
- [KvDS96] KOENDERINK J., VAN DOORN A., STAVRIDIS M.: Bidirectional reflection distribution function expressed in terms of surface scattering modes. In *Computer Vision—ECCV '96* (1996), Springer Verlag, pp. 28–39.
- [LBAD*06] LAWRENCE J., BEN-ARTZI A., DECORO C., MATUSIK W., PFISTER H., RAMAMOORTHY R., RUSINKIEWICZ S.: Inverse shade trees for non-parametric material representation and editing. *ACM Transactions on Graphics (SIGGRAPH 2006)* 25, 3 (2006).
- [LFTG97] LAFORTUNE E. P. F., FOO S.-C., TORRANCE K. E., GREENBERG D. P.: Non-linear approximation of reflectance functions. In *Proceedings of ACM SIGGRAPH 1997* (1997), pp. 117–126.
- [LKG*01] LENSCH H. P. A., KAUTZ J., GOESELE M., HEIDRICH W., SEIDEL H.-P.: Image-based reconstruction of spatially varying materials. In *Eurographics Workshop on Rendering* (2001), pp. 63–70.
- [LKG*03] LENSCH H. P. A., KAUTZ J., GOESELE M., HEIDRICH W., SEIDEL H.-P.: Image-based reconstruction of spatial appearance and geometric detail. *ACM Transactions on Graphics* 22, 2 (2003), 234–257.
- [LM01] LEUNG T., MALIK J.: Representing and recognizing the visual appearance of materials using three-dimensional textons. *International Journal of Computer Vision* 43, 1 (2001), 29–44.
- [LRR04] LAWRENCE J., RUSINKIEWICZ S., RAMAMOORTHY R.: Efficient BRDF importance sampling using a factored representation. *ACM Transactions on Graphics (SIGGRAPH 2004)* 23, 3 (2004).
- [MAA01] MCCOOL M. D., ANG J., AHMAD A.: Homomorphic factorization of BRDFs for high-performance rendering. In *Proceedings of ACM SIGGRAPH 2001* (2001), pp. 185–194.

- [McA02] MCALLISTER D.: *A Generalized Surface Appearance Representation for Computer Graphics*. PhD thesis, UNC Chapel Hill, 2002.
- [MMS*04] MÜLLER G., MESETH J., SATTLER M., SARLETTE R., KLEIN R.: Acquisition, synthesis and rendering of bidirectional texture functions. In *Eurographics 2004, State of the Art Reports* (September 2004), INRIA and Eurographics Association, pp. 69–94.
- [MPBM03a] MATUSIK W., PFISTER H., BRAND M., MCMILLAN L.: A data-driven reflectance model. *ACM Transactions on Graphics (SIGGRAPH 2003)* 22, 3 (2003), 759–769.
- [MPBM03b] MATUSIK W., PFISTER H., BRAND M., MCMILLAN L.: Efficient isotropic BRDF measurement. In *Eurographics Workshop on Rendering* (2003), pp. 241–247.
- [MWAM05] MARSCHNER S. R., WESTIN S. H., ARBREE A., MOON J. T.: Measuring and modeling the appearance of finished wood. *ACM Transactions on Graphics (ACM SIGGRAPH 2005)* 24, 3 (2005), 727–734.
- [NDM05] NGAN A., DURAND F., MATUSIK W.: Experimental analysis of BRDF models. In *Proceedings of the Eurographics Symposium on Rendering* (2005), Eurographics Association, pp. 117–226.
- [NRH77] NICODEMUS F. E., RICHMOND J. C., HSIA J. J.: Geometrical considerations and reflectance. *National Bureau of Standards* (October 1977).
- [ON94] OREN M., NAYAR S. K.: Generalization of Lambert's reflectance model. In *Computer Graphics (SIGGRAPH 1994)* (1994), pp. 239–246.
