Image-Based Relighting using Room Lighting Basis

Antoine Toisoul
Imperial College London
ast13@imperial.ac.uk

Abhijeet Ghosh
Imperial College London
ghosh@imperial.ac.uk

ABSTRACT

We present a novel and practical approach for image-based relighting that employs the lights available in a regular room to acquire the reflectance field of an object. The lighting basis includes diverse light sources such as the house lights and the natural illumination coming from the windows. Once the data is captured, we homogenize the reflectance field to take into account the variety of light source colours to minimise the tone difference in the reflectance field. Additionally, we measure the room dark level corresponding to a small amount of global illumination with all lights switched off and blinds drawn. The dark level, due to some light leakage through the blinds, is removed from the individual local lighting basis conditions and employed as an additional global lighting basis. Finally we optimize the projection of a desired lighting environment on to our room lighting basis to get a close approximation of the environment with our sparse lighting basis. We achieve plausible results for diffuse and glossy objects that are qualitatively similar to results produced with dense sampling of the reflectance field including using a light stage and we demonstrate effective relighting results in two different room configurations. We believe our approach can be applied for practical relighting applications with general studio lighting.

CCS Concepts

•**Computing methodologies** → Image-based rendering;

Keywords

Image-based relighting; Rendering

1. INTRODUCTION

Measurement-based methods are a popular topic in computer graphics and have been extensively used for accurate modelling of geometry or appearance for photoreaslitic renderings. Image-based relighting is a well known example of such a technique that relies on captured data to produce realistic renderings. The approach takes as an input a set of basis images of an object captured under known illumination and computes a linear combination of these to render the object in an arbitrary lighting environment. The basis images correspond to the reflectance field of the object, i.e., slices of an eight dimensional function that captures how an object surface reflects light from any incoming light directions towards any given viewing direction.

Conventionally, the reflectance field is captured by sequentially illuminating the object from known directions over the whole sphere and by recording the reflected radiance for each of these directions towards a camera. In order to capture a reflectance field correctly, a dense sampling of the incoming light directions is usually required. As a result controlled measurement setups are commonly used to get fast and accurate measurements. An example of such a setup is the light stage [4] where typically a dome or a rotating arc of light sources allow a camera to automatically record the outgoing radiance of an object for any of the sampled lighting directions. Moreover, free-form acquisition with a hand-held light source has also been employed for ease of reflectance field acquisition [9]. Both of these approaches work well, but also have their drawbacks. A hand-held acquisition takes time and capturing a dense sampling of the incoming light directions is difficult. However a highly specialised device such as a light stage is often very expensive and requires a laboratory setup to perform the capture.

In this work, we propose a novel and practical method to capture the reflectance field of an object that only relies on the lighting conditions available in a regular room. The reflectance field that is captured does not recover the entire surface reflectance but corresponds to a sparse sampling of the reflectance function described by Debevec et al. [4]. We employ all the available light sources in the room (e.g., windows and house lights) as a lighting basis for our image-based relighting approach. Such a lighting basis has two main characteristics that make standard image-based relighting algorithms produce poor results. First, it is heterogeneous as each light source is different. Indeed in a regular room the light sources are non uniformly distributed, have different solid angles and colours and objects and walls diffuse the light in the room adding a bit of indirect illumination. Such specificities do not exist in a laboratory measurement setup. We propose a preprocessing step on the measured reflectance field to tackle these issues. Secondly, our lighting basis is sparse as it typically only contains a few (nine to twelve in our experiments) lighting conditions

(a) (b) (c) (d)

Figure 1: **(a)** Room lighting basis. **(b), (c)** Decorative objects relit with room lighting: **(b)** a wooden Bird object in the Pisa Courtyard, and **(c)** a decorative Egg in the Grace Cathedral. **(d)** Relighting with dense free form sampling according to Masselus et al. [9]

.

compared to several hundred in a typical light stage dataset [1]. Hence, the direct projection of a desired lighting environment on to a room lighting basis tends to produce a suboptimal approximation. We tackle this issue by employing an optimisation process to obtain a linear combination of the projection that closely approximates the desired lighting environment. Such a method could be employed in practice for image-based relighting applications using typical studio lighting setups without requiring a highly specialized setup such as a light stage.

The rest of the paper is organized as follows: in Section 2 we first review some related works before describing our approach and method in Sections 3 and 4. In Section 5, we compare our results to dense sampling of the reflectance field with free-form acquisition according to the method of Masselus et al. [9] and show that the method can be used to produce convincing renderings with a sparse room lighting basis. We also compare our technique using simulations to relighting results obtained with a light stage. Finally in Section 6 we present relighting results achieved in another room with a different configuration of house lights and windows to demonstrate the generality and flexibility of the approach.

