

Time-Varying BRDFs

Bo Sun, Kalyan Sunkavalli, Ravi Ramamoorthi, Peter Belhumeur, and Shree Nayar

Abstract—The properties of virtually all real-world materials change with time, causing their BRDFs to be time-varying. However, none of the existing BRDF models and databases take time variation into consideration; they represent the appearance of a material at a single time instance. In this work, we address the acquisition, analysis, modeling and rendering of a wide range of time-varying BRDFs. We have developed an acquisition system that is capable of sampling a material’s BRDF at multiple time instances, with each time sample acquired within 36 seconds. We have used this acquisition system to measure the BRDFs of a wide range of time-varying phenomena which include the drying of various types of paints (watercolor, spray, and oil), the drying of wet rough surfaces (cement, plaster, and fabrics), the accumulation of dusts (household and joint compound) on surfaces, and the melting of materials (chocolate). Analytic BRDF functions are fit to these measurements and the model parameters’ variations with time are analyzed. Each category exhibits interesting and sometimes non-intuitive parameter trends. These parameter trends are then used to develop analytic time-varying BRDF (TVBRDF) models. The analytic TVBRDF models enable us to apply effects such as paint drying and dust accumulation to arbitrary surfaces and novel materials.

Index Terms—BRDFs, time-varying phenomena, measurement, natural phenomena.

I. INTRODUCTION

The appearance of essentially all real-world materials changes with time, often dramatically. Indeed, there are so many different phenomena that give rise to time-varying visual appearance that it is difficult to write down an exhaustive list. Examples include aging of human skin, decaying of flora, corrosion of metals, weathering of surfaces, and aging of materials. In this paper, we focus on those that can be visually described by a time-varying BRDF. In this domain, we explore three categories: drying of paints (watercolor, spray, and oil), drying of wet rough surfaces (fabrics, plaster, and cement) and dust accumulation (household and joint compound). These phenomena are particularly interesting as they are commonplace, are often visually dramatic, and have many practical applications. For example, artistic effects of watercolors, oil and spray paints are often provided by commercial products such as Fractal Design Painter. Drying models are used in vision applications to identify wet regions in photographs [15]. Dust simulation is very popular in driving simulators, games and visualization of interacting galaxies [3], [13].

While there has been a good deal of work on physics-based techniques for simulating time-varying effects due to weathering and aging [6]–[8], this work largely focuses on temporal changes in the diffuse (not specular) texture pattern, developing explicit models for specific effects. These models

require a thorough understanding of the underlying physical processes. The time-varying properties of materials with both specular and diffuse reflectance – such as those considered in this paper – are difficult to model with physics-based techniques because the underlying interactions are often too complex, or not fully understood.

Consider Figure 2. Depending on the specific properties of the medium and its particles, light can be refracted by the liquid-air interface (b, c, d) and reflected by the underlying surfaces (a, b, c). Further, it can also be scattered by dust particles (a), attenuated/reflected by pigments (c), or forward scattered by water droplets (d). Exact simulation of the light transport in these cases, based on the properties of the scattering particles, is too complex in terms of computation, even for a single time instance. Yet, for each case shown in the figure, the material changes with time: the dust layer thickens (a); the paint layer thins (b, c); or the water droplets evaporate (d).

In each of these cases, the change of the BRDF – the directional dependence of reflectance on lighting and viewpoint – cannot be ignored. While there has been considerable work on measuring the BRDFs of real world materials such as [5], [22], [23], [29], these previous efforts only represent the appearance of a material at a single time instance.

In contrast, our work explicitly addresses the acquisition and modeling of *time-varying* BRDFs (TVBRDFs). Central to this work is the measurement of a material’s surface reflectance as it undergoes temporal changes. To record these measurements, we have built a simple robotic rig (Figure 3) to acquire the first time-varying BRDF database, as conventional BRDF measurement devices are too slow to capture the temporal material changes. The system provides very fine sampling along the incident light plane and covers four viewpoints from head-on to angles near grazing (grazing angle specularities may still be missed). It enables us to complete the measurement of each material for one time instance within 36 seconds. The same measurement process is repeated automatically for subsequent time instances.

