

Measurement-based Editing of Diffuse Albedo with Consistent Interreflections

Bo Dong¹ Yue Dong² Xin Tong² Pieter Peers¹

¹College of William & Mary ²Microsoft Research



Figure 1: Changing the diffuse albedo with consistent interreflections for a decorative table centerpiece and a bathroom scene. The difference images (inset) between the initial scene (left column) and the recolored results (middle and right) highlight the complex changes in interreflections due to a change in diffuse albedo of the planter and middle towel respectively.

Abstract

We present a novel measurement-based method for editing the albedo of diffuse surfaces with consistent interreflections in a photograph of a scene under natural lighting. Key to our method is a novel technique for decomposing a photograph of a scene in several images that encode how much of the observed radiance has interacted a specified number of times with the target diffuse surface. Altering the albedo of the target area is then simply a weighted sum of the decomposed components. We estimate the interaction components by recursively applying the light transport operator and formulate the resulting radiance in each recursion as a linear expression in terms of the relevant interaction components. Our method only requires a camera-projector pair, and the number of required measurements per scene is linearly proportional to the decomposition degree for a single target area. Our method does not impose restrictions on the lighting or on the material properties in the unaltered part of the scene. Furthermore, we extend our method to accommodate editing of the albedo in multiple target areas with consistent interreflections and we introduce a prediction model for reducing the acquisition cost. We demonstrate our method on a variety of scenes and validate the accuracy on both synthetic and real examples.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Color, shading, shadowing, and texture

Keywords: Diffuse Albedo Editing, Interreflection Decomposition, Computational Illumination

1 Introduction

Typical photo-editing operations directly manipulate pixel values without any notion of the underlying physics of light transport through the scene, and significant artistic skill is required to produce plausible results. One such operation is changing the color or albedo of a surface in a photograph. While changing the directly reflected appearance of the target surface is easy, correctly changing the effects of indirect lighting on other surfaces, as well as interreflections back to the target surface, is non-trivial.

One possible strategy for physically-accurate editing of surface albedo and corresponding interreflections is to employ an inverse rendering approach to infer material properties, geometry, and lighting of the scene, then alter the target surface’s albedo, and finally re-render the scene. Not only does the quality of the edited result depend on the accuracy of each of the estimated separate components, it is also very labor intensive compared to the relative modest change in appearance in the photograph.

An alternative strategy is to start from just a single photograph, and by exploiting different heuristics, approximate the effects of physical light transport. Typically, these methods impose restrictions on the materials, light, and/or geometrical components that comprise the scene. While such methods can produce visually pleasing results, the accuracy depends on the heuristics used and how well the assumptions are met.

In this paper, we present a novel method for recoloring surfaces with accurate interreflections that strikes a balance between acquisition cost and flexibility. Our data-driven approach refrains from fully characterizing the scene properties, and only captures the components necessary to model the change in light transport when editing the surface albedo of a target area. We introduce a novel measurement-based method for computing the portion of outgoing radiance that interacts a specified number of times with the target area with an acquisition complexity *linearly* proportional to the maximum number of interactions modeled. The final recolored result is then a weighted linear combination of the estimated interaction components. We acquire the interaction components using a camera-projector pair and conventional methods only, without relying on frequency-based separation, time-of-flight imaging, or transient rendering. Key to our method is the recursive application of the transport operator by appropriately reprojecting the observed radiance, in conjunction with rewriting the transport operator in terms of the relevant interaction components. This allows us to formulate the estimation of the interaction components as a per-pixel linear system. Furthermore, we extend our method to recoloring of multiple target areas with consistent interreflections between the different target areas. However, this comes at the cost of a polynomial acquisition complexity. To reduce the acquisition cost, we propose to build, on-the-fly, an acquisition-prediction model, and only capture photographs when the prediction is insufficiently accurate. While our method is limited to changing the albedo of diffuse surfaces only, it imposes no further restrictions on lighting, geometry, or material properties outside the target area. We demonstrate our method on a variety of scenes, and validate the accuracy of editing and decomposition on synthetic and real scenes.

2 Related Work

Inverse Rendering To support artistic direction while alleviating the high computational costs of forward rendering pipelines, a number of methods (e.g., [Sun et al. 2007; Hašan and Ramamoorthi 2013; Ben-Artzi et al. 2008]) rely on reuse and precomputation to efficiently visualize the effects of altering the albedo or the reflectance properties in a synthetic scene. Such methods can, in conjunction with inverse methods such as [Yu et al. 1999], be used to alter albedo and reflectance properties in photographs of *physical* scenes. However, the quality of the edited result strongly depends on the accuracy of the estimated components, including the unaltered ones. Furthermore, such inverse methods can be labor-intensive and computationally expensive compared to the modest change in the resulting photograph. Instead, the proposed method aims to reduce acquisition and computation cost by only capturing and modifying the affected light transport components.

Recoloring Recoloring is a natural application of intrinsic images [Barrow and Tenenbaum 1978] that represent a decomposition of an image into its shading or illumination component and its reflectance component. Recoloring is achieved by altering the reflectance component to the desired color before recomposing the image. Beigpour and van de Weijer [2011] constrain intrinsic image decomposition to a single-colored object and propose a Multi-Illuminant Dichromatic Reflection model that permits object recoloring under multiple light sources and which accounts for specular reflections. Carroll et al. [2011] recolor images with consistent *first-order* diffuse interreflections by decomposing the illumination component of an intrinsic image decomposition in direct lighting and indirect diffuse lighting for each material. However, a key assumption is that interreflections are a linear mix of a sparse set of independent colors, and hence it cannot disambiguate between interreflections of surfaces with similar colors. Similar to Carroll et al., we also focus on recoloring diffuse surfaces with consistent interreflections. However, whereas Carroll et al. aim to produce plausible results from a single photograph, we target physically-accurate recoloring and take higher-order interreflections in account, which in contrast to first-order interreflections, can alter the appearance of the target area itself in a non-linear way (e.g., light reflected multiple times from the target area). Furthermore, we impose no restrictions on the incident lighting or on the colors in the scene. However, this increased flexibility and accuracy comes at the cost of more input images.

Light Transport Decomposition The proposed technique is also closely related to measurement-based light transport decomposition methods and inverse light transport theory. Nayar et al. [2006] and O’Toole et al. [2012] decompose light transport into direct and indirect lighting. While these methods only require a few measurements, they can only separate the direct and indirect components for uniform lighting. Reddy et al. [2012] estimate direct, near indirect, and far indirect lighting under any initial projector lighting condition. While these methods could potentially be adapted to recoloring, they would only be applicable to full scene recoloring and they would only correctly account for *first-order* interreflections.

Wu et al. [2012] and O’Toole et al. [2014] rely on time-of-flight cameras to tease apart different lighting components. However, similarly to the direct-indirect separation methods, the decompositions are not suitable for recoloring.