2. RELATED WORK

In image-based relighting the illumination in a scene is modelled by the product between the light transport matrix and the illumination condition vector [12]. The light transport matrix describes how each pixel of an image behaves given a lighting condition. Hence it maps an input lighting condition to an image. Measuring the entries of the light transport matrix has been a very active area of research. We present literature on controlled and uncontrolled measurement setups as well as basis representations.

Controlled measurement setup. The idea of relighting using a linear combination of basis images goes back to Nimeroff et al. [13]. However, it became practical only in 2000 when Debevec et al. [4] defined the concept of a reflectance field and proposed the usage of a light stage to acquire it. They also presented realistic renderings of faces in arbitrary environments as well as novel viewpoint renderings. The algorithms and light stage itself have improved. In 2004, Hawkins et al. [7] proposed a new light stage made of an arc of lights. It measures the reflectance field in eight seconds and allow a multiview capture as well as the capture of several facial expressions. Light stage 3 [5] was the first dome of RGB light sources that allowed live performance to be captured and composited in a virtual environment. Light stage 5 [17] allowed the reflectance of a live performance of an actor to be captured and modified in post production using high-speed cameras. In 2007, Peers et al. [14] used a reflectance transfer approach to practically relit faces in post-production. They use a database of viewpoint dependent faces relit in the desired lighting environment and capture a new actor under known illumination. They then transfer the reflectance field of the known actor onto the new actor using the quotient image.

In this paper our approach does not require any special device to capture the reflectance field and belongs to the category of uncontrolled measurement setups. It is a cheap and practical technique that can be used by anyone.

Uncontrolled measurement setup. In 2002, Masselus et al. [9] showed that image-based relighting can be achieved using a dense sampling of the incoming light directions with a hand-held light source. Each picture in the reflectance field contains four diffuse spheres used to estimate the incoming light direction automatically. Besides the non-uniform sampling is solved by computing appropriate weights for each direction based on a Voronoi diagram. When projected on the new lighting environment each Voronoi cell contains the pixels that are the closest to a given direction. Recently in 2015, Ren et al. [15] used machine learning and neural networks as a tool for image-based relighting. The main idea is that the light transport matrix can be approximated in segments from a small basis of input images due a non linear local coherence in the matrix. Their method can render complex objects under any arbitrary illumination and only requires a small set of images which can be acquired with free-form acquisition. While not as powerful as this approach in terms of resolving complex reflectance functions, our proposed method has its advantages of simplicity and requires no prior training data.

Relighting and basis representation. In 2003, Masselus et al. [10] proposed a technique to relight objects with 4D incident light fields instead of regular environment maps. It allows the rendering of effects that are only captured with light fields such as shadows and spot lights. However the measurement setup restricts lighting to a narrow solid angle and does not cover the full sphere of directions. In 2004, Masselus et al. [11] proposed a compact and smooth representation of the reflectance function using different basis functions. Their method is fast to compute and preserves high frequency variations but still requires a significant number of basis coefficients for high quality results. More recently, Tunwattanapong et al. [16] proposed a method combining spherical harmonics and local lights to reduce the number of lighting conditions that are required for image-based relighting. The basis is chosen to capture both the low frequency lighting with the spherical harmonics and the high frequencies with the local lights. They employ an optimisation to find the weights that best approximate the target environment map given the lighting basis. Their technique provides good results with only twenty lighting conditions. Our method employs a similar optimisation procedure (albeit with some differences such as not employing PCA) to improve the approximation of the target environment in our sparse lighting basis. Although their approach uses few lighting conditions, it still requires an LED sphere lighting device to illuminate a subject with spherical harmonic illumination. In contrast, our approach is well suited to general studio lighting which is much more commonly accessible.

3. REFLECTANCE FIELD CAPTURE

This section discusses the specific features of capturing a reflectance field with a regular room as well as the preprocessing step that is required before computing the relighting algorithm.

3.1 Lighting basis

Any room can be chosen to create the lighting basis as long as the light sources can be individually controlled. In this section, we present the capture of the reflectance field with a regular office room as a first example. We later discuss measurements with a another room in section 6.1. Our first example, an office, has two broad windows and two sets of house lights making a total of four lighting conditions which is usually not enough to produce correct relighting. As a result, we increase the number of illumination conditions by using the windows full opened and half opened (blind halfway closed) as two separate illumination conditions. The half opened window condition is then subtracted from the fully opened window condition in a post-process in order to avoid overlaps in the lighting basis and to obtain smaller window lights. The office room did not have any windows to the sides. Hence, we added two more lighting conditions from the sides by using a smartphone flashlight (iPhone 5S). Finally the room with the blinds drawn on the windows and the house lights switched off was also used as a lighting condition (referred to as the dark room condition). In the end our lighting basis has nine different illumination conditions including the dark room. We record the radiance values of an object under these nine illumination conditions with high-dynamic range (HDR) photographs [6]. Each lighting configuration is also captured with a HDR light probe as shown in figure 1a. and will be used as a measurement of the incident light direction in the relighting phase.