Our time-varying BRDF database includes the drying of paints (*watercolors*, *oil paint*, and *spray paint*); the drying of wet rough surfaces (*fabrics*, *plaster*, *cement*, and *clay*); the accumulation of dusts (*joint compound* and *household dust*); and miscellaneous time-varying effects such as melting (*chocolate*) and staining (*red wine*). In all we have acquired data for 41 samples. (All of our data in their raw and processed forms is available from www1.cs.columbia.edu/CAVE/databases/tvbrdf/tvbrdf.php.)

For each time instance, our data is carefully fit to the appropriate analytic BRDF functions such as Oren-Nayar, Torrance-Sparrow and a modified Blinn’s dust model, producing a compact set of time-varying parameter curves for each process. We analyze the underlying trends in the parameter curves

The authors are with Columbia University, 500 west 120th street, 450 computer science building, New York, NY 10027. Corresponding e-mail: bosun@cs.columbia.edu



Fig. 1. Rendered images of dust accumulating in a tea set scene, leading to effects such as the diffusing and fading of specular highlights and the shifting of the diffuse component resulting in overall changes in color saturation and hue. The teapot and teacup are rendered with our acquired data, and the table with a novel material showing the same characteristic time-varying behavior. Please refer to Figure 15 for enlarged insets and Section VI for more details.

to derive the first ever set of analytic time-varying BRDF models. These time-varying BRDF models are controlled by a handful of intuitive parameters and are easily integrated into any of the existing rendering packages. Furthermore, we show how the time-varying appearance of one material can be transferred to another, significantly increasing the impact of the data and models presented here. Finally, in addition to temporal variations, we have shown that our model can be combined with simple physics-based control mechanisms to create compelling spatial variations such as dust shadows under occluders, fine dust gradients on curved surfaces and spatial drying patterns as can be observed in Figures 1, 11 and 15.

To summarize, our primary contributions are twofold:

- 1) We introduce an efficient BRDF acquisition system that allows for the capture of time-varying BRDFs. We use this system to acquire the first time-varying BRDF database.
- 2) From our measurements, we develop a set of analytic models for time-varying BRDFs. These models allow time-varying reflectance effects to be incorporated within standard rendering software, transferred to novel materials or controlled spatially by environmental factors.

The rest of the paper is organized as follows. In the next section, we discuss how this work relates to previous work. In Section III, we describe our acquisition rig, the fitting of our data using parametric reflectance models and the time-varying database. In Sections IV, V, and VI, we analyze the trends for the drying of paints, the drying of wet surfaces, and dust accumulation, and respectively develop analytic TVBRDF models. In Section VII, we validate the accuracy of our acquisition and the TVBRDF models. In Section VIII, we compare our work in greater detail with two contemporaneous works [11], [28]. Finally, in Section IX, we present our conclusions and a discussion of future work.

II. PREVIOUS WORK

There is a significant body of research that is closely related to the work presented in this paper. However, the area of time-varying BRDFs has remained largely unexplored. The current paper is a more detailed and extended version of [26], with a more thorough validation (Section VII) and comparison to contemporaneous work (Section VIII). The current paper also explains in depth the modifications to the standard Blinn’s reflectance model for dusty surfaces [1] in Section VI.

Time-Varying Texture Patterns: Time-varying texture patterns have been studied at various levels over the past two decades. Becket et al. [30] modeled surface imperfections through texture specification and generation. Koudelka [16] and Enrique et al. [9] considered a class of data-driven time-varying textures and developed simple algorithms for synthesis and controllability. Others have explicitly modeled the underlying physical/chemical processes such as the formation of metallic patinas [7], aging of stone [6], and appearance changes [8]. Most recently, Lu et al. [19] studied the drying histories of objects based on surface geometries and exposure. Yet, all these methods only focus on the temporally changing spatial pattern of the diffuse albedo and do not address specular reflection of glossy surfaces.