Inverse light transport methods [Bai et al. 2010; Seitz et al. 2005; Ng et al. 2012] employ an interreflection or “bounce” cancellation operator to remove the effects of a single interreflection. Repeated application of the cancellation operator decomposes the full transport matrix in its separate “bounces”. Such decomposition is ideally suited for recoloring with consistent higher-order interreflections. However, all of these methods require full knowledge of the transport matrix which can be costly to acquire at high resolution. In contrast, the proposed method computes a related decomposition with significantly less measurements, at the cost of “baking” in the initial lighting condition.

O’Toole and Kutulakos [2010] propose optical computing to invert the light transport operator on a given lighting condition by repeatedly re-emitting the observations. Similarly, we also re-emit the observations, but instead of inverting the light transport operator we estimate the different interreflection components.

3 Measurement-based Recoloring

Our goal is to decompose a photograph of a scene under arbitrary lighting such that we can effectively modify the albedo of a user-selected target area, while keeping the interreflections consistent




Figure 2: *Left:* diffuse lighting $L_d(x)$ that models the outgoing radiance at x on the target area (marked in cyan). Note that this is the result of all sub-paths from the light source that end at x . This includes all sub-paths of varying length (examples of length one (blue), two (green) and three (red) shown) that do not interact with the target area except at the end point x . Furthermore, it includes the diffuse surface reflection at x (marked in light blue). *Middle:* interaction operator $\mathbf{R}(y, x)$ that relates the outgoing radiance at point y on the target area to the outgoing radiance at point x . *Right:* camera operator $\mathbf{C}(u, y)$ that models all sub-paths from y on the target area to a pixel u .

both inside *and* outside the target area. Our measurement-based solution builds on two key ideas: an “interaction decomposition” suitable for recoloring, and the *recursive* evaluation of light transport through a scene to obtain sufficient observations to estimate the interaction components.

Interaction Decomposition Consider a light path starting at the light source, bouncing several times around the scene while interacting n times with the target area before reaching the camera. Scaling the albedo of the target surface by a factor κ , results in a change of the path’s contribution by κ^n . Hence, a suitable decomposition will partition the light transport in “interaction components” that aggregate the contributions from all light paths that interacted a specified number of times with the target area. Appropriately weighting (by κ^n) and summing the different components yields the recolored result.

Recursive Acquisition To estimate the different interaction components, we repeatedly re-emit the observed radiance (normalized by direct lighting), in effect recursively evaluating light transport through a scene. Our key observation is that the recursive evaluation of light transport can be expressed as a weighted sum of the original light transport’s interaction components. Intuitively, light paths encoded in the observations are “concatenated” with the light paths that constitute the light transport through the scene. The resulting concatenated light paths are still valid light paths through the scene, and thus part of the original interaction components. Given the recursive observations, we then formulate the recovery of the interaction components as a linear system.

The remainder of this section is organized as follows. We first introduce a novel light transport formulation that supports albedo-recoloring of diffuse target surfaces (Section 3.1) and that allows us to specify the relation between the recursive evaluation of light transport and the interaction components (Section 3.2). Finally, we discuss practical considerations for the recoloring of physical scenes in Section 3.3.

3.1 Light Transport Formulation

Assume we have a photograph of a scene under (unknown) incident lighting and a user-selected directly visible *diffuse* target area A for which we would like to change the albedo by a factor κ . Light transport along a light path starting at the light source, bouncing through the scene before reaching the camera, remains unaltered except for interactions at the target area. To better reflect the effects of recoloring on light transport, we “chop” the light paths in sub-

paths at each interaction with the target area, and group these sub-paths based on whether they start or end (or both) at the target area.

Define the diffuse lighting $L_d(x)$ for $x \in A$ as the outgoing radiance at x from all sub-paths that start at the light source and which do not interact with the target area A , except for the last interaction at x (Figure 2-Left). Since the sub-paths only interact once with the target A at x , scaling the albedo of the target area by a factor κ implies that the corresponding “recolored” diffuse lighting equals $\kappa L_d(x)$.

Similarly, define the single-interaction operator $\mathbf{R}(y, x)$ for $x, y \in A$ that relates outgoing radiance at y to outgoing radiance at x over all sub-paths from y to x of any length that do not interact with the target area A , except at the last interaction at x ¹. However, these sub-paths can interact with any number of surface points outside the target area (Figure 2-Middle). Similarly, since all of the sub-paths only interact with the target area A at x , scaling the albedo of the target area by κ implies that the corresponding recolored operator is equal to $\kappa \mathbf{R}(y, x)$.

Finally, define the camera operator $\mathbf{C}(u, y)$ which relates the outgoing radiance at $y \in A$ to the observed radiance at a pixel u , over all sub-paths of any length that do not interact with the target area (Figure 2-Right). The operator $\mathbf{C}(u, y)$ is not affected by changes in the target surface’s albedo, since the sub-paths in $\mathbf{C}(u, y)$ only leave the target area at y without interacting with it.

To combine these operators to describe the light transport through a scene, let us first consider the light transport between two points x and y on the target area. The transport between both points is the result of sub-paths that interact once with the target area, plus the sub-paths that interact twice, plus the sub-paths that interact three times, etc. However, the operator \mathbf{R} only models single interactions². To model higher-order interactions, observe that the operator \mathbf{R} aggregates *all* single-interaction sub-paths between two surface points x and y on the target area, and thus any collection of sub-paths that interact twice with the same set of surface points must be a concatenation of the respective single-interaction sub-paths. Hence, the outgoing reflectance of all sub-paths that start at y , interact at any point $z \in A$, and then end at x is equal to $\int_A \mathbf{R}(y, z) \mathbf{R}(z, x) dz$. Using the discrete matrix equivalent of the operator $\mathbf{R}(y, x)$, we

¹The sub-paths start at y , and hence do not interact with the target area A at y .

²The operator \mathbf{R} does not describe a single “bounce” of light, but models a single interaction with the target area which can be the result of sub-paths of any length. This is a key difference with prior work on inverse light transport theory [Bai et al. 2010; Seitz et al. 2005; Ng et al. 2012].

can write the two-interaction operator as: $\mathbf{R}\mathbf{R} = \mathbf{R}^2$. This trivially generalizes to i interactions as: \mathbf{R}^i .

Consider first the case of directly visible surface points on the target area. To express the full transport, we connect the multiple interaction operators \mathbf{R}^i to the camera via \mathbf{C} and the lighting via L_d . Using the discrete matrix and vector equivalents of the operators \mathbf{R} , \mathbf{C} , and L_d , we can then express the observed radiance I as:

$$I = \mathbf{C} \left(\sum_{i=0}^{\infty} \mathbf{R}^i L_d \right), \quad (1)$$

$$= \sum_{i=0}^{\infty} \mathbf{R}^i L_d. \quad (2)$$

We can drop the operator \mathbf{C} because the points on the target surface are directly visible, and hence the sub-paths between camera and surface points on the target area do not interact at any other surface points inside or outside the target area.

