

Variable Ring Light Imaging: Capturing Transient Subsurface Scattering with An Ordinary Camera

Ko Nishino^{1,3}, Art Subpa-asa², Yuta Asano²,
Mihoko Shimano³, and Imari Sato³

¹ Kyoto University

² Tokyo Institute of Technology

³ National Institute of Informatics

`kon@i.kyoto-u.ac.jp`, `{art.s.aa,asano.y.ac}@m.titech.ac.jp`,
`{miho,imarik}@nii.ac.jp`

Abstract. Subsurface scattering plays a significant role in determining the appearance of real-world surfaces. A light ray penetrating into the subsurface is repeatedly scattered and absorbed by particles along its path before reemerging from the outer interface, which determines its spectral radiance. We introduce a novel imaging method that enables the decomposition of the appearance of a fronto-parallel real-world surface into images of light with bounded path lengths, i.e., transient subsurface light transport. Our key idea is to observe each surface point under a variable ring light: a circular illumination pattern of increasingly larger radius centered on it. We show that the path length of light captured in each of these observations is naturally lower-bounded by the ring light radius. By taking the difference of ring light images of incrementally larger radii, we compute transient images that encode light with bounded path lengths. Experimental results on synthetic and complex real-world surfaces demonstrate that the recovered transient images reveal the subsurface structure of general translucent inhomogeneous surfaces. We further show that their differences reveal the surface colors at different surface depths. The proposed method is the first to enable the unveiling of dense and continuous subsurface structures from steady-state external appearance using ordinary camera and illumination.

1 Introduction

Subsurface scattering, the light transport inside a surface near its external interface, has a significant effect on the appearance of real-world surfaces. Once incident light penetrates into the subsurface, it bounces off particles (e.g., pigments) in the medium multiple times. In addition to the absorption by the surface medium itself, each of these bounces causes absorption and scattering, leading to the unique color and direction of light reemerging from the interface.




Fig. 1: We propose an imaging method that enables the decomposition of the appearance of subsurface scattering into a set of images of light with bounded path lengths, and proportionally a limited number of bounces with particles inside the surface, i.e., transient light transport in the surface. These transient images collectively reveal the internal structure and color of the surface.

The longer the light travels the more bounces it experiences. If we can decompose the appearance of a subsurface scattering surface into a set of images, each of which captures light that has traveled a limited range of distance (i.e., path length) or probabilistically proportionally a limited number of bounces with particles in the surface, the results would tell us a great deal about the otherwise unobservable internal surface structure underneath the external interface.

Imaging path length or bounce decomposed light has been studied primarily for interreflection as part of scene-wide global light transport analysis. Subsurface scattering is distinct from interreflection as it concerns light transport inside a surface and is governed by the statistically distributed particles rather than much larger-scale 3D scene geometry. As we review in Sec. 2, past methods fundamentally rely on frequency and geometric characteristics that are unique to interreflection, and cannot be applied to decompose subsurface scattering.

This paper introduces the first method for subsurface scattering decomposition, the recovery of path length (and proportionally bounce) dependent light that make up the appearance of real-world translucent surfaces. Our key contribution lies in the use of a natural lower bound on the shortest path length defined by the external interface of a surface. When lit and viewed orthographically from the top, the surface defines a halfspace and the shortest path length of light incident at a single surface point, referred to as impulse illumination, observed at distance r from the point of incidence would naturally be r .

We exploit this lower bound on light path length by observing the same surface point while varying the distance of an impulse illumination. Instead of using impulse illuminations at an arbitrary point, we use a radius- r ring light around each surface point, i.e., all surface points at surface distance r illuminated with impulse illuminations. We obtain variable ring light images, each of which encodes the appearance of each surface point captured with ring lights of varying radii. We show that taking the difference of two ring light images

of slightly different radii bounds the light path length observed at each surface point. In other words, from steady-state appearance captured as variable ring light images, we compute surface appearance with varying degrees of subsurface scattering, i.e., transient images of subsurface light transport. We also show empirical evidence based on simulation and analytical approximation that the path length-bounded transient images may also be interpreted as approximately and proportionally bounded n -bounce images of the surfaces.