- [PH04] PHARR M., HUMPHREYS G.: *Physically Based Rendering : From Theory to Implementation*. Morgan Kaufmann, 2004.
- [Pow92] POWELL M.: The theory of radial basis function approximation in 1990. In *Advances in Numerical Analysis II: Wavelets, Subdivision Algorithms and Radial Functions* (Oxford, U.K., 1992), Light W., (Ed.), Oxford Science Publications, pp. 105–210.
- [PvBM*06] PEERS P., VOM BERGE K., MATUSIK W., RAMAMOORTHY R., LAWRENCE J., RUSINKIEWICZ S., DUTRÉ P.: A compact factored representation of heterogeneous subsurface scattering. *ACM Transactions on Graphics (SIGGRAPH 2006)* 25, 3 (2006).
- [Rus98] RUSINKIEWICZ S.: A new change of variables for efficient BRDF representation. In *Eurographics Workshop on Rendering* (1998), pp. 11–22.
- [TJBB04] TEH Y., JORDAN M., BEAL M., BLEI D.: Hierarchical Dirichlet processes. *Journal of the American Statistical Association* (2004), to appear.
- [TOS*03] TSUMURA N., OJIMA N., SATO K., SHIRAISHI M., SHIMIZU H., NABESHIMA H., AKAZAKI S., HORI K., MIYAKE Y.: Image-based skin color and texture analysis/synthesis by extracting hemoglobin and melanin information in the skin. *ACM Transactions on Graphics (SIGGRAPH 2003)* 22, 3 (2003), 770–779.