As noted before, changing the albedo by a scale factor κ is equivalent to scaling the light paths that interact i times with the target area by a factor κ^i . Observe that each term in Equation (2) interacts $i + 1$ times with the target area (i times for \mathbf{R}^i , and 1 time for L_d). Thus, the observed radiance when recoloring the target area is:

$$I' = \sum_{i=0}^{\infty} \kappa^{i+1} \mathbf{R}^i L_d. \quad (3)$$

To formulate the observed radiance at a point outside the target area, we need to define an additional function $\bar{L}_d(u)$ that represents the observed radiance at a camera pixel u from sub-paths that start at the light source and that do not interact with the target area. Similarly as for the camera operator, $\bar{L}_d(u)$ is not affected by scaling the albedo in the target area. The observed radiance outside the target area after recoloring is then:

$$\bar{I}' = \bar{L}_d + \mathbf{C} \left(\sum_{i=0}^{\infty} \kappa^{i+1} \mathbf{R}^i L_d \right). \quad (4)$$

While the summations in Equations (2) and (4) are infinite, practically, $\mathbf{R}^i \approx 0$ when $i > N$ for some large enough value of N . In the remainder of this paper, we will assume that this approximation is exact, and limit the summation to $N + 1$ terms.

3.2 Interaction Component Estimation

Observe that in order to compute the recolored images I' (Equation (3)) and \bar{I}' (Equation (4)), we only need to estimate the $N + 1$ terms $\mathbf{R}^i L_d$ and $\mathbf{C}\mathbf{R}^i L_d$, and the additional term \bar{L}_d , instead of fully characterizing \mathbf{C} , \mathbf{R} , L_d separately. While there are typically more terms than operators, each term is only a vector (i.e., image), whereas \mathbf{R} is a matrix. The length of the vector (and thus dimensions of the matrix) is significantly larger than $N + 1$, and hence the $N + 1$ terms require less effort to fully characterize than \mathbf{R} .

We first focus on recoloring the directly visible points inside the target area, and hence only the $N + 1$ terms $\mathbf{R}^i L_d$, $0 \leq i \leq N$ are required. Unfortunately, the terms $\mathbf{R}^i L_d$ cannot be directly measured. Instead, we seek a sequence of observations, generated by altering the direct lighting L_d , that allows us to infer the desired terms $\mathbf{R}^i L_d$. Our key insight is that the recursive sequence I_r , the result of repeatedly applying the transport operator $\mathbf{T} = \sum_i \mathbf{R}^i$ on I_{r-1} :

$$I_r = \mathbf{T}I_{r-1}, \quad (5)$$

Table 1: Relation between $\mathbf{T}^{r+1}L_d$ and $\mathbf{R}^n L_d$. Each $\mathbf{T}^{r+1}L_d$ is equal to the weighted sum of the weights (in the same row) times the corresponding $\mathbf{R}^n L_d$ terms (listed in the columns). The weights listed correspond to the figurate numbers $P_r(n + 1)$.

	$\mathbf{R}^0 L_d$	$\mathbf{R}^1 L_d$	$\mathbf{R}^2 L_d$	$\mathbf{R}^3 L_d$	$\mathbf{R}^4 L_d$
$\mathbf{T}^1 L_d$	1	1	1	1	1
$\mathbf{T}^2 L_d$	1	2	3	4	5
$\mathbf{T}^3 L_d$	1	3	6	10	15
$\mathbf{T}^4 L_d$	1	4	10	20	35
$\mathbf{T}^5 L_d$	1	5	15	35	70

with $I_0 = I = \mathbf{T}L_d$, the original observation (Equation (2)), yields a sequence of polynomials in terms of \mathbf{R}^i only:

$$I_r = \mathbf{T}^{r+1}L_d = \sum_{n=0}^N P_r(n + 1)\mathbf{R}^n L_d, \quad (6)$$

where $P_r(n) = \binom{n+r-1}{r}$ are figurate numbers (see Table 1 and Appendix A).

Given the sequence $\{I_r\}_{r=0}^N$ of $N + 1$ observations (obtained by replacing the outgoing radiance L_d at the target area with the previous observation of the scene), it follows from Equation (6) that the $N + 1$ unknown terms $\mathbf{R}^i L_d$ can be computed by solving a linear system per pixel (Equation (6)). To ensure a physically plausible solution, we constrain the solution to: $0 \leq \mathbf{R}^i L_d \leq \mathbf{R}^{i-1} L_d$.

The corresponding sequence of observed radiance outside the target area yields a similar recursive sequence:

$$\bar{I}_r = \mathbf{C}\mathbf{T}I_{r-1}, \quad (7)$$

$$= \mathbf{C}(\mathbf{T}^{r+1}L_d), \quad (8)$$

$$= \sum_{n=0}^N P_r(n + 1)\mathbf{C}\mathbf{R}^n L_d, \quad (9)$$

where $\bar{I}_0 = \bar{I}$, the original observation. Thus $\{\bar{I}_i\}_{i=0}^N$ represents the observed radiance outside the target area when setting the outgoing radiance at the target area to match the previously observed radiance *inside* the target area. Observe that only the initial observation \bar{I}_0 contains the term \bar{L}_d , but not any of the N subsequent observations $\{\bar{I}_i\}_{i=1}^N$. Given the sequence $\{\bar{I}_i\}_{i=0}^N$ of $N + 1$ observations, we can compute \bar{L}_d and the unknown terms $\mathbf{C}\mathbf{R}^i L_d$ by solving a linear system per pixel, using a similar constraint as before and enforcing $\bar{L}_d \geq 0$. Note, we can only solve for up to N interactions $\mathbf{C}\mathbf{R}^i L_d$, compared to $N + 1$ interactions for points inside the target area because of the additional unknown direct lighting \bar{L}_d . If we know a priori that $\bar{L}_d = 0$, then an additional interaction component can be estimated.

3.3 Acquisition

To practically acquire the sequences $\{I_i\}_i$ and $\{\bar{I}_i\}_i$, we first note that both I_r and \bar{I}_r can be observed in a single photograph $\mathcal{I}_r = I_r \cup \bar{I}_r$ of the scene since in both cases the lighting is identical, and where I_r are the observed pixels on the target area, and where \bar{I}_r are the pixels outside the target area.

To alter the outgoing radiance L_d at the diffuse target area, we employ a camera-projector system. We assume that the camera and the projector are radiometrically calibrated, and that the correspondences between camera and projector pixels are known. We subtract \mathcal{I}_0 from subsequent observations to remove the light transport

input : Target Areas: A_1, A_2 , albedo scales κ_1, κ_2 , #interactions N
output: Recolored image $I'' \cup \tilde{I}''$

- 1 initialization: $\mathcal{I}'_0 = \text{Recolor}(L, \kappa_1, A_1)$;
- 2 **for** $r=1$ **to** N **do**
 - 3 $L_r = (\mathcal{I}'_{r-1} \cap A_2)/D$;
// New lighting to be emitted onto A_2
 - 4 $\mathcal{I}'_r = \text{Recolor}(L_r, \kappa_1, A_1)$;
// Recoloring from albedo change on A_1 under
// initial lighting L_r
- 5 compute: $I'' \cup \tilde{I}''$ from $\{\mathcal{I}'_r\}_{r=0}^N$ (for A_2 and κ_2);

Algorithm 1: Recoloring of two target areas.

effects of the initial (unknown and uncontrolled) lighting condition and only retain the effects from the emitted lighting.