We implement variable ring light imaging by capturing N impulse illumination images, where N is the number of surface points (i.e., projector pixels in an observed image region). As shown in Fig. 1, we then synthetically compute variable ring lights images and transient images from their differences. We show extensive experimental results of applying variable ring light imaging to both synthetic and real-world complex surfaces. The results show that the method successfully decomposes subsurface scattering and demonstrate that the recovered transient images reveal the internal surface structure including its color variation and lateral and along-depth configurations for general inhomogeneous surfaces. We further show that the differences of the transient images reveal the colors of the surface at different depths.

Our method is simple requiring only a regular projector and camera. Most important, it is model-free, making no restrictive assumptions about the subsurface light transport. To our knowledge, the proposed method is the first to enable the dense recovery of the continuous variation of inner surface structures from external observation of the steady-state appearance using ordinary imaging components. Variable ring light imaging complements past transient imaging methods especially those for bounce decomposition of scene-wide global light transport, by providing an unprecedented means for deciphering subsurface scattering and consequently for visually probing intrinsic surface structures. We believe the method opens new avenues for richer radiometric understanding of real-world scenes and materials.

2 Related Work

Recently, a number of approaches have been proposed to capture the transient light transport of real-world scenes. Many of these transient imaging methods directly sample the temporal light transport with time-of-flight or ultra-fast cameras in the pico- to femtosecond range [3, 5, 11, 21, 20], which enable separation of large-scale light transport events including reflection and interreflection. It is, however, unclear how such temporal sampling can be applied to disentangle subsurface scattering, where the light interaction happens at a much smaller scale and is not openly accessible from an external view.

Other approaches leverage ordinary or moderately high-speed cameras and programmable illumination (e.g., projectors) to control the light transport itself for capturing its distinct components. Seitz et al. [17] recover n -bounce images of a scene with successive multiplications of “interreflection cancellation operators” to the observed image. Here, a bounce is defined as light interaction between

scene geometry, not with minute particles in a surface that we are interested in. The method specifically requires diffuse interreflection so that the light transport matrix is full-rank and can be inverted to construct the cancellation operator. For subsurface scattering, the light transport viewed and illuminated from the outside will be rank-deficient or worse yet rank-one for completely homogeneous surfaces. For this, the inverse light transport theory [17, 1] does not apply to subsurface scattering.

Nayar et al. [9] introduced the use of spatial high-frequency illumination patterns to decouple 1-bounce and 2-and-higher bounces of light in the scene which are referred to as direct and global light transport, respectively. The method relies on the fact that the first light bounce is essentially an impulse function in the angular space resulting in all-frequency light transport, and the rest of the scene-wide light bounces are of comparatively lower frequency. This is generally true for direct reflection and other light transport components including diffuse interreflection and subsurface scattering. The method, however, can only separate the first bounce (direct) from the rest, and cannot further decompose subsurface scattering or interreflection into their components. Although various subsequent techniques extend the original high-frequency illumination method [18, 6, 2], for instance, to decompose the global light transport into spatially near- and far-range components [15], they are all limited by this fundamental frequency requirement of separable components and cannot be applied to decompose subsurface scattering.

O’Toole et al. [14, 13] introduced the idea of modulating both the illumination and image capture to realize selective probing of the light transport matrix in a coaxial imaging setup, which was later extended to exploit epipolar geometry in non-coaxial systems for realtime direct-indirect imaging [12] and further combined with rolling shutter for efficiency [10]. Our goal is to decompose the indirect light into its constituents based on path length (and proportional number of bounces) of subsurface scattering, and primal-dual imaging cannot be used for this as the intensity of each bounce would be too low to probe with pixel masking. Most important, subsurface scattering does not obey the geometric two-view constraint cleverly exploited in these methods. Similarly, high-frequency temporal light transport sampling can only separate subsurface scattering as a whole and cannot decompose it further [21]. When viewed in light of the categorization of [12], our novel imaging method can be viewed as successive co-axial “long-range indirect imaging” applied to subsurface scattering. The key ideas lie in varying this “range” and to bound path lengths by taking their differences.

A few recent works concern light transport analysis of subsurface scattering. Mukaigawa et al. [7] achieves n -bounce imaging of subsurface scattering by directly observing a side slice of the surface illuminated with high-frequency illumination on its top interface. The parameters of an analytical scattering model is recursively estimated from the observation to compute n -bounce scattering. The method requires physical slicing of the surface to gain a view of the actual light propagation in depth and is restricted to a specific analytical scattering model. As a result, the method cannot be applied to general surfaces including

even homogeneous ones of unknown materials. Our method does not assume either—the imaging setup does not require any modification to the surface of interest and it is model-free, enabling the application to arbitrary surfaces.