Existing BRDF Models and Databases: Models for surface reflection date as far back as Lambert [17], with numerous models having been developed over the last four decades, e.g., Phong, Torrance-Sparrow, Oren-Nayar, Ward (anisotropic), LaFortune, and Blinn (dust). However, these models treat a material’s reflectance as static – not a function of time. Likewise, BRDF databases have been acquired for real world materials, e.g., CUReT (BRDF) [5], Ward [29], Marschner’s skin measurements [22], and MIT/MERL [23]. However, the materials in these databases were acquired at a fixed time instance and their BRDFs were treated as temporally static.

Paints, Wet Surfaces and Dust: Paints have been well studied in pigmented material modeling. Haase et al. [12] applies the Kubelka-Munk theory of pigment mixing to computer graphics to improve image synthesis. Curtis et al. [4] simulated watercolors with an ordered set of translucent glazes that are generated using a shallow-water simulation. However, these methods do not consider the dynamic drying effects of various paints and cannot capture their specular changes and diffuse color shifts.

For wet materials, the popular L&D model [18] works best for rough solid surfaces, such as blackboards and asphalt. In computer vision, Mall et al. [15] applied the L&D model to the problem of wet surface identification. In computer graphics, Jensen et al. [14] presented a refined optical model incorporating this theory and rendered wet materials using a Monte Carlo raytracer. In addition, other work focuses on specific effects such as wet roadways [24]. However, none of these techniques address “partially wet” surfaces or how drying influences surface appearance.

Dust on diffuse surfaces has also been studied. Blinn [1] introduced a reflectance model for dusty surfaces to the graphics




Fig. 2. In the presence of dust, watercolor fluid, pigmented medium and water, internal light attenuation/reflection/refraction/scattering due to liquid medium and micro particles heavily influence the light paths and completely change the appearance of the material. Moreover, the thickness of the layer of dust, watercolor fluid, pigmented medium and water changes with time.




Fig. 3. A photograph and a diagram of our TVBRDF acquisition system. The blue and green circles show the positions of the four cameras and the light source, respectively. The red and yellow boxes show the sample plate and the robot arm, respectively.

community. Hsu et al. [13] studied dust accumulation using cosine functions and “dust maps” to simulate dust adherence and scratching effects. Chen et al. [2] modeled dust behavior for the purpose of driving simulations. However, the effect of dust on the appearance of glossy surfaces remains unexplored. We show that, unlike diffuse colors, glossy highlights attenuate at a faster exponential rate with increases in the thickness of the dust layer. Moreover, from this exponential decay, one can determine the optical properties of different dust particles.

III. ACQUISITION AND TVBRDF DATABASE

In this section, we describe our acquisition setup in more detail, explain the BRDF models and the fitting algorithm used to fit the raw measurements, and present the time-varying BRDF database.

A. Acquisition System

A key consideration in capturing time-varying BRDFs is to sample the time domain finely enough so as not to miss any important temporal variations. In this respect, previously developed scanning (gantry-type) systems for BRDF acquisition are not suitable as they take a significant amount of time for a single BRDF measurement. Moreover, the angular domain also has to be densely sampled to ensure that high frequency changes due to specularities are captured. Multi-light (dome-type) systems only sparsely sample the lighting directions and hence do not satisfy our sampling requirement. In addition, a practical problem is the influence of gravity on the dusts and liquids that are involved in our time-varying processes.

This makes it difficult to use homogenous spherical samples to expedite the acquisition, as done in [21], [23].