When emitting a projector pattern L , the resulting direct outgoing radiance L_d at the target area can be modeled by: $D(x)L(x)$, where D is a (per pixel) scale factor that depends on the diffuse surface albedo, foreshortening, attenuation, etc. Instead of separately modeling each of these factors, we will estimate D directly. Observe when $L = 1$, that the *direct* outgoing radiance $L_d = D$. This corresponds exactly to the direct image from [Nayar et al. 2006]. Once the direct image D is estimated, acquisition of \mathcal{I}_r is achieved by capturing a photograph of the scene when emitting $(\mathcal{I}_{r-1} \cap A)/D$ from the projector to the target area – we abuse notation to indicate with $\mathcal{I}_{r-1} \cap A$ that we only select the pixels in \mathcal{I}_{r-1} that overlap with the target area A .

Summary The above algorithm allows us to estimate an N -interaction decomposition for a single target area in $O(N)$ acquisition complexity. While the material properties in the target area are limited to diffuse only, no restrictions are imposed on the material properties outside the target area or on the initial incident lighting which does not need to originate from the projector.

4 Multiple Target Areas

We now extend our method to multiple target areas. We first consider the two-target area case, where the albedo of the first target area A_1 is scaled by a factor κ_1 , and the albedo of the second area A_2 by a factor κ_2 . We will also assume that the user first finalizes the adjustment to A_1 (i.e., κ_1 is fixed thereafter), before adjusting the albedo in A_2 .

Denote $\text{Recolor}(L, \kappa_1, A_1)$ the resulting image of recoloring target area A_1 under the original incident lighting L . This recolored image also contains the altered interreflections outside A_1 (and thus in A_2). Consequently, when recoloring target area A_2 , we can consider $\text{Recolor}(L, \kappa_1, A_1)$ as the initial observation in Equations (5) and (9). However, naively running the acquisition process again for the second target area A_2 , starting from this “initial” observation, will not produce the correct result because the albedo in target area A_1 has not been altered physically. Consequently, *observed* physical interreflections from light paths between surface points on A_2 that pass i times through A_1 will not be scaled by the factor κ_1^i (Figure 8 Bottom-Left). Hence, each time we illuminate target area A_2 with some novel lighting condition L' , we cannot directly use the observed photograph, but we need to compute the recolored image $\text{Recolor}(L', \kappa_1, A_1)$ that accounts for the scaled albedo of target area A_1 . Thus, for each new lighting condition on A_2 , we need to perform a decomposition on A_1 under this new lighting condition in order to perform the recoloring. Algorithm 1 summarizes the two-target area recoloring procedure.

Discussion Each call in the for-loop to $\text{Recolor}(L_r, \kappa_1, A_1)$ (line 1 & 4) in Algorithm 1 captures $O(N)$ photographs of the scene, resulting in a total acquisition complexity of $O(N^2)$. In total, $N + 1$ recolored images are computed (i.e., N times for target area A_1 , and once for target area A_2). Note, that once the recolored *observations* for $\{\mathcal{I}'_r\}_{r=0}^N$ are computed, all recolored interreflections from target area A_1 are accounted for, and thus we do not need to take A_1 into account anymore when computing the final recolored image (line 5).

The lighting condition L_r (line 3) is based on the recolored observations \mathcal{I}'_r (i.e., recolored according to κ_1 on A_1). Because recoloring has a non-linear effect on the observed image (and thus new lighting), the recolor scale factor κ_1 for the first target area A_1 needs to be fixed a priori. However, the recolor factor κ_2 for the second target area A_2 can still be modified after acquisition.

The above algorithm trivially generalizes to 3 or more target areas; recoloring is recursively cascaded to the previous target areas. The total acquisition complexity for recoloring K target areas is $O(N^K)$. When altering a large number of target areas, it might be more efficient to first capture the light transport matrix \mathbf{T} , and generate the sequences in Equations (5) and (9) numerically (cf. image-based relighting). The transition point for K at which it is more efficient to first capture the transport matrix depends on N and on the acquisition complexity of the method used to capture \mathbf{T} .

5 Adaptive Capture

Depending on the number of interactions N and the number of target areas K , the acquisition cost can be significant. In order to reduce the acquisition cost, we employ a prediction model that given the projector lighting pattern, produces an approximation of the observed photograph. If the predicted approximation is accurate enough, then we can use this instead of capturing the photograph.

Linear Prediction Denote $L_r = (\mathcal{I}_{r-1} \cap A)/D$, $r < N$: the projector lighting that generates the r -th observation \mathcal{I}_r . Our goal is to predict the observations based on the lighting. For this, we first approximate the lighting L_r as a weighted linear combination of the previously emitted lighting conditions $\{L_i\}_{i=1}^{r-1}$ (excluding the initial *unknown* lighting):

$$\underset{w_i}{\operatorname{argmin}} \left| L_r - \sum_{i=1}^{r-1} w_i L_i \right|_2. \quad (10)$$

Due to linearity of light transport, we can replace each L_r by \mathcal{I}_r , yielding the following linear prediction of the observations:

$$\mathcal{I}_r \approx \tilde{\mathcal{I}}_r = \sum_{i=1}^{r-1} w_i \mathcal{I}_i, \quad (11)$$

where the weights w_i are the same as for the lighting.

We validate the quality of the linear predictor based on the relative error between \mathcal{I}_r and its prediction $\tilde{\mathcal{I}}_r$. If this falls below some user-set threshold, then we halt the acquisition, and use the predictor for subsequent observations: $\{\mathcal{I}_i\}_{i=r+1}^N = \{\tilde{\mathcal{I}}_i\}_{i=r+1}^N$, as subsequent lighting conditions will be linear combinations of the first $r - 1$ lighting conditions.

Recolored Prediction Ideally, we would like to control the error on the *recolored* result instead of the error on the observations. When the scaling factor $\kappa \approx 1$, then the above linear predictor suffices. When $\kappa \ll 1$, then the magnitude of the interreflections are

reduced, and the error on the observations will be too conservative. However, when $\kappa \gg 1$, then the influence of interreflections increases, and the error on the observations will tend to underestimate the error on the recolored image, especially for higher-order interactions. To correctly account for such cases, we adapt the linear prediction scheme to predict the *recolored* observations \mathcal{I}'_r given the incident lighting.

Observe that any \mathcal{I}_r is a weighted sum of the same interreflection components as \mathcal{I}_0 weighted by the appropriate figurate number (Equations (6) and (9)). Hence, given a decomposition of \mathcal{I}_0 we can easily compute the recolored versions of any observation \mathcal{I}_r . However, this poses a dilemma as we need *all* N observations $\{\mathcal{I}_r\}_{r=0}^N$ to decompose \mathcal{I}_0 , while we only have a partial sequence $\{\mathcal{I}_r\}_{r=0}^M$, with $M < N$ available. We circumvent this dilemma by computing an approximate decomposition on the first M observations supplemented by $N - M$ linear predictions as in Equation (11).