Tanaka et al. [19] propose a method for decomposing the appearance of a surface seen from above into a few layers at different depths. The surface is captured with high-frequency illumination of varying pitches whose direct component images [9] are then deblurred with known kernels corresponding to each depth. Subsurface light transport captured with an external illuminant, however, cannot be expressed with the radiative transfer model they use [8] which does not account for scattering of incident light. Furthermore, depth-dependent appearance of subsurface scattering neither results in distinctive texture patterns nor is its component frequencies distinct enough to disentangle with high-frequency illumination. More recently, Redo-Sanchez et al. [16] demonstrated similar discrete layer extraction with sub-picosecond imaging using terahertz time-domain spectroscopy. In contrast, our method is model-free and does not assume any frequency characteristics of the subsurface light transport, and is specifically aimed at decomposing the continuous appearance variation both across the surface extent and along its depth including its colors with an ordinary camera-projector imaging setup.

Diffuse optical tomography temporarily samples transmitted light through a scattering volume, whereas our method bounds the light path length by spatial differentiation on a flat surface, which makes them complementary both in application and methodology.

Our method, for the first time to the best of our knowledge, provides access to the transient light transport of subsurface scattering. The resulting bounded path length images, or proportionally approximate n -bounce images, would likely find use in a wide range of applications. As we demonstrate, they reveal the complex internal subsurface structures including their colors and provide novel means for real-world surface appearance understanding.

3 Variable Ring Light Imaging

In this section, we describe the intuition and theoretical underpinnings of our novel imaging method.

3.1 Preliminaries

A surface at a microscopic scale, i.e., smaller than an image pixel but larger than the wavelength, consists of a medium and particles. An incident light ray, after penetrating through the outer interface, travels straight in the medium and every time it hits a suspended particle it is scattered until eventually reemerging from the interface. Each of these bounces causes the light ray to change its direction while some portion of its energy is absorbed. The degree of absorption by both the medium and the particles vary across the wavelength, causing the light ray to accumulate a unique color depending on its traveled path. The appearance of

a surface point consists of such light rays that have undergone various number of bounces in the subsurface.

Due to this subsurface scattering, light travels on a zigzag path inside the surface. As a result, the *surface distance* of a light ray, i.e., the Euclidean distance on the surface between its entry and exit points, is not the same as its *light path length*, the actual geometric distance the light ray traveled inside the surface. Note that the *optical* path length is the product of the light path length and the index of refraction of the medium. The incident light may have traveled straight right beneath the interface or it may have traveled deep through the surface before reemerging from the interface.

We consider an imaging setup consisting of an orthographic camera and directional light collinear with each other and both perpendicular to the surface of interest. The following theoretical discussion is valid for non-parallel line of sight and illumination, and in practice the camera and illuminant can be tilted with respect to the surface as they will only scale the path length range corresponding to each transient image. Nevertheless, for simplicity both in theory and practice, we will assume a collinear and perpendicular imaging setup.

The surface of interest is assumed to be flat to safely ignore interreflection at its external interface. The surface distance for a non-flat surface would be the geodesic distance which reduces to the geometric distance for a flat surface. It is important to note that, in this paper, our focus is on real-world surfaces that exhibit subsurface scattering. The path length lower bound will not necessarily apply to general scattering without a surface that defines a halfspace, such as smoke and general scene interreflection.

3.2 Impulse Subsurface Scattering

Let us consider a directional light source illuminating a surface point corresponding to a single pixel in the image. We refer to this as impulse illumination [17]. In practice, impulse illumination is achieved by calibrating the correspondences between projector pixels and image pixels and the finest resolution we use is usually the projector pixel (which equates roughly 4 image pixels in our later real experiments). Fig. 2a depicts the isophotes of a subsurface scattering radiance distribution inside the surface.