Such a lighting basis contains wide area light sources and light sources of different kind. As a result the captured data has to be preprocessed before being used in a regular image-based relighting pipeline.

3.2 Preprocessing the data

The lighting basis that we use for relighting is complex compared to the standard basis acquired with a light stage. Three main differences have to be noticed : the ambient illumination in the room when no light is switched on, the overlapping illumination conditions and the different light colours. We deal with these three differences in this section.

The room used for the capture was never completely dark when the windows were covered with the blinds and the house lights were switched off. This ambient illumination is not negligible as the objects in the room and especially the white interior walls cause diffuse interreflection of any light bleeding into the room from gaps in the window blinds. As a result an illumination factor is added to every picture in the reflectance field that we captured. This illumination factor is removed from the reflectance field by subtracting the dark room measurement from each photograph in the reflectance field.

The lighting basis that is considered is also not orthogonal as certain lighting conditions overlap. For instance, this is the case of the windows half-opened and fully opened. If these two conditions were directly used in the relighting algorithm the final rendering would not conserve the energy as half of the window would be used twice in the linear combination. Hence, we subtract the half-opened condition from the fully opened condition to replace the two overlapping conditions with two non overlapping half opened windows (top half of the window and bottom half in our case). This process is applied to the corresponding photographs of the reflectance field as well as the light probe measurements to create orthogonal basis conditions.

Finally, the light sources in the lighting basis have different colours (figure 1a.). For instance the house lights are fluorescent with an orange colour compared to the daylight coming in from the windows. This effect is removed by performing white balancing so that every light in the room has the same colour temperature. This is done by taking two pictures of the white diffuse wall in the room : one with the house lights only and one with the windows open only (natural illumination is taken as a reference). Then we calculate the average RGB colour of the two images, denoted $(R_{win}, G_{win}, B_{win})$ and (R_{hl}, G_{hl}, B_{hl}) for the window and the house light illumination respectively. We finally calculate a RGB scaling factor $(\alpha_R, \alpha_G, \alpha_B)$ and multiply each color channel of the pictures taken under the house lights by this factor. We perform similar white balancing for data acquired with the phone flash.

$$(\alpha_R, \alpha_G, \alpha_B) = (\frac{R_{win}}{R_{hl}}, \frac{G_{win}}{G_{hl}}, \frac{B_{win}}{B_{hl}}) \qquad (1)$$

This last step makes the photographs of the reflectance field have a similar overall colour that matches the data captured with the window condition i.e under natural illumination. After these three steps, the reflectance field is ready to be used in a regular image-based relighting pipeline.

4. RELIGHTING

4.1 Environment map partition

Regular image-based relighting is a three steps algorithm. First each of the N lighting conditions in the lighting basis is associated with a direction using the light probe measurements. Secondly each light direction is mapped as a point on a latitude longitude map of the desired new environment. Then a Voronoi diagram is computed from these directions [9]. A Voronoi diagram is made of cells, referred to as Voronoi cells, that cover the entire latitude longitude map. Each cell represents a unique unique light direction. As a result, each pixel in the target environment falls in a cell that contains all the pixels that are the closest to that specific direction. Hence a cell can be seen as the solid angle of how much a light source covers on the sphere of directions. Finally each Voronoi cell is integrated by summing all the pixels that belong to the cell. The result is a set of RGB weights $(\omega_i)_{1\leq i\leq N}$ that are associated to pictures of the object taken under the N conditions (reflectance field). In our case, we have $N = 9$ lighting conditions in the office room, hence nine Voronoi cells in our diagram and nine weights RGB $(\omega_i)_{1\leq i\leq 9}$. The integration for cell i is given by equation where p is a pixel in cell i (equation 2).

$$\omega_i = (\omega_{i,R}, \omega_{i,G}, \omega_{i,B}) = \sum_{p\in cell_i} (p_R, p_G, p_B) \qquad (2)$$

After the integration the final relighting is computed by a linear combination of the pictures in the reflectance field $(X_i)_{1\leq i\leq N}$ (equation 3). Note that the product $\omega_i X_i$ in the sum is done for each color channel.

$$X = \sum_{i=1}^{N} \omega_i X_i \qquad (3)$$

We adapt the first and second step to our specific lighting basis by using the center of a large light source as the light direction. The Voronoi diagram can then be computed using these directions. However computing such a diagram with a sparse lighting basis leads to large Voronoi cells. Indeed the smaller the sampling of incoming light directions in the reflectance field, the lower the number of cells and the larger the size of the cells as they cover the entire environment map. As a result, during the integration step, large cells associate radiance from far away directions with a given light source. This can lead to inaccurate relighting results. In order to avoid such inaccuracies, we partition the environment by making masks that cover the solid angle of a given light source. These masks correspond to our area of integration for each light source. Note that the size of the masks can be chosen inversely proportional to the solid angle of a light source. That way, wide light sources have smaller solid angles which compensates for the higher energy in the reflectance field. Finally the non-covered area of the environment (marked in gray) is associated with the dark room condition. The final projection of a target environment on to the lighting basis is shown in figure 2.