As a result, we are forced to make a trade-off between the time efficiency and the angular density of our acquisition system. To this end, we do not capture all lighting and viewing directions but instead densely measure the BRDF along a single incidence plane and for a small number of viewpoints. One of these viewpoints lies on the incidence plane, which guarantees that the specular highlight is well captured. The remaining viewpoints lie outside the incidence plane. One limitation of this design is that the grazing-angle specularities may be missed. However, for most of our samples, we anticipate little grazing-angle specular effect and hence can minimize the artifacts. Clearly, this approach does not result in a complete (4D) BRDF measurement. To fill in the missing data, the acquired data is fit to analytic BRDF functions. The use of analytic BRDF functions also has the advantage that the TVBRDF of a sample can be compactly represented as a small number of time-varying BRDF parameters.

As shown in Figure 3, our system is composed of four key components: Four remote-head Dragonfly color cameras mounted on an aluminum frame, a sample plate with adjustable tilt, a programmable Adept robot, and a light arm holding a halogen light source and a diffuser. The four cameras lie in a vertical plane. Each camera is 94 cm from the center of the sample plate. In the viewing plane, the cameras have viewing angles of 0° , 25° , 50° and 75° with respect to the vertical axis. All the camera optical axes pass through the center of the sample plate, which is 16.26 cm by 12.19 cm in dimension and has four extensible legs to adjust its height and tilt. All sample materials are prepared as planar patches and placed on a 5.08 cm by 5.08 cm square tray on the plate, as shown in the inset of Figure 3. The light source has a stable radiant intensity and the diffuser is used to make the irradiance uniform over the entire sample. The robot moves the light source around the sample plate along a circle of radius 44 cm.

All the cameras are rigidly fixed and their positions are calibrated. The cameras are also radiometrically calibrated by measuring the radiance of both a Kodak standard color chart and a Gray Spectralon sample, as done in [5]. The cameras are connected to a computer via firewire interface and are synchronized with respect to each other. Additionally, the robot is synchronized with the cameras via a RS232 serial cable and the computer so that the light source position can be determined from the time stamps recorded by the cameras.




Fig. 4. Rendered spheres with time-varying BRDF data captured using our system. To fully illustrate the time-varying phenomena, the renderings use complex natural lighting from the St. Peters environment map.




Fig. 5. Example fits for three of our acquired samples. Each row shows the measurement (red solid lines) obtained in the blue color channel of the top camera for three different time instances, and the results of fitting appropriate analytic BRDF functions (blue dotted lines). Though the changes in BRDFs across time are dramatic, all our fits are found to be fairly accurate with a maximum RMS error of 3.8%. The small gap in the original measurement indicated by the red arrow (around the incidence angle of 12°) is due to the occlusion of the sample by the light source. However, because the sample plane is tilted, the occlusion is shifted away from the peak of the specular lobe and does not affect the robustness of the fitting.

As mentioned earlier, our goal is to capture sharp specularities using the top-most camera that lies on the incidence plane. However, if the sample is placed horizontally, a large part of the specular highlight will be occluded by the light source. To avoid this, we incline the plate by 12° , as shown in the inset of Figure 3. This shifts the specular peak by about 24 degrees with respect to the vertical axis, enabling us to capture the most important portion of it, as shown in Figure 5.

A single scan (circular motion) of the light source takes about 12 seconds, during which time around 360 color images (lighting direction increments of 0.5 degrees) are recorded by each camera. To obtain high dynamic range (HDR) measurements, two more scans are done with all cameras automatically switching to different exposures ranging from 0.2 milliseconds to 32 milliseconds. Therefore, the measurement corresponding to a single time instance of the TVBRDF takes about 36 seconds. To capture a complete TVBRDF, the robot and the cameras are programmed to repeat the above acquisition at a preset time interval (which ranges from 1 minute to 5 minutes in our experiments).

B. Data Fitting

In this section, we focus on the fitting of analytic BRDF functions to our acquired data to obtain a compact set of time-varying model parameters.