The radiance of a surface point at distance r , which we denote with $E(r)$, from an impulse illumination point is the sum of the radiance of each light ray, denoted with $L(\cdot)$, that traveled through the surface before reemerging from the interface at surface distance r . Instead of considering each individual light ray ℓ , we will instead distinguish each by its path length $|\ell|$:

$$E(r) = \sum_{|\ell_i| \in \mathfrak{L}(r)} L(|\ell_i|), \quad (1)$$

where $\mathfrak{L}(r)$ is the set of path length of light rays that we observe at surface distance r ,

$$\mathfrak{L}(r) = \{|\ell_i| : |H(\ell_i)| = r\}. \quad (2)$$




Fig. 2: (a) Isophotes of subsurface scattering. (b) Surface appearance of subsurface scattering consists of light rays that have traveled various path lengths which are naturally bounded by the surface distance between the points of incident and observation. (c) Illuminating the surface with variable ring light, a circular impulse illumination centered on the surface point of interest of different radii, and taking their differences lets us bound the path length of light observed at each surface point.

Here, $II(\ell)$ is an operator that returns the entry and exit surface points of a light ray ℓ , $|II(\ell)|$ is the surface distance the light ray traveled, and although it is not denoted explicitly, the set only includes light rays whose paths are contained inside the halfspace defined by the surface. In other words, as Fig. 2b depicts, the light ray set $\mathfrak{L}(r)$ consists of those light rays that entered a surface and exited from a surface point of distance r from entry. We assume that the impulse illuminant is of unit radiance and the camera is calibrated for linear radiance capture.

We are now ready to state a key observation of subsurface scattering:

$$\min \mathfrak{L}(r) \geq r. \quad (3)$$

The proof of this observation is trivial; the shortest path between two points on a surface is the geodesic between them, which we have defined as the surface distance r for our flat surface, and the observed light cannot have traveled a shorter distance to reemerge from the surface.

We rewrite Eq. 2 to make this property explicit:

$$\mathfrak{L}(r) = \{|\ell_i| : |\ell_i| \geq r, |II(\ell_i)| = r\}. \quad (4)$$

3.3 Ring Light Subsurface Scattering

Given the lower bound on the light path length, observations of a surface point with impulse illuminations would likely let us extract light of bounded path lengths (i.e., transient light). Before we discuss how exactly we achieve this transient image computation in Sec. 3.4, let us first address fundamental issues of using impulse illumination for it. Using single impulse illumination is inefficient for two reasons. First, the radiance of subsurface scattering decays exponentially with the light path length and the transient light computed from impulse subsurface scattering will have low fidelity. Second, the light paths of these transient lights, especially for inhomogeneous surfaces would be different

depending on which surface point at distance r is used for impulse illumination. Preserving this directionality of subsurface scattering might be of interest for some surfaces, but achieving this with high fidelity is challenging and beyond the scope of our current work.

We resolve these inefficiencies with a special illuminant-observation configuration which we refer to as variable ring light. As depicted in Fig. 2c, consider impulse illuminating all surface points at distance r from the surface point of interest. The normalized radiance of the observed surface point can then be written as the linear combination of all the radiance values due to each of the impulse illuminants (Eq. 1)

$$E(r) = \frac{1}{K} \sum_k^K E(r_k) = \frac{1}{K} \sum_k^K \sum_{|\ell_i| \in \mathcal{L}(r_k)} L(|\ell_i|), \quad (5)$$

where K is the number of impulse illuminant points (i.e., projector pixels corresponding to image pixels lying at distance r as depicted in Fig. 2c right). We are abusing the notion E to avoid clutter by representing both the normalized ring light radiance and the unnormalized impulse illumination radiance, but since $r_k = r$ for all k this notational overwrite should be intuitive. Ring light illumination improves the SNR by K , where K increases with larger r .

Ring light subsurface scattering also effectively averages all the possible light paths in the spatial vicinity of the point of observation. This leads to a more robust observation of internal structures surrounding the observed point as it does not rely on a single arbitrary selected point of incident impulse illumination. Note that any point other than the center of the ring light (i.e., the point of observation) would have superimposed impulse subsurface scattering of different distances. It is only the radiance of the center point that we consider. In practice, we virtually construct ring lights of varying radii around each of the surface points as we discuss in Sec. 4.