4.2 Optimisation

Such a sparse lighting basis gives a suboptimal approximation of the target environment after integration of the cells. In order to improve the approximation we solve a convex optimisation problem. The optimisation procedure associates an intensity scaling factor to each lighting condition and uses a constrained optimisation (the scaling factors have to be positive) in intensity space to find the scaling factors that best approximate the environment map given the lighting basis. The intuitive idea is to reduce and increase the importance of a lighting condition in the lighting basis so that the overall approximation of the environment map is improved.

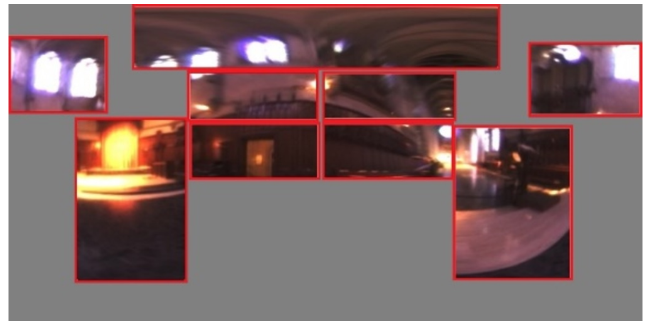

Figure 2: Lighting basis projected on the Grace Cathedral environment. The grey area is associated to the dark room condition.

In formal terms, let y denote a one dimensional vector of the environment map intensities (average of the RGB color channels). This vector has a dimension n (1024×512 in our case) corresponding to the pixels of the environment map represented as a single column. Let A denote the $n \times N$ matrix ($N = 9$ is the number of illumination conditions in the lighting basis) that projects the integrated weights $(\omega_i)_{1\leq i\leq N}$ on our lighting basis. Each row in A corresponds to one pixel in the environment map and is made of zeros in every column but one. The column j that has a non zero value, has a value $I_j = \frac{(\omega_{j,R}+\omega_{j,G}+\omega_{j,B})}{3}$ where j is the index of lighting condition the pixel belongs to. The vector x is a column vector of dimension $N = 9$ that contains the nine scaling factors to be optimised. The final formulation of the optimisation problem is given in equation 4.2. We computed it in C++ using the function *find_min_box_constrained* of Dlib library [8].

$$\begin{aligned} \min_x \quad & \|Ax - y\|_2 \\ \text{subject to} \quad & \forall i \in 1,..,9 \;\; x_i > 0 \end{aligned}$$

The solution to this problem is a vector denoted x_{opt}. We then multiply the original weights $(\omega_i)_{1\leq i\leq N}$ by their corresponding scaling factor in x_{opt} before the relighting. A comparison of the results with and without the optimisation is shown in figure 3. Such an optimisation procedure is similar to the work of Tunwattanapong et al. [16]. However, they optimise in the space given by the principal component analysis of the projection matrix A, whereas our optimisation is done in original space for convexity guarantees.

A pseudo code that explains how to compute the objective function is given in algorithm 1. The input variables are EM an image of the target environment map, ω an RGB vector that contains the weights after integration and x the vector of scaling factors. The function *find_lighting_condition_of_pixel* returns the index of the lighting condition the input pixel is associated to. Note that the optimisation is convex due to the convexity of the linear least squares problem. As a result, any initialisation of the x vector will lead to a global minimum.

5. RESULTS

5.1 Comparison to dense free-form sampling

We present relightings in several environments [2] for two

(a)

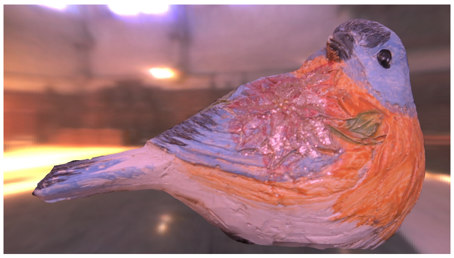

(b)

Figure 3: Relighting of the bird in the Grace Cathedral. **(a)** Without optimisation **(b)** With optimisation