Drying Paints and Rough Surfaces: We fit a combination of the Oren-Nayar diffuse model [25], denoted as ρ_d , and the Torrance-Sparrow specular model [27], denoted as ρ_s , to the BRDF measurement obtained from drying paints and drying

d, s	Subscripts for diffuse, specular components
r, g, b	Superscripts for r, g and b color channels
ω_i	Incoming direction
ω_o	Outgoing direction
θ_i, ϕ_i	Elevation and azimuth angles for incident ray
θ_o, ϕ_o	Elevation and azimuth angles for exitant ray
θ_h, ϕ_h	Elevation and azimuth angles for half angle
ρ	Radiance
K_d	Diffuse component albedo
K_s	Specular component albedo
γ	Angle between light source and viewing ray
σ_d	Roughness parameter for diffuse component
σ_s	Roughness parameter for specular component
τ	Optical thickness for dust
w	Dust albedo
Φ	Phase function for dust particles
g	Parameter for the Henyey-Greenstein function

Fig. 6. Notation used in the paper.

wet surfaces. These two models are appropriate because they have a handful of parameters that have physical significance, and are widely used to model glossy or rough surfaces. This combined BRDF model can be written as:

$$\rho(\omega_i, \omega_o; \sigma_d(t), \sigma_s(t), K_d^{r,g,b}(t), K_s(t)) = \rho_d(\omega_i, \omega_o; \sigma_d(t), K_d^{r,g,b}(t)) + \rho_s(\omega_i, \omega_o; \sigma_s(t), K_s(t)) \quad (1)$$

where ω_i and ω_o are the incoming and outgoing directions that are defined in a coordinate frame aligned with the surface normal, σ_s and σ_d are roughness parameters for the specular and diffuse components, respectively, and K_s and $K_d^{r,g,b}$ are the amplitudes of the specular and diffuse components, respectively. More details on the Oren-Nayar and Torrance-Sparrow models are given in Appendices A and B.

The above model has 6 time-varying parameters, namely, the amplitudes of the diffuse and specular components $[K_d^{r,g,b}(t), K_s(t)]$, and the diffuse and specular roughness parameters $[\sigma_d(t), \sigma_s(t)]$. These time-varying parameters exhibit interesting temporal trends which will be discussed later in Sections IV and V.

Dust Accumulation: Blinn's reflectance model [1] appears to be the most relevant for dusty surfaces and relatively simple. The model is also controlled by a few intuitive parameters. We have modified Blinn's model to fit our dust samples. This model can be written as:

$$\begin{aligned} & \rho(\omega_i, \omega_o; g, w_{r,g,b}; \sigma_d, K_d^{r,g,b}; \sigma_s(\tau), K_s(\tau)) \\ &= (1 - T(\tau)) \cdot \rho_{dust}(\omega_i, \omega_o; g, w_{r,g,b}) \\ & \quad + T(\tau) \cdot \rho_d(\omega_i, \omega_o; \sigma_d, K_d^{r,g,b}) \\ & \quad + \rho_s(\omega_i, \omega_o; K_s(\tau), \sigma_s(\tau)), \end{aligned} \quad (2)$$

where:

$$T = e^{-\tau(\frac{1}{\cos \theta_i} + \frac{1}{\cos \theta_r})}. \quad (3)$$

Again, ω_i and ω_o are the incoming and outgoing directions, g is the parameter used in the Henyey-Greenstein phase function, $w_{r,g,b}$ are the dust albedos in the different color channels, K_s and σ_s are the amplitude and roughness for the specular component, $K_d^{r,g,b}$ and σ_d are the amplitudes and roughness

for the diffuse component, and τ is a dimensionless quantity called *optical thickness* which represents the attenuating power of the dust layer. A detailed discussion on the above model and how it relates to Blinn’s model [1] is given in Section VI-A.

Fitting Algorithm: The Levenberg-Marquardt non-linear least-squares optimization algorithm [20] is used to fit the above analytic models to the measured TVBRDF data. For all of our 41 samples, the fits are found to be accurate with a maximum RMS error of 3.8%, as seen from Table I.