3.4 Transient Images from Ring Light Images

Let us consider illuminating the surface with another ring light of slightly larger radius $r + \Delta r$

$$\mathcal{L}(r + \Delta r) = \{|\ell_i| : |\ell_i| \geq r + \Delta r, |II(\ell_i)| = r + \Delta r\}. \quad (6)$$

The actual instantiation of the light rays in these two sets Eq. 4 and 6 are disjoint, because even though they are observed at the same surface point, the entry points are slightly different.

Let us instead assume that the two sets Eq. 4 and 6 approximately overlap for all light with path length longer than $r + \Delta r$

$$\mathcal{L}(r) \cap \mathcal{L}(r + \Delta r) \approx \mathcal{L}(r + \Delta r). \quad (7)$$

This key assumption implies that the radiances of light rays with path length longer than $r + \Delta r$ observed at the same surface point with an impulse illumination at surface distance r are the same as those with an impulse illumination

at surface distance $r + \Delta r$. In other words, the only difference is the light rays of path length shorter than $r + \Delta r$ due to lower bounds on the two sets. This property holds exactly for homogeneous surfaces, as the light rays would accumulate the same spectral radiance regardless of their incident points from the observed surface point for the same path length. Strictly speaking the number of bounces a light ray experiences for a given path length would vary, but statistically speaking (i.e., on average with a tight variance), we may safely assume this to be true.

When Eq. 7 holds, the difference of the normalized observed radiance at the same surface point but illuminated with ring lights of slightly different radii becomes

$$E(r) - E(r + \Delta r) \approx \sum_{|\ell_i| \in \mathfrak{L}(r) \setminus \mathfrak{L}(r + \Delta r)} L(|\ell_i|), \quad (8)$$

where

$$\mathfrak{L}(r) \setminus \mathfrak{L}(r + \Delta r) = \{|\ell_i| : r + \Delta r > |\ell_i| \geq r, |H(\ell_i)| = r\}. \quad (9)$$

Note that \setminus denotes set difference. This means that, as Fig. 1 depicts, by illuminating each surface point with a variable ring light of incrementally increasing radius, which are then assembled into ring light images each of whose imaged surface points encode the radiance for a given ring light radius, and then by taking the difference of them across radius increments, we recover images of light with bounded path lengths, i.e., transient images of subsurface light transport.

When the surface is volumetrically inhomogeneous, i.e., consists of different material regions distributed both in space and depth, Eq. 7 would not necessarily hold, since the path length difference between the two can include integration of subsurface scattering in additional material regions. The discrepancy of spectral radiance of light with path length longer than $r + \Delta r$ between the two observations can, however, be minimized by keeping Δr small—sufficiently smaller than the minimum size of a distinct material region in either the depth or spatial direction. In this case, the resolution of the bounded path lengths would be smaller than the distinct material regions. In other words, the discrete resolution of transient light transport recovered from the ring light image differences is finer than the material composition of the surface volume. As a result, the transient images would still capture important subsurface scattering events sufficiently fine enough to reveal the intrinsic surface structure. It is also important to note that the radiance of longer path length light decays exponentially and thus the discrepancy is usually small in brightness in any case.

Fig. 3 shows transient images computed from variable ring light images and path-length limited renderings for two synthetic inhomogeneous surfaces. Each transient image is normalized in its brightness independently to show its intricate structure, as otherwise they become exponentially darker for longer path lengths. We rendered these synthetic surfaces using Mitsuba renderer [4]. For variable ring light imaging, we first render impulse illumination images using an idealized single-ray light source orthogonal to the surface pointed at each pixel center. By reassembling these impulse illumination images, we compute ring light images




Fig. 3: Top: Transient images computed from synthetic ring light images for two example scenes rendered with Mitsuba [4]. Bottom: Bounded path length images rendered by Mitsuba. The ring light transient images visually match the rendered bounded path length images except for discrepancies expected from its theory and also stemming from rendering noise and aliasing (the dark rim around the white glow in the red bead case).

of varying radii of a pixel increment, and take the difference of ring light images of adjacent radii to compute the transient images. Mitsuba only includes limited bounce rendering, where each bounce occurs at sampled rendering nodes, not statistically distributed particles in the surface. To achieve path-length limited rendering with Mitsuba, we modified its bidirectional path tracing integrator to limit the longest path length of light in the rendering. By taking the difference of these images of upper-bounded path length light, we obtain bounded path length images that can be considered as ground truth.