Data: Image(width,height) EM, RGBVector[9] w, Vector[9] x
Result: Computation of the objective function
sum = 0.0;
for *k in range(0,width)* **do**
  for *l in range(0,height)* **do**
    pixelAvg = average color channel EM(k,l);
    j = find_lighting_condition_of_pixel(k,l)
    I = average color weight w[j]
    sum += $(Ix[j] - pixelAvg)^2$
  end
end
Output: $\sqrt{sum}$

Algorithm 1: How to compute the objective function

decorative figurines (a bird and an egg). The results are compared to the method of Masselus et al. [9] that uses a dense hand-held acquisition of the reflectance field with a smartphone flashlight (iPhone 5S). Our lighting basis for this dense capture contains 142 lighting conditions that cover important areas of the sphere around the object. The light directions and the corresponding Voronoi diagram are plotted on a latitude longitude map on figure 4. Our results, shown in figure 5 look plausible and are comparable to the one obtained with the dense acquisition. Additional relightings are presented in section 6. Note that some differences in the colour tone are due to the fact that high-dynamic range [6] data was employed for the room relighting basis while low dynamic range imaging had to be employed for the free-form acquisition. This also highlights an advantage of our approach over that of Masselus et al. for acquiring specular reflectance functions which tend to have a large dynamic range.

5.2 Comparison to light stage relighting

No light stage reflectance field data is available for the egg and bird objects shown in the previous section. Hence a direct comparison between the room relighting and the light stage relighting is not possible for these objects. However USC ICT [1] captured a few objects with their light stage 6 and made them available online. As a result, a comparison between the room relighting and the light stage relighting can be made on these objects using a two step simulation. We first used the light stage data to relight a plant dataset and a helmet dataset in our office room lighting basis conditions using our captured light probes. Here, we start from 253 pictures of the objects taken with the light stage illumination and end with nine relit results in our office room lighting basis. These relit results correspond to our simulated reflectance field data. Then during a second step we compute the room relighting algorithm and compare it to a direct relighting with the densely sampled light stage data.

A few results are shown in figures 6 and 7. Although the overall results differ slightly, the room relighting successfully captures some of the directional illumination. Additional results can be found in supplementary material. These results verify the plausibility of the relit results given the sparse room lighting basis conditions.

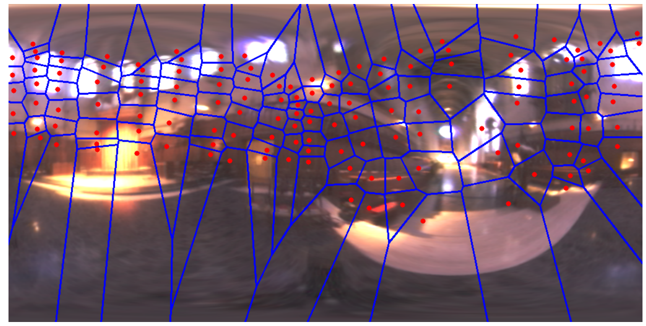

Figure 4: Latitude longitude map of the Grace Cathedral with the 142 lighting conditions of the dense acquisition plotted as red dots. The corresponding Voronoi diagram is shown in blue.

6. ADDITIONAL RESULTS

Additional relighting results of the bird figurine in the office room lighting basis are presented in figure 8. We also show additional comparisons of the office room relighting to dense free-form acquisition for the egg in figure 9.

6.1 Relighting with a different room lighting basis

We applied our technique in another room with a different lighting configuration to demonstrate its practicality and flexibility. This room is a bedroom that has one set of house lights and five windows. The windows have different shapes and the solid angle they cover vary. Indeed the first, fourth and fifth windows are broader than the second and third windows (see figure 10). We again capture the reflectance field by splitting the window illumination into two halves (fully and half opened) and apply the same preprocessing step as explained in section 3. The final lighting basis has twelve non-overlapping lighting conditions. The projection of the lighting basis is shown in figure 11. Given the greater distribution of windows in the bedroom, we did not add any flash lighting conditions in this case. Renderings of the egg with this lighting basis are shown in figure 12. Although the basis has a different layout of lights, we still achieve very

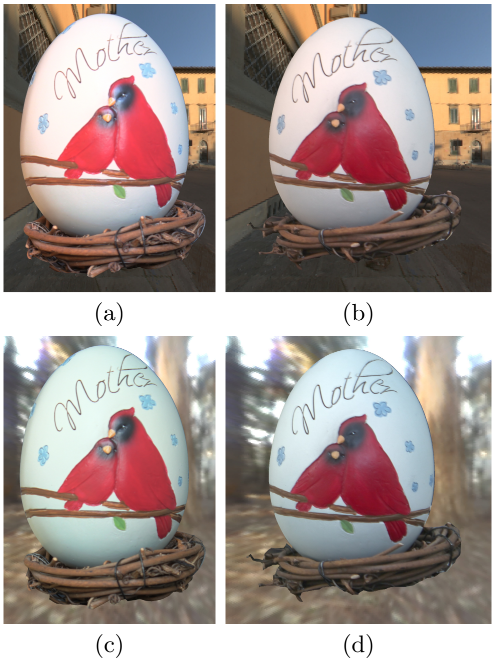

Figure 5: Comparison of the room relighting and the relighting using dense free-form acquisition of the reflectance field in the Pisa Couryard (top) and the Eucalyptus Grove (bottom). **(Left column)** Room relighting. **(Right column)** Method of Masselus et al. [9]

plausible results.