C. Database

As shown in Figure 4, we have acquired a variety of samples including watercolors, spray paints, oil paints, fabrics, cement, clay, plaster, joint compound dust, household dust and chocolate. A complete list of our 41 samples and the models used to fit their data is given on the left side of Table I. On the right side of the table are the time intervals, number of temporal samples and the RMS errors in the BRDF fits. The estimated parameter values are not included for lack of space. The minimum time interval between consecutive scans is set to be 1 min because most of the time-varying phenomena that we measured are long-lasting, ranging from 10 mins to 200 mins. In addition, since each single scan takes 36 seconds, continuously scanning the material samples is less meaningful and won’t improve the sampling resolution. All of our measurements and parameter estimates are available for download from www1.cs.columbia.edu/CAVE/databases/tvbrdf/tvbrdf.php.

This database can be directly used in a variety of computer graphics and vision applications.

Figure 5 shows the accuracy of three example fits for our acquired samples. The RMS errors within the some group of paint materials are varying because the different colored paints can have dramatically different makes and hence distinct appearance.

Our goal is to use this database to first identify temporal trends in the estimated parameter values that are associated with each type of time-varying phenomenon (drying paint, drying wet surface, dust accumulation, Sections IV, V, VI). Next, we propose analytic functions that model these temporal trends in parameter values. These models enable us to “apply” several of the above physical processes to novel materials. We validate the accuracy of these analytic TVBRDF models in Section VII. In Section VIII, we compare our work with two contemporaneous publications [11], [28].

IV. DRYING OF PAINTS

Existing scattering theories related to pigmented materials, such as the Kubelka-Munk theory, do not address how the appearance of the material changes as it dries. In this section, we explore the temporal behaviors of the BRDF parameters of our drying paint samples. Based on our analysis, we propose simple analytic models for the parameter variations over time. These models allow us to achieve two effects: We can predict the TVBRDF of a paint of the same type but with a different color as well as the TVBRDF when the paint is applied to a novel surface.

Sample Name and BRDF Model	Interval (mins)	Scans	RMS (%)
Paints - TS+ON			
Krylon Spray Paint			
<i>Flat / White</i>	1	24	0.90
<i>Satin / Green</i>	1	27	1.73
<i>Glossy / Blue</i>	1	40	1.36
<i>Glossy / Red</i>	1	40	0.67
<i>Satin / Dove-Teal</i>	1	30	1.63
Rust-Oleum Spray Paint			
<i>Flat / Yellow</i>	1	40	1.34
Crayola Watercolor			
<i>Blue</i>	1	21	1.27
<i>Red</i>	1	30	1.26
<i>Green</i>	1	30	3.11
<i>Purple</i>	1	40	0.51
<i>Orange</i>	1	40	0.82
<i>Light Green</i>	1	40	2.32
<i>Yellow</i>	1	40	1.20
Daler-Rowney Oil Paint			
<i>Prussian Green</i>	1	10	1.87
<i>Prussian Red</i>	1	10	0.98
<i>Permanent Light Green</i>	1	40	0.66
<i>Cadmium Yellow</i>	1	40	0.26
Drying - TS+ON			
Fabrics			
<i>Alme Grey Blue Fabric</i>	5	30	3.08
<i>Idemo Beige Fabric</i>	5	40	0.35
<i>Ingebo Dark Red Fabric</i>	5	39	0.41
<i>Pink Denim Fabric</i>	1	30	0.10
<i>Orange Cotton Fabric</i>	1	41	0.08
<i>Beige Cotton Fabric</i>	3	40	0.38
<i>Pink Cotton Fabric</i>	3	40	0.22
White Plaster	1	40	0.31
Cement	5	30	2.55
Terracotta Clay	5	30	0.55
Dust - TS+Blinn			
Joint Compound Powder			
<i>Electric Red Exterior Paint</i>	-	15	1.04
<i>Satin / Red Spray Paint</i>	-	11	0.66
<i>Satin / Dove-Teal Paint</i>	-	15	0.26
<i>Flat / Yellow Spray Paint</i>	-	15	0.20
<i>Almas Red Fabric</i>	-	13	0.23
<i>Green Grey Metallic Paint</i>	-	15	2.94
Household Dust			
<i>Electric Red Exterior Paint</i>	-	10	1.25
<i>Satin / Red Spray Paint</i>	-	10	3.62
<i>Satin / Dove-Teal Paint</i>	-	09	0.28
<i>Flat / Yellow Spray Paint</i>	-	11	0.09
<i>Almas Red Fabric</i>	-	10	0.07
<i>Green Grey Metallic Paint</i>	-	10	3.84
Miscellaneous - TS+ON			
Hershey’s Chocolate Melting	1	30	0.48
Red Wine on White Fabric	3	47	0.10