Unfortunately, rendering subsurface scattering, especially for a complex inhomogeneous surface and also with limits on the path length is challenging as they require almost unrealistically large numbers of samples to properly simulate their light transport. We used 262,144 samples for the path-limited rendering, which takes a day to render and still suffers from excessive noise. For this, the variable ring light imaging results also become noisy. The results in Fig. 3, however, still show that the recovered transient images match the path-length-limited renderings well for both surfaces. For the layered color surface, the integrated light color in each transient image agree, and for the surface with an immersed red bead, the light propagation from the bead (white light expanding out from the red bead due to specular reflection at its surface, followed by red light glowing out of the bead after bouncing inside it⁴) match between the variable ring light and the ground truth. There is a small but noticeable discrepancy for the red bead surface, where the first few images from path-limited renderings do not show the full internal red bead, while the variable ring light result includes light from the entire red bead from the first transient image. For the ring light images the first few ring light images have red light from the bead contributing less for the larger radius ($r + \Delta r$), since the bead is spherical resulting in residual red light in their differences. In contrast, red light would only gradually be observed in the first few path-length limited renderings as light has not reached the red bead. Other than this expected difference and noise, these synthetic results show that variable ring light imaging captures transient subsurface light transport.

⁴ The white glow in the path length limited renderings are difficult to see due to contrast with the dark noise, but they are present.

If the surface is homogeneous in which the scatterers are uniformly distributed in the surface volume, the number of particles along the path is directly proportional to the light path length. Even if the surface is inhomogeneous in its composition, statistically speaking, the number of bounces will monotonically increase with the path length. As a result, we may safely assume that the mean number of bounces a light ray experiences is proportional to the path length. As a result, the radiance difference of two ring light images of different radii bound the mean number of bounces of light. The bound is, however, on the mean, and if the variance of the number of bounces of light for a given path length is large, variable ring light images would not be able to sift out the progression of subsurface scattering. This variance is, however, usually small for general real-world surfaces. In the supplemental material, we provide empirical evidence based on Monte Carlo simulation of the ratio of radiance of n -bounce light of ring light observation as well as theoretical analysis using its analytical approximation that show that this variance is indeed tight and we can safely assume that the recovered transient images approximately correspond to bounded n -bounce light.


4 Experimental Results

In addition to the synthetic surfaces in Fig. 3, we experimentally validate our theory of variable ring light by examining recovered transient images of real-world surfaces.


4.1 Spatial Light Transport in Transient Images

We implement variable ring light imaging with a USB camera (Point Grey Grasshopper3) and a DLP projector (TI DLP Light Commander). The lines of sight and projection are roughly aligned and made perpendicular to the target surface by placing together the projector and camera sufficiently far from the surface. We use telephoto lenses for both the camera and the projector to realize orthographic projection and directional light. We calibrate the projector and camera to associate each projector pixel to image pixels. In our experiments, roughly one projector pixel corresponds to a 2×2 image pixel region. We capture an impulse illumination image for each projector pixel and compute variable ring light images from these images.

The path lengths of light in each transient image is determined by the thickness and spacing of the variable ring light radii. The thicker the ring light, the brighter the observed radiance and thus better fidelity in the recovered transient images. Thick ring lights, however, directly compound the bound on the path length (i.e., r will have a large range in Eq. 9). For this, we set the thickness to be a single projector pixel (e.g., Fig. 2c) to ensure sharp bounds on the path length. The spacing between the different ring light radii, Δr , controls the range of light path lengths we capture in one transient image. We also set this to one projector pixel to recover the highest possible path length resolution. Note that,



(a) Spatial propagation of light in the subsurface



(b) Color variation in subsurface depth

Fig. 4: The transient images, in their pixels that encode bounded integration of light interactions in the subsurface, reveal the volumetric composition of the surface and their colors that are unclear from its outer interface appearance. The difference of the transient images, i.e., the difference of the difference of variable ring light images, reveal the surface color at different depth.

for inhomogeneous surfaces, the condition on Δr regarding Eq. 7 that we discussed can be controlled by ensuring that this single projector pixel is sufficiently smaller than the finest material region. This can be achieved by changing the distance of the projector from the surface.