To predict \mathcal{I}'_r , we recombine the recolored components $\kappa^{i+1} \mathbf{R}^i L_d$ according to Equations (6) and (9). However, unlike the linear predictor, the lighting conditions are not linearly related to the predicted recolored images. Consequently, a relative error between the recolored prediction $\tilde{\mathcal{I}}'_r$ and the recolored observation \mathcal{I}'_r below the user-set error threshold does not guarantee that all subsequent predictions will fall below the threshold. Instead, we will use the relative error on \mathcal{I}'_r to decide whether to capture photograph \mathcal{I}_{r+1} , and repeat the process again for $r + 1$ with either the captured or predicted image. Hence, the final captured sequence will be a mix of captured and predicted photographs. In our implementation, we set the error threshold to 1%.

Discussion The number of successful predictions depends on the specifics of the scene and the recolor scale κ . A smaller κ reduces the influence of higher-order interreflections, and thus allows us to terminate the acquisition earlier. In contrast, a larger scale factor κ typically requires more acquisitions to obtain an accurate recolored photograph. In general, the accuracy of the prediction for \mathcal{I}_r improves for increasing size of the sequence $\{\mathcal{I}_i\}_{i=0}^{r-1}$, and thus increases the likelihood of reducing the acquisition cost. However, there is no requirement that only observations from the same sequence can be used. In the case of multiple target areas, decompositions of several lighting conditions on initial target areas are sought, and thus we can trivially increase the size of $\{\mathcal{I}_i\}_i$ by including the observations from prior decompositions (on the same target area).

Finally, an interesting application of the *linear* prediction model is to automatically determine the optimal decomposition order in a regular (non-recoloring related) interaction decomposition by not setting a maximum interaction count N , but instead only terminating the acquisition when the relative error on the linear prediction falls below the error threshold.

6 Results

Acquisition & Calibration The results in this paper were captured with a DLP projector and a Nikon D700 camera. We directly use the RAW image format of the camera to ensure a radiometrically linear response. The gamma curve of projector is measured by emitting each intensity value and recording the response. We emit and capture each color channel separately to avoid color crosstalk between the projector and the camera at the cost of tripling the acquisition cost. The camera and projector are not co-axial (but positioned nearby), and correspondences between camera and projector are estimated using gray coded structured light [Inokuchi et al. 1984]. We increase the dynamic range of the projector by exploit-




Figure 3: Three recolored results using a 3rd-order decomposition of a plaster statue, initially illuminated by a full-on projector pattern, using recoloring scales $\kappa = (1.3, 0.6, 0.2)$ (2nd column), $\kappa = (0.2, 1.2, 1.2)$ (3rd column), and $\kappa = (0.1, 0.1, 0.1)$ (last column).




Figure 4: Two-target area recolored result of a scene containing four balls, using a 3rd-order decomposition, under two different initial lighting conditions: fluorescent office lighting (left) and an incandescent light source positioned on the right side (right).

ing linearity of light transport and apply the following two steps:

1. To maximize the available dynamic range and to avoid oversaturation, we scale the projector image by $(\tau / \max \text{ImageIntensity})$, where τ is the desired maximum intensity to emit. In our implementation, we simply set τ to 1. After capture, we undo the scale by dividing the observed image by the same scale factor.
2. Instead of re-emitting I_r directly, we re-emit $(I_r - I_{r-1})$, which effectively drops the direct component $\mathbf{R}^0 L_d$, resulting in an increase of dynamic range. By exploiting Equation (5), and the linearity of light transport, the expected observation \mathcal{I}_{r+1} can be computed as:

$$\mathcal{I} = \mathbf{T}(I_r - I_{r-1}), \quad (12)$$

$$= \mathbf{T}I_r - \mathbf{T}I_{r-1}, \quad (13)$$

$$= \mathcal{I}_{r+1} - \mathcal{I}_r, \quad (14)$$

$$\text{or, } \mathcal{I}_{r+1} = \mathcal{I} + \mathcal{I}_r.$$

Simulated results in this paper are generated using the same pipeline, except that instead of recording a photograph, an image is rendered with path tracing [Pharr and Humphreys 2010]. We perform the synthetic experiments with a simulated co-axial projector-camera configuration to avoid potential bias due to inaccuracies in the correspondences. Furthermore, we modify PBRT to output the ground truth interaction components by keeping track of how many times a ray interacts with the target surface.



Figure 5: Recolored results (using a 3rd-order decomposition) of a wooden toy car and a road roller with recoloring scales $\kappa_{car} = (0.8, 0.1, 0.1)$ and $\kappa_{roller} = (0.9, 0.3, 0.4)$. The proposed measurement-based recoloring method can only modify the albedo of visible surfaces; invisible surfaces (e.g., the underside of the roller) and their corresponding (visible) interreflections cannot be recolored.

Results Figure 3 shows recolored results, using a 3rd-order decomposition, of a plaster statue on a specular floor initially illuminated by a full-on projector pattern. Observe the coherent interreflections from the ear to the head, from the mouth to the nose, and the correct recoloring of the specular reflection on the floor. Furthermore, observe that reducing the recolor scale factor, results in more muted interreflections. Conversely, by increasing the scale factor, interreflections become more pronounced. Recoloring (with any factor κ) is simply a linear weighted combination of the different interaction components that only need to be acquired once. This recoloring example would be difficult to achieve with the method of Carroll et al. [2011] because there is little color difference between the target surface and other surfaces in the scene.

Figure 4 shows a scene containing four balls of which the middle two are recolored with a 3rd-order decomposition using $\kappa_1 = (1.03, 0.6, 0.1)$ and $\kappa_2 = (1.2, 0.1, 0.8)$ for the left and right ball respectively. We demonstrate that our method can produce consistent recolored images under different lighting conditions (i.e., fluorescent office lighting, and an incandescent light source positioned on the right side) that did not originate from the acquisition system’s projector. Note that we can only recolor surface points which are visible from both the camera and the projector. In this particular case the outer rim of the balls are visible from the camera, but not from the projector, and thus cannot be altered. To minimize the number of “non-editable” surface points, projector and camera should be positioned as close together as possible; a co-axial setup would fully eliminate this issue.

Figure 1 highlights two scenes recolored using a 3rd-order decomposition. The initial unaltered scenes (under fluorescent lighting) are shown in the left column, and two recolored examples are shown in the middle and right column. We further highlight the complex interreflection changes in the difference images (insets) between the recolored images and the initial scene. For illustration purposes, we scaled the difference images by a factor 4 (top row) and 2 (bottom row). In the *decorative table centerpiece* scene (top row), the color of the textured planter is scaled by $\kappa_{middle} = (0.5, 1.0, 0.4)$ and $\kappa_{right} = (1.1, 0.3, 0.4)$ respectively. As illustrated by the complex interactions with the glass beads, the proposed measurement-based recoloring method is agnostic to the light transport and geometrical complexity in the scene since each pixel is processed separately. For the *bathroom* scene (bottom row), the diffuse albedo of the middle towel is changed by $\kappa_{middle} = (0.6, 0.6, 1.0)$ and $\kappa_{right} = (1.0, 0.6, 0.6)$ respectively with consistent diffuse higher-order interreflections (between the towels), specular reflections (on the glass dish and the soap dispenser), and caustics (though the

glass dish on the floor).