TABLE I

The complete list of 41 samples and their associated effects that are included in our TVBRDF database. “TS” and “ON” stand for the Torrance-Sparrow and Oren-Nayar models, respectively. “Interval” is the time interval between consecutive scans (time instances) and “Scans” is the number of total scans. The RMS errors show the accuracy of the model fits to the acquired measurements over all time instances. The maximum RMS error (over all samples) is found to be 3.84 %.

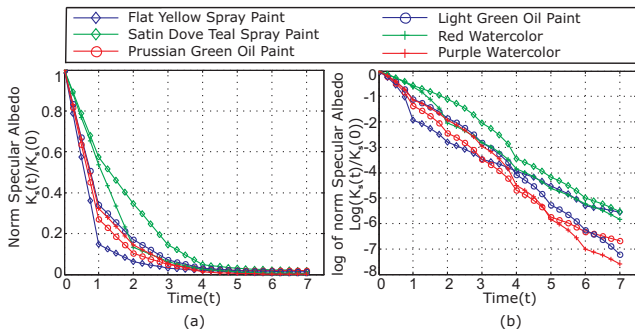


Fig. 7. (a) The fall-off with time of K_s (normalized by the initial value $K_{s,wet}$) of various paint samples. (b) This plot is similar to the one in (a) except that the dry value is first subtracted and then the natural log is applied. Note that K_s in the case of paints attenuates exponentially with time.

A. Temporal Specular Trends

Materials with wet paint applied are highly specular due to strong surface reflection at the liquid-air interface (Figure 2(b) and (c)). As the material dries and the liquid layer thins, the specular component diffuses out and eventually disappears in some cases. This effect is characteristic of the paint-drying process and must be captured by the TVBRDF.

In the Torrance-Sparrow model, the glossiness of a material is governed by two parameters: the specular roughness σ_s and the specular amplitude K_s . Specular highlights of different materials fall off at different rates. In our paint measurements, we observed two important temporal effects. As shown in the linear and log plots in Figure 7(a), K_s (normalized) falls off exponentially from its initial value $K_{s,wet}$ to the value $K_{s,dry}$. After subtracting $K_{s,dry}$, we plot the attenuation of K_s in log scale in Figure 7(b). Note that the temporal variation in this plot is more or less linear, indicating that K_s decreases exponentially with time. The rate of the decrease is given by the slope of the plot, which varies between the paints.

On the other hand, σ_s (after normalization) increases from its initial value $\sigma_{s,wet}$ to $\sigma_{s,dry}$, as shown in Figure 8(a). We plot $1/\sigma_s$ on a linear scale in Figure 8(b) and see that it falls off exponentially with time. In Figure 8(c), $1/\sigma_{s,dry}$ is subtracted from $1/\sigma_s$ and the negative of the log of this quantity is plotted. Note that these plots are more or less straight lines, indicating that σ_s increases exponentially with time. Qualitatively, this agrees with our intuition that as the paint on the material dries, the specularity broadens and fades away.

The exponential forms of K_s and σ_s are strongly coupled and have a rather stable linear relation across different materials. As shown in Figure 8(d), the average slope of this linear relation is around 1. The above observation can be used to develop the following simple analytic model for the temporal variation of the specular parameters of paints:

$$K_s(t) = (K_{s,wet} - K_{s