As shown in Fig. 4a, for spatially inhomogeneous subsurface structures, the recovered transient images capture interesting spatial propagation of light as the corresponding ring light radius increases. For these results and all others, the brightness of each transient image is normalized independently to visualize the details independent of decreasing radiance of longer path length light. The results show, for a red translucent bead inserted in white plastic clay, concentric propagation of white light specularly reflecting off the red bead followed by red light glowing out from the bead, and for a star shape plastic embedded in pink polymer, star shaped light emanating and interreflecting creating petal-like light propagation shapes as the path length increases. These subsurface structures are unclear from the outer surface as shown in the left most images taken with ordinary room light.

4.2 Color from Transient Images

When the surface has subsurface structures that vary across its depth, the recovered transient images will provide appearance samples of its continuous depth variation. Furthermore, as the bounded path length light directly encodes the




Fig. 5: Variable ring light imaging reveals complex intrinsic surface structures in its recovered transient images in which light that has traveled deeper and longer is captured with larger radius ring light. For the skin, wrinkles fade out, deeper red of skin appear, and the arteries become prominent, and for the marble pendant, spatial extents of different color regions emerge as the radius increases.

accumulated color along its path, each of the transient images would reveal the integrated color variation along depth. This spectral integration of transient light suggests that we may reveal the true color at each corresponding surface depth by taking the difference of light with incrementally increasing path lengths. In other words, by taking the difference of the difference of variable ring light images, we are able to reconstruct the color of the subsurface at different depth.

Fig. 4b shows transient images and chromaticity values recovered from their difference for surfaces made of various color layers in depth. The chromaticity trajectories of the transient images reveal the spectral integration of subsurface scattering. We manually selected three transient images at equal spacings, so that each of them likely corresponds to light path length roughly of each layer depth, and computed the mean chromaticities of each image and then their differences to recover the colors at each depth. The recovered colors shown in the right most column match the ground truth shown in the left most image very well. Note that these color variations are not accessible from simple external observation, and worse yet, the same outer surface color could have been made of infinite color combinations inside the subsurface.

4.3 Transient Images of Complex Real Surfaces

Figs. 1 and 5 show recovered transient images of real surfaces with complex subsurface structures and color compositions including those of natural objects. The transient images of each example reveal the subsurface composition both in




Fig. 6: Variable ring light imaging of whole objects with parallel impulse illumination capture. Each transient image is about 800 pixels in width.

its volumetric structure as well as its color variation that are otherwise invisible from the outer surface. Notice the colorless surface reflection and early scattering as well as the complex volumetric propagation and richer color of longer path length light in transient images of larger radius. The supplemental material contains more results and movies of transient images.

Our current implementation of variable ring light imaging is suboptimal in speed, as it requires impulse illumination capture. Note that HDR capture significantly adds to image capture time. This limits us to the 64×64 transient images for reasonable computation time (i.e., around an hour for all image capture and computation). For high resolution image capture of larger surface regions, we can speed up the imaging with parallel impulse illumination capture of surface points at far distances but by limiting the maximum ring light radius to avoid accidentally picking up light from a concurrently illuminated point. Fig. 6 shows example results of variable ring light imaging for large object regions in high resolution. In future work, we plan to fundamentally speed up our implementation by combining spatial and dynamic range multiplexing with ring light imaging.

5 Conclusion

In this paper, we introduced variable ring light imaging for disentangling subsurface scattering light into varying bounded path lengths, and proportionally number of bounces, from external observations. The method does not assume any restrictive analytical scattering model and can be applied to arbitrary real-world surfaces that exhibit subsurface scattering. The theory and its experimental validation demonstrate the effectiveness of variable ring light imaging for unveiling subsurface structures and their colors. We believe the method has significant implications for visual probing of surfaces and for deciphering light transport for richer scene understanding. We plan to derive a radiometrically and computationally efficient implementation in our future work.

Acknowledgements This work was supported in part by JSPS KAKENHI Grant Numbers JP15K21742, JP15H05918, and JP17K20143.