Figure 5 demonstrates the recoloring using a 3rd-order decomposition of a wooden toy car ($\kappa_{car} = (0.8, 0.1, 0.1)$) and road roller ($\kappa_{roller} = (0.9, 0.3, 0.4)$) that both cast strong interreflections on the floor and on neighboring vehicles. A limitation of our method is that we can only recolor visible surfaces. Invisible surfaces, such as the underside of the road roller, remain unaltered and can potentially cast visible interreflections (e.g., the residual yellow interreflections under the road roller).

Accuracy Figure 6 compares the recoloring of the right wall from $(0.1, 0.8, 0.1)$ to $(0.2, 0.1, 0.8)$ ($\kappa = (2.0, 0.125, 8.0)$) on a simulated Cornell Box-like scene that contains a wide variety of material properties and which is originally lit by an area light source. In addition, the separated 4th-order diffuse interreflections are shown and compared to the ground truth decomposition. The decomposition visualizations are obtained by merging the corresponding pixels from $\mathbf{R}^i L_d$ and $\mathbf{C}\mathbf{R}^{i-1} L_d$ for inside and outside the target area respectively. Consequently, the pixels on the target area are “temporally ahead” of the pixels outside the target area, and thus appear darker. We opt for this visualization because it allows us to show both $L_d = \mathbf{R}^0 L_d$ and \bar{L}_d in the first image. Both the recoloring and the decomposed interreflections closely match the ground truth results with a relative error of 0.6%. The main sources of error are: (1) Monte Carlo noise ($1M$ samples per pixel), especially at higher-order interactions, (2) at the edge of the target area due to discretization of the projected target area that causes a mismatch with the wall edge, and (3) light paths that retain significant energy after the 4th-order interaction.

We further validate the accuracy of our method on a *physical* scene of a V-shape wedge (Figure 7) and where only the target area is illuminated by uniform lighting from the projector. We decompose the interreflections from a diffuse colored paper attached to the right side of the wedge. We also capture a ground truth recolored result by carefully replacing the colored paper. We compute the recolor scale factor based on the ratio of the direct components [Nayar et al. 2006] of each of the pieces of colored paper. The recolored result closely matches the ground truth measurement with a relative error of 5.8% (3.7% on the left wall of the V-shape wedge) for a 3rd-order decomposition. A 2nd-order decomposition yields a $1.2\times$ ($1.6\times$ on the left wall) larger error. For this particular scene and setup, a 4th-order decomposition did not yield a noticeable improvement compared to a 3rd-order decomposition due to the impact of measurement noise.




Figure 6: Recoloring of the right wall ($\kappa = (2.0, 0.125, 8.0)$) from green to blue with consistent interreflections in a synthetic scene. **Top Row:** ground truth recolored result and a 4th-order decomposition. The brightness of the decomposition is adjusted for illustration purposes by a factor of 1, 4, 16, and 64 respectively. **Bottom Row:** 4th-order decomposition and the resulting recoloring obtained with the proposed method. **Middle Row:** False color difference between the ground truth and the estimated decomposition.




Figure 7: Recoloring using a 3rd-order decomposition of a V-shape wedge scene, where only the target area (right wall) is illuminated from the projector. The ground truth recolored result is obtained by placing a diffuse red paper over the target area.

Multiple Targets Figure 8 shows a simulated two-target area recoloring on both sides of a V-shaped wedge using a 4th-order decomposition. We recolor the left side first from purple $(0.6, 0.2, 0.6)$ to yellow $(0.9, 0.9, 0.3)$ with $\kappa_{left} = (1.5, 4.5, 0.5)$. Next, we recolor the right side from light blue $(0.4, 0.7, 0.8)$ to light purple $(0.8, 0.0875, 0.96)$, with $\kappa_{right} = (2.0, 0.125, 1.2)$. The recolored result closely matches the ground truth recolored result with a relative error of 1.2%. Capturing all

images, without using the selective capture method from Section 5, results in a total of 16 images for a 4th-order decomposition. When employing the adaptive acquisition algorithm, we only capture 11 images with approximately the same relative error (1.2%). When using a 5th-order decomposition the adaptive acquisition captures exactly the same number of images (11) versus 25 for the brute force acquisition. This indicates that a 5th or higher order decomposition does not contribute additional information for this particular combination of scene and recoloring. In general, increasing the albedo requires a more accurate decomposition (and thus higher-order interactions) as it increases the magnitude of the interreflections. Furthermore, increasing the albedo is further complicated by the amplification of measurement noise.

Target areas do not need to be connected and can exhibit interreflections between surface points inside the target area. Figure 9 demonstrates this on the V-shape wedge scene from Figure 8 by comparing a 4th-order decomposition and recoloring with two target areas with identical recoloring scales ($\kappa_1 = \kappa_2 = (1.5, 1.2, 0.25)$) and a recoloring with a single target area that is the union of the previous two areas. Both images match the ground truth with a relative error of 8.9% and 6.6% for the single-target area and two-target area recoloring respectively. The slightly higher relative error (mostly located on the ground plane) for the single-target area recoloring is due to the relatively shorter light paths, and thus relative higher importance of the separate decompositions.

As noted before, our method trivially extends to more than two target areas. Figure 10 demonstrates a three-target area recoloring. Starting from the initial scene (left-top), we first alter the red area to yellow using $\kappa_1 = (0.86, 6.0, 1.0)$ (middle-top), then the blue




Figure 8: A V-shape wedge scene with two target areas covering the left (1st) and right (2nd) wall respectively. **Top-left:** The initial scene illuminated from the projector with uniform white. **Bottom-middle:** Ground truth recolored result with $\kappa_{left} = (1.5, 4.5, 0.5)$ and $\kappa_{right} = (2.0, 0.125, 1.2)$. **Bottom-left:** Naive recoloring of each target area separately fails to produce the correct result as highlighted by the large error in the false color difference image. **Top-right:** brute-force recoloring with an $O(N^2)$ acquisition complexity on a 4th-order decomposition (16 images) produces results that closely match the ground truth (1.2% relative error). **Bottom-right:** adaptive 4th-order decomposition with a 1% relative error threshold only captures 11 images with an almost identical error.




Figure 9: A comparison between a single-target area recoloring and a two-target area recoloring, where the single target area is the union of the two target areas, and where the recolor weight is the same for all target areas $\kappa = (1.5, 1.2, 0.25)$. Both recolored results are comparable to the ground truth result with a relative error of 6.6% and 8.9% for the two-target area and single-target area recoloring respectively.

area to green using $\kappa_2 = (2.0, 4.5, 0.5)$ (right-top), and finally the beige area to white using $\kappa_3 = (2.0, 2.0, 4.0)$ (middle-bottom). Note that even though the first target area and the second target area are in the same plane, there is still indirect light transport via the third target area between both. The error on the brute-force (64 captured images) and the adaptive capture (38 captured images) is virtually identical (1.3%) for a 4th-order decomposition.