7. LIMITATIONS

Our technique has several limitations. First it is dependent on the room used as the number of independent light sources is very important with a sparse lighting basis. In practice, it is usually difficult to find more than five or six independent light sources in a regular room. As a result it is important to divide wide light sources into several smaller ones as we did for the windows.

Secondly symmetry in the lighting basis (around the viewing direction) plays a very important role. Indeed a non-symmetric lighting basis tends to give too much importance to the part of the environment whose light sources cover the biggest solid angle. We captured the reflectance field in our second room (see section 6.1) at a 45 degrees orientation from the first capture viewpoint. This causes the lighting basis to mostly cover the left part of the environment and not the right part (see figure 13). As a result a relighting in the Grace Cathedral leads to an incorrect (biased) result as the left part of the light probe is overfitted. Hence the viewpoint (orientation) to capture an object should be chosen in such a way that the light sources cover an equal area on each side of the light probe.

Thirdly, our technique relies on natural illumination that is time dependent. Indeed in our rooms, the incident light from the window varied quickly depending on the weather. As an HDR capture of a reflectance field takes time, changes of illumination between pictures can sometimes be observed. This can affect the reflectance field which is why the capture has to be done quickly. Besides the measurements cannot be made for animated objects due to the time it takes to switch from one light condition to another.

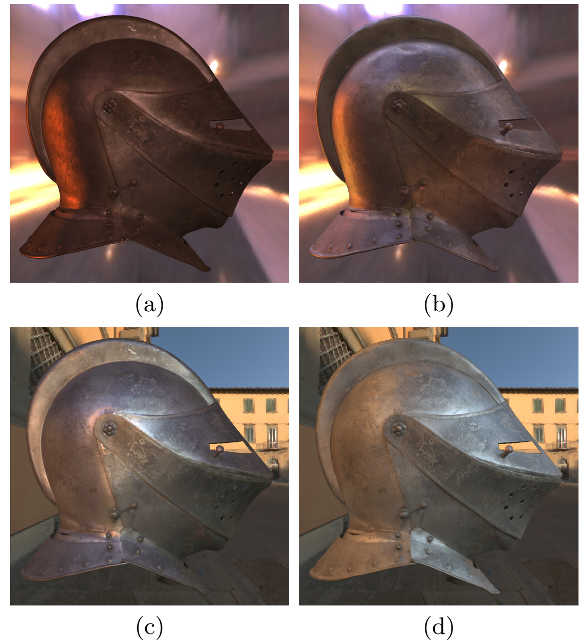

Figure 6: Comparison of the relighting of the helmet in the Grace Cathedral (top) and the Pisa Courtyard (bottom) environments **Left column :** Room relighting. **Right column :** Light stage relighting.

Finally we found that the results tend to be better for environment maps with low frequency lighting. Indeed with high frequency lighting, the angular changes in illumination are hard to capture correctly with a sparse lighting basis and large integration areas. Despite these limitations, the relighting results are plausible even for high frequency lighting environments such as the Grace Cathedral.

8. CONCLUSION

We presented a novel approach to capture the reflectance field of an object using the light sources available in a regular room. The data is not directly usable and has to be preprocessed to take into account the room dark level, the diversity and distribution of light sources and the sparsity of the lighting basis. A direct projection based relighting gives suboptimal results that can be improved using an optimisation method. While the results are qualitatively good for glossy and diffuse objects and compare well to relighting produced with dense acquisition of the reflectance field, quantitatively the error is dependent on the position and solid angle of the light sources in the room. However we believe that such a simple and cheap method could be effectively employed with studio lighting conditions which are more commonly accessible in practice than a controlled light stage setup.

9. ACKNOWLEDGMENTS

This work was partly supported by a Royal Society Wolfson Research Merit Award and an EPSRC Early Career Fellowship EP/N006259/1.