References

1. Bai, J., Chandraker, M., Ng, T.T., Ramamoorthi, R.: A Dual Theory of Inverse and Forward Light Transport. In: European Conference on Computer Vision (2010)
2. Gu, J., Kobayashi, T., Gupta, M., Nayar, S.K.: Multiplexed Illumination for Scene Recovery in the Presence of Global Illumination. In: IEEE International Conference on Computer Vision (2011)
3. Heide, F., Hullin, M.B., Gregson, J., Heidrich, W.: Low-budget Transient Imaging Using Photonic Mixer Devices. *ACM Transactions on Graphics* **32**(4), 45:1–45:10 (2013)
4. Jakob, W.: Mitsuba: Physically based renderer 0.5.0. <http://www.mitsuba-renderer.org/> (2014)
5. Jarabo, A., Masia, B., Marco, J., Gutierrez, D.: Recent advances in transient imaging: A computer graphics and vision perspective. *Visual Informatics* **1**(1), 65–79 (March 2017)
6. Li, D., Nair, A.S., Nayar, S.K., Zheng, C.: AirCode: Unobtrusive Physical Tags for Digital Fabrication. In: ACM Symposium on User Interface Software and Technology (2017)
7. Mukaigawa, Y., Yagi, Y., Raskar, R.: Analysis of light transport in scattering media. In: IEEE Computer Vision and Pattern Recognition (2010)
8. Narasimhan, S., Nayar, S.: Shedding Light on the Weather. In: IEEE Conference on Computer Vision and Pattern Recognition. pp. 665–672 (2003)
9. Nayar, S.K., Krishnan, G., Grossberg, M.D., Raskar, R.: Fast separation of direct and global components of a scene using high frequency illumination. *ACM Trans. on Graphics* **25**(3), 935–944 (July 2006)
10. O’Toole, M., Achar, S., Narasimhan, S.G., Kutulakos, K.N.: Homogeneous Codes for Energy-Efficient Illumination and Imaging. *ACM Trans. on Graphics (SIGGRAPH)* **34**(4), 35:1–35:13 (2015)
11. O’Toole, M., Heide, F., Xiao, L., Hullin, M.B., Heidrich, W., Kutulakos, K.N.: Temporal Frequency Probing for 5D Transient Analysis of Global Light Transport. *ACM Transactions on Graphics* **33**(4), 87:1–87:11 (2014)
12. O’Toole, M., Mather, J., Kutulakos, K.N.: 3D Shape and Indirect Appearance By Structured Light Transport. In: IEEE Conference on Computer Vision and Pattern Recognition (2014)
13. O’Toole, M., Mather, J., Kutulakos, K.N.: 3D Shape and Indirect Appearance by Structured Light Transport. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **38**(7), 1298–1312 (2016)
14. O’Toole, M., Raskar, R., Kutulakos, K.N.: Primal-Dual Coding to Probe Light Transport. *ACM Trans. on Graphics (SIGGRAPH)* (2012)
15. Reddy, D., Ramamoorthi, R., Curless, B.: Frequency-Space Decomposition and Acquisition of Light Transport under Spatially Varying Illumination. In: European Conference on Computer Vision (2012)
16. Redo-Sanchez, A., Heshmat, B., Aghasi, A., Naqvi, S., Zhang, M., Romberg, J., Raskar, R.: Terahertz time-gated spectral imaging for content extraction through layered structures. *Nature Communications* (2016)
17. Seitz, S.M., Matsushita, Y., Kutulakos, K.N.: A theory of inverse light transport. In: IEEE Int’l Conference on Computer Vision (2005)
18. Subpa-asa, A., Fu, Y., Zheng, Y., Amano, T., Sato, I.: Direct and Global Component Separation from a Single Image Using Basis Representation. In: Asian Conference on Computer Vision. pp. 99–114 (2016)

19. Tanaka, K., Mukaigawa, Y., Kubo, H., Matsushita, Y., Yagi, Y.: Recovering Inner Slices of Layered Translucent Objects by Multi-frequency Illumination. *IEEE Transaction on Pattern Analysis and Machine Intelligence* **39**(4), 746–757 (Apr 2017)
20. Velten, A., Wu, D., Jarabo, A., Masia, B., Barsi, C., Joshi, C., Lawson, E., Bawendi, M.G., Gutierrez, D., Raskar, R.: Femto-Photography: Capturing and Visualizing the Propagation of Light. *ACM Transactions on Graphics* **32**(4), 44:1–44:8 (2013)
21. Wu, D., Velten, A., O’Toole, M., Masia, B., Agrawal, A., Dai, Q., Raskar, R.: Decomposing Global Light Transport Using Time of Flight Imaging. *International Journal of Computer Vision* **107**, 123–138 (2014)