Figure 11 shows the results of a two-target area recoloring on a physical scene containing a V-shape wedge using a 4th-order decomposition. We first alter the left side from light purple to green ($\kappa_1 = (0.672, 1.18, 0.34)$), and then the right side from yellow to red ($\kappa_2 = (0.49, 0.098, 0.25)$). Observe the consistent recoloring of the reflections on the specular floor, and the difference in color (on both the ground truth as well as the recolored results) on the yellow and green sides when altering the albedo on the other side. Similarly as in Figure 7, we compare the recolored result to a ground truth recolored result where we physically alter the albedo of the target areas. Additionally, we show a false color difference image for the full two-target area recolored result. The relative error, for the brute-force capture, after altering the first target is 6.9%




Figure 10: Three-target area recoloring using a 4th-order decomposition. **Top Row:** Initial scene, first target area recolored ($\kappa_1 = (0.86, 6.0, 1.0)$), and second target area recolored result ($\kappa_2 = (2.0, 4.5, 0.5)$). **Middle-bottom:** ground truth three-target area recolored result ($\kappa_3 = (2.0, 2.0, 4.0)$). **Left-bottom:** 64-image brute-force capture recoloring. **Right-bottom:** 38-image adaptive capture recoloring (1% relative error threshold).

(mostly located at the edges), and 6.5% after altering the second target area. The relative error for the adaptive capture is virtually identical while only capturing 9 images (compared to 16 for the brute-force capture).

Discussion and Limitations The proposed method is agnostic to the albedo variation present in the target area, and it will uniformly scale the texture by the recolor scale factor while maintaining consistent interreflections. This is illustrated in Figure 12 on a recolored result (1.1% relative error) of a scene using a 4th-order decomposition and using a recolor scale factor of $\kappa = (2.0, 0.125, 1.2)$. Altering the texture of a surface, besides scaling, is difficult; each texel would correspond to a target area. An interesting avenue for future research would be to exploit the low frequency nature of indirect lighting and adaptively subdivide the target area based on the texture and the error on the interreflections.

A limitation of our method, shared by prior work on inverse light transport [Bai et al. 2010; Seitz et al. 2005; Ng et al. 2012], is that only interreflections from diffuse target areas can be accurately decomposed. However, very few physical materials are perfectly Lambertian. To better understand the impact of specular “pollution” on the target area, we compute the error on a 4th-order decomposition of a V-shape wedge scene for varying degrees of specular albedo and specular roughness (Figure 13). To compare the different decompositions, we average the absolute error for each interaction normalized by the peak value. We employ a Cook-Torrance BRDF model and vary specular roughness on a logarithmic scale from very rough (1.0) to highly specular (0.009). Similarly, we also vary the diffuse-specular ratio on a logarithmic scale and ensure that the sum of specular and diffuse equals 1.0. Examples of varying parameters are shown on the right in Figure 13. The impact of specular roughness is minor compared to the impact of specular albedo. Overall, our method produces an error of less than 5% when the specular albedo is less than 10% of the total albedo.

Similarly, recoloring a translucent target area would produce plausible but slightly incorrect results. Currently when re-emitting the observed radiance, we assume that there exists a one-to-one relation between incident and exitant radiance at each point on the target area. This allows us to “undo” the effects of the first interaction (foreshortening, diffuse albedo, etc.) when projecting an image on the scene by a simple division by D . In case of a translucent mate-




Figure 11: Two-target area recoloring with a 4th-order decomposition of a V-shape wedge scene where the incident lighting originates from the projector covering the full scene. **Top-row:** Initial scene, ground truth first target area recoloring, and ground truth second target area recoloring. **Bottom-middle:** Result after recoloring the first target area. **Bottom-right:** Result after recoloring both target areas. **Bottom-left:** A false color difference of the resulting two-target area recolored result compared to the ground truth.




Figure 12: A recolored result of a textured diffuse surface using a 4th-order decomposition and $\kappa = (2.0, 0.125, 1.2)$.

rial, this one-to-one relation is broken due to the subsurface scattering in general, producing a slightly (more) blurred outgoing radiance L_d .

The accuracy of our method is also impacted by measurement noise. First, when increasing the albedo of a surface (i.e., $\kappa > 1$), noise in the decompositions gets amplified, especially when the original surface albedo is very low, and the scale factor is large. This is further exacerbated for higher-order terms due to the raised scale factor κ^i . Second, measurement noise affects how accurately we can recover the interaction components. The magnitude of the interreflections decreases with increasing order, and thus it will fall below the noise level at some point. Figure 14 plots the effects of noise on the accuracy for a 5th-order decomposition (using the same scene as in Figure 6) by adding camera noise with varying magnitude to the ground truth measurements. We then compute the absolute error on each of the decomposed interactions, normalized by the peak value of the corresponding ground truth decomposition. We vary the standard deviation of Gaussian noise from 0.05% to an extreme of 10%. As can be seen, the first-order direct component is not very sensitive to measurement noise. This is further confirmed by the relative error for this component, which does not rise above 3%. While the absolute magnitude of the absolute error decreases




Figure 13: Impact of specular albedo and roughness on the accuracy of a 4th-order interaction decomposition. **Right:** Examples of the scene with varying specular albedo and roughness on the left wall. **Left:** False color plot of the absolute error (normalized by peak value) averaged over the four interaction components.

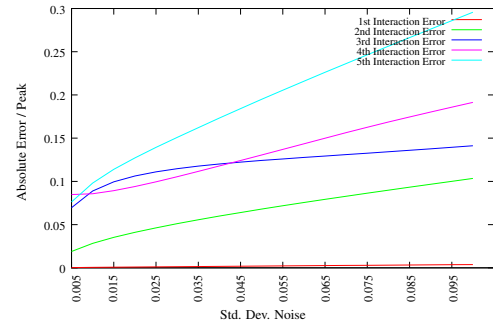


Figure 14: Impact of measurement noise on the accuracy of a 5th-order decomposition of the scene in Figure 6.

with increasing order, the dynamic range (cf. peak) of the interaction components decreases more rapidly. The resulting normalized error gradually increases. Overall, the proposed decomposition is robust to moderate degrees of camera noise. However, this analysis does not include the impact of camera noise on the re-emitted lighting. In general, due to the low-pass behavior of diffuse interreflections, the decomposition of pixels outside the target area are less affected by noise in the emitted lighting. In contrast, incident lighting on the target area is mostly reflected directly back, effectively doubling the camera noise (once from the observation, and once from the lighting). In practice, the impact of noise is further complicated due to the dynamic range maximization applied before emitting the projector pattern, and due to the subtraction of the projector black level.