- [TS66] TORRANCE K. E., SPARROW E. M.: Off-specular peaks in the directional distribution of reflected thermal radiation. *Transactions of the ASME* (1966).
- [VT04] VASILESCU M. A., TERZOPOULOS D.: TensorTextures: Multilinear image-based rendering. *ACM Transactions on Graphics (SIGGRAPH 2004)* 23, 3 (2004).
- [War92] WARD G. J.: Measuring and modeling anisotropic reflection. In *Computer Graphics (Proceedings of ACM SIGGRAPH 92)* (1992), pp. 265–272.
- [ZERB05] ZICKLER T., ENRIQUE S., RAMAMOORTHY R., BELHUMEUR P.: Reflectance sharing: Image-based rendering from a sparse set of images. In *Proceedings of the Eurographics Symposium on Rendering* (2005), Eurographics Association, pp. 253–264.
- [ZERB06] ZICKLER T., ENRIQUE S., RAMAMOORTHY R., BELHUMEUR P.: Reflectance sharing: Predicting appearance from a sparse set of images of a known shape. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 28, 8 (2006), 1287–1302.

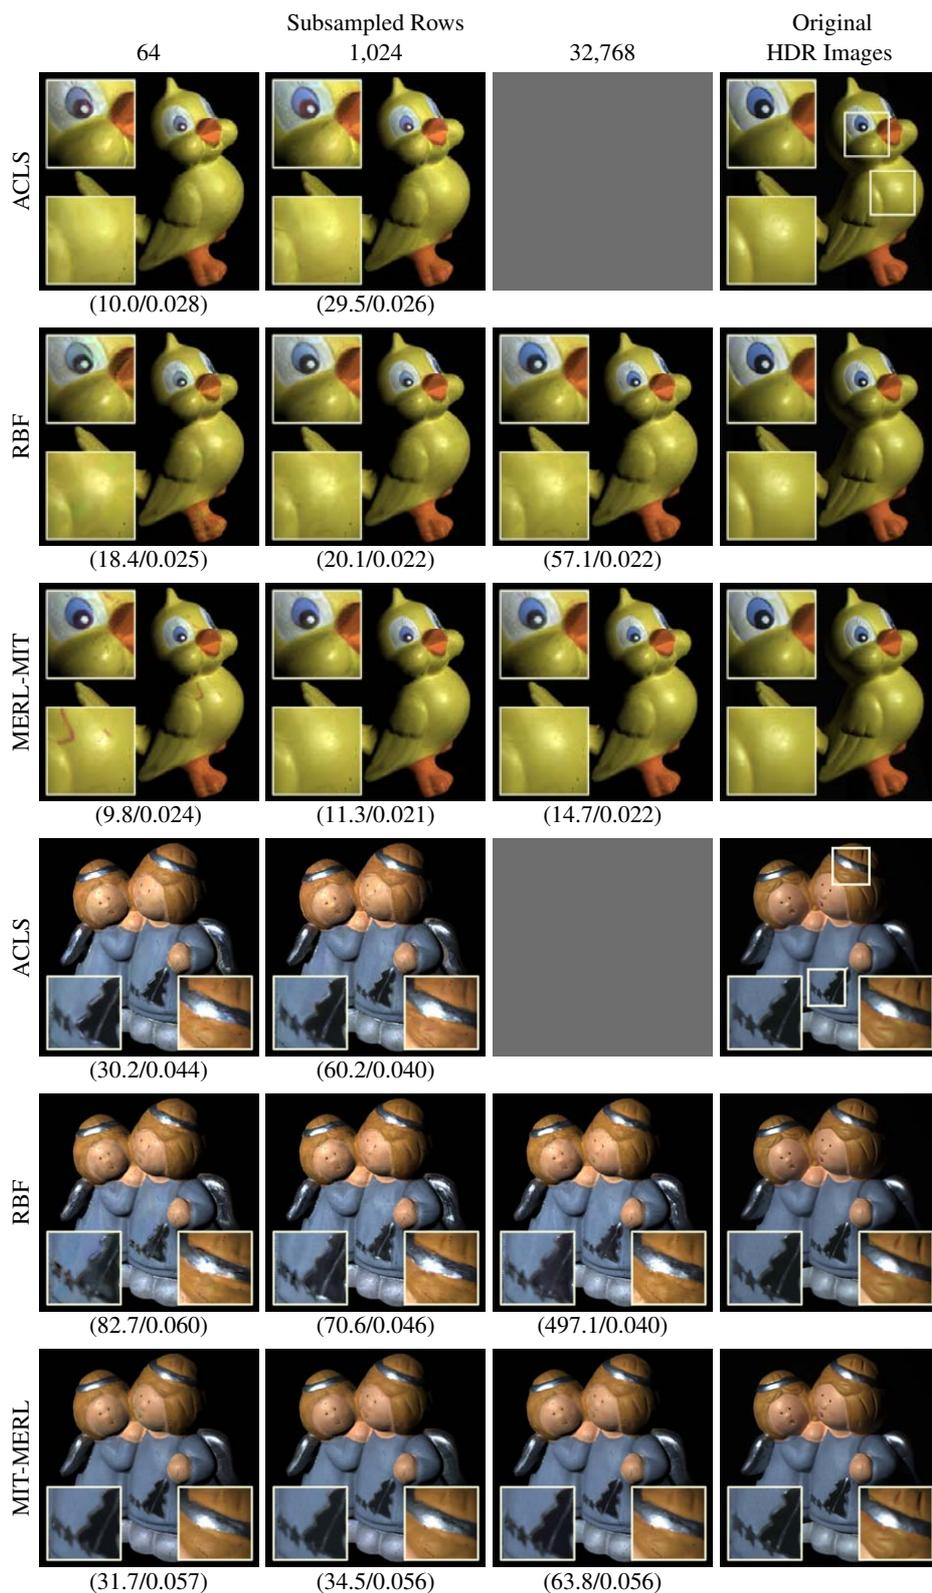


Figure 13: Visual comparison to ACLS for two secondary bases applied to the *bird* and *angel* datasets. Below each image we report (running time in minutes / RMS error). For 32,768 rows, the space and time requirements of ACLS exceed available resources by an order of magnitude and were not included (shown as gray boxes).