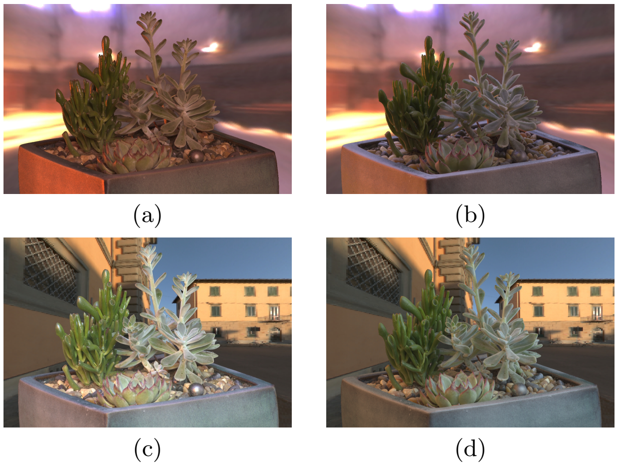

Figure 7: Comparison of the relighting of the plant in the Grace Cathedral (top) and the Pisa Courtyard (bottom) environments **Left column :** Room relighting. **Right column :** Light stage relighting.

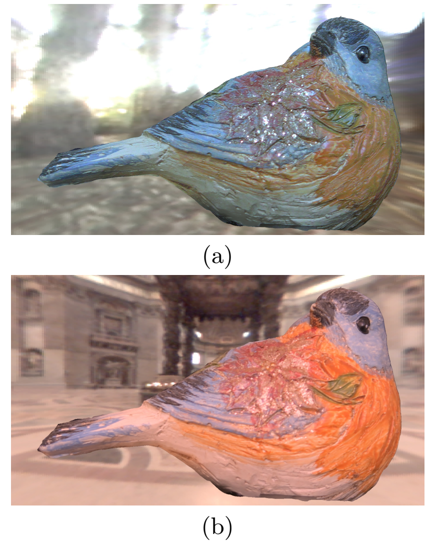

Figure 8: Relighting of the bird in two lighting environments. **(a)** Eucalyptus Grove. **(b)** St. Peter's basilica.

10. REFERENCES

[1] P. Debevec. Light stage data gallery. http://gl.ict.usc.edu/Data/LightStage/.

[2] P. Debevec. High-resolution light probe image gallery. http://gl.ict.usc.edu/Data/HighResProbes/, 2001.

[3] P. Debevec. The light stages and their applications to photoreal digital actors. *SIGGRAPH Asia Technical Briefs*, 2, 2012.

[4] P. Debevec, T. Hawkins, C. Tchou, H.-P. Duiker, W. Sarokin, and M. Sagar. Acquiring the reflectance field of a human face. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '00, pages 145–156, New York, NY, USA, 2000. ACM Press/Addison-Wesley Publishing Co.

[5] P. Debevec, A. Wenger, C. Tchou, A. Gardner, J. Waese, and T. Hawkins. A lighting reproduction approach to live-action compositing. *ACM Trans. Graph.*, 21(3):547–556, July 2002.

[6] P. E. Debevec and J. Malik. Recovering high dynamic range radiance maps from photographs. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '97, pages 369–378, New York, NY, USA, 1997. ACM Press/Addison-Wesley Publishing Co.

[7] T. Hawkins, A. Wenger, C. Tchou, A. Gardner, F. Göransson, and P. Debevec. Animatable facial reflectance fields. In *Proceedings of the Fifteenth Eurographics Conference on Rendering Techniques*, EGSR'04, pages 309–319, Aire-la-Ville, Switzerland, Switzerland, 2004. Eurographics Association.

[8] D. E. King. Dlib-ml: A machine learning toolkit. *Journal of Machine Learning Research*, 10:1755–1758, 2009.

[9] V. Masselus, P. Dutré, and F. Anrys. The free-form light stage. In *ACM SIGGRAPH 2002 Conference Abstracts and Applications*, SIGGRAPH '02, pages 262–262, New York, NY, USA, 2002. ACM.

[10] V. Masselus, P. Peers, P. Dutré, and Y. D. Willems. Relighting with 4d incident light fields. *ACM Trans. Graph.*, 22(3):613–620, July 2003.

[11] V. Masselus, P. Peers, P. Dutré, and Y. D. Willems. Smooth reconstruction and compact representation of reflectance functions for image-based relighting. In *Rendering Techniques 2004*, pages 287–298, Norrkoping, Sweden, June 2004. Eurographics Association.

[12] R. Ng, R. Ramamoorthi, and P. Hanrahan. All-frequency shadows using non-linear wavelet lighting approximation. *ACM Trans. Graph.*, 22(3):376–381, July 2003.

[13] J. S. Nimeroff, E. Simoncelli, and J. Dorsey. Efficient re-rendering of naturally illuminated environments. In *Photorealistic Rendering Techniques*, pages 373–388. Springer, 1995.