Finally, prior work on inverse light transport [Bai et al. 2010; Seitz et al. 2005; Ng et al. 2012] start from the rendering equation, and derive a cancellation operator that cancels out a single “bounce” of light transport. However, as noted by Seitz et al. [2005], this is only a true single-bounce cancellation operation between *visible diffuse* surface points; interactions with invisible points are ignored. The derivation of our method, on the other hand, starts by explicitly modeling these “invisible” paths, enabling us to precisely characterize the light transport effects on surfaces outside the target area. However, in the case where the target area covers the full (diffuse) scene, both decompositions align. Inverse light transport methods require full knowledge of the light transport operator which can be costly to acquire. In contrast, the proposed method decomposes light transport more efficiently in $O(N)$ acquisition complexity for a single target area. However, this comes at the cost of a priori fixing the lighting for which to compute the decomposition.

7 Conclusions

In this paper we presented a measurement-based method for changing the albedo of diffuse surfaces in photographs with consistent interreflections. Key to our method is a novel diffuse interreflection decomposition algorithm that can estimate an N -order decomposition for a single target area in $O(N)$ measurements under any initial lighting condition. We believe that our decomposition algorithm is applicable to other problems besides recoloring such as light transport analysis and the decomposition of the full transport matrix. In addition, we demonstrated how our theory can be applied to editing multiple target areas at the cost of a polynomial acquisition complexity. Finally, we introduced a prediction model for reducing the number of required photographs.

Although our method still produces plausible results for a wide range of non-diffuse target surfaces, our method is only exact for recoloring diffuse surfaces. Extending our method to non-diffuse surfaces is an interesting avenue for future research.

Acknowledgements

We wish to thank the reviewers for their constructive feedback, and Victoria Cooper and Kathleen Moore for proofreading. This work was partially funded by NSF grants: IIS-1217765 and IIS-1350323, and a gift from Google.

A Derivation of Equation (6)

First observe that \mathbf{T}^r is a polynomial only in terms of \mathbf{R}^i due to \mathbf{T} being a polynomial only in terms of \mathbf{R}^i . Denote the i -th coefficient of \mathbf{T}^r as $c_r(i)$:

$$\mathbf{T}^r = \sum_i c_r(i) \mathbf{R}^i. \quad (15)$$

We can then derive an expression for \mathbf{T}^{r+1} :

$$\mathbf{T}^{r+1} = \mathbf{T} \mathbf{T}^r, \quad (16)$$

$$= \sum_j \mathbf{R}^j \sum_i c_r(i) \mathbf{R}^i, \quad (17)$$

$$= \sum_j \sum_i c_r(i) \mathbf{R}^{i+j}, \quad (18)$$

$$= \sum_n \sum_{i=0}^n c_r(i) \mathbf{R}^n. \quad (19)$$

Since \mathbf{T}^{r+1} is also a polynomial with coefficients c_{r+1} , the following relation follows from Equation (19):

$$c_{r+1}(n) = \sum_{i=0}^n c_r(i), \quad (20)$$

$$= c_{r+1}(n-1) + c_r(n), \quad (21)$$

where $c_{r+1}(0) = c_r(0)$. Furthermore, from Equation (3), it follows that $c_1(n) = 1$ (and thus $c_r(0) = 1$). Under these conditions, Equation (21) expresses the numbers on the diagonal of Pascal's triangle. With appropriate renumbering this yields:

$$c_{r+1}(n) = P_r(n+1) = \binom{n+r}{r}. \quad (22)$$

References

BAI, J., CHANDRAKER, M., NG, T.-T., AND RAMAMOORTHI, R. 2010. A dual theory of inverse and forward light transport. In *ECCV*, 294–307.

BARROW, H. G., AND TENENBAUM, J. M. 1978. Recovering Intrinsic Scene Characteristics from Images. In *Computer Vision Systems*, vol. 27, 3–26.

BEIGPOUR, S., AND VAN DE WEIJER, J. 2011. Object recoloring based on intrinsic image estimation. In *IEEE ICCV*, 327–334.

BEN-ARTZI, A., EGAN, K., DURAND, F., AND RAMAMOORTHI, R. 2008. A precomputed polynomial representation for interactive brdf editing with global illumination. *ACM Trans. Graph.* 27, 2 (May).

CARROLL, R., RAMAMOORTHI, R., AND AGRAWALA, M. 2011. Illumination decomposition for material recoloring with consistent interreflections. *ACM Trans. Graph.* 30, 4 (July).

HAŠAN, M., AND RAMAMOORTHI, R. 2013. Interactive albedo editing in path-traced volumetric materials. *ACM Trans. Graph.* 32, 2 (Apr.).

INOKUCHI, S., SATO, K., AND MATSUDA, F. 1984. Range imaging system for 3-D object recognition. In *International Conference on Pattern Recognition*, 806–808.

NAYAR, S. K., KRISHNAN, G., GROSSBERG, M. D., AND RASKAR, R. 2006. Fast separation of direct and global components of a scene using high frequency illumination. *ACM Trans. Graph.* 25, 3 (July), 935–944.

NG, T.-T., PAHWA, R. S., BAI, J., TAN, K.-H., AND RAMAMOORTHI, R. 2012. From the rendering equation to stratified light transport inversion. *Int. J. Comput. Vision* 96, 2 (Jan.), 235–251.

O'TOOLE, M., AND KUTULAKOS, K. N. 2010. Optical computing for fast light transport analysis. *ACM Trans. Graph.* 29, 6 (Dec.), 164:1–164:12.

O'TOOLE, M., RASKAR, R., AND KUTULAKOS, K. N. 2012. Primal-dual coding to probe light transport. *ACM Trans. Graph.* 31, 4 (July), 39:1–39:11.

O'TOOLE, M., HEIDE, F., XIAO, L., HULLIN, M. B., HEIDRICH, W., AND KUTULAKOS, K. N. 2014. Temporal frequency probing for 5d transient analysis of global light transport. *ACM Trans. Graph.* 33, 4 (July), 87:1–87:11.

PHARR, M., AND HUMPHREYS, G. 2010. *Physically Based Rendering, Second Edition: From Theory To Implementation*, 2nd ed. Morgan Kaufmann Publishers Inc.

REDDY, D., RAMAMOORTHI, R., AND CURLESS, B. 2012. Frequency-space decomposition and acquisition of light transport under spatially varying illumination. In *ECCV*, 596–610.

SEITZ, S. M., MATSUSHITA, Y., AND KUTULAKOS, K. N. 2005. A theory of inverse light transport. In *IEEE ICCV*, 1440–1447.

SUN, X., ZHOU, K., CHEN, Y., LIN, S., SHI, J., AND GUO, B. 2007. Interactive relighting with dynamic brdfs. *ACM Trans. Graph.* 26, 3.

WU, D., O'TOOLE, M., VELTEN, A., AGRAWAL, A., AND RASKAR, R. 2012. Decomposing global light transport using time of flight imaging. In *IEEE CVPR*, 366–373.

YU, Y., DEBEVEC, P., MALIK, J., AND HAWKINS, T. 1999. Inverse global illumination: Recovering reflectance models of real scenes from photographs. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '99*, 215–224.