[14] P. Peers, N. Tamura, W. Matusik, and P. Debevec. Post-production facial performance relighting using reflectance transfer. *ACM Trans. Graph.*, 26(3), July 2007.

[15] P. Ren, Y. Dong, S. Lin, X. Tong, and B. Guo. Image based relighting using neural networks. *ACM Trans. Graph.*, 34(4):111:1–111:12, July 2015.

[16] B. Tunwattanapong, A. Ghosh, and P. Debevec. Practical image-based relighting and editing with spherical-harmonics and local lights. In *Proceedings of the 2011 Conference for Visual Media Production*, CVMP '11, pages 138–147, Washington, DC, USA, 2011. IEEE Computer Society.

[17] A. Wenger, A. Gardner, C. Tchou, J. Unger, T. Hawkins, and P. Debevec. Performance relighting and reflectance transformation with time-multiplexed illumination. *ACM Trans. Graph.*, 24(3):756–764, July 2005.

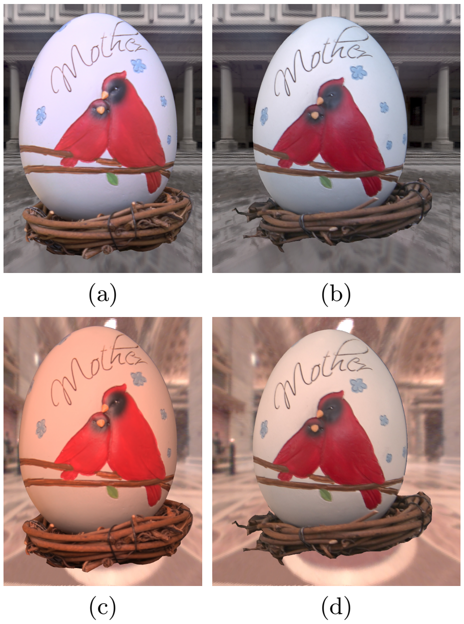

Figure 9: Comparison of the relighting in the Uffizi Gallery (top) and St. Peter's basilica (bottom) environments. **Left column :** Room relighting. **Right column :** Dense free-form acquisition.

Figure 10: Lighting basis of the second room. The top-left light probe is the dark room. The next ten correspond to the five windows full opened and half opened. The bottom right light probe corresponds to the house lights condition.

Figure 11: Lighting basis of the second room projected on the Grace Cathedral environment. The grey area is associated to the dark room condition.

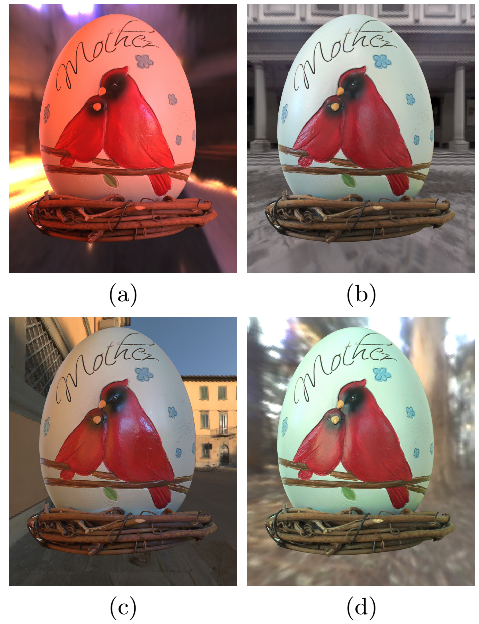

Figure 12: Relighting of the egg using the lighting basis of the second room. **(a)** Grace Cathedral. **(b)** Uffizi Gallery. **(c)** Pisa Couryard. **(d)** Eucalyptus Grove.

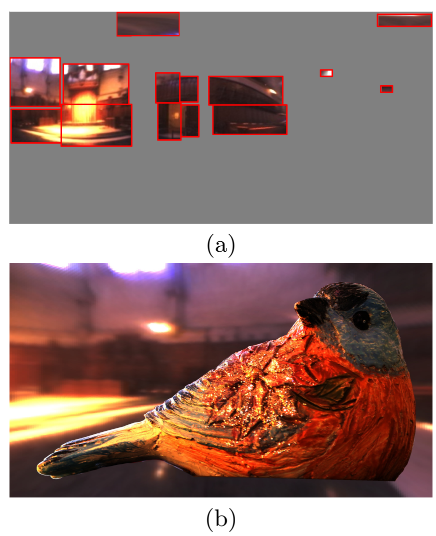

Figure 13: Failure case due to a non-symmetric light distribution. The light sources in the room were mostly orientated at the left of the light probe making this area of the environment dominant in the rendering. **(a)** Projection of the non-symmetric lighting basis. **(b)** Rendering of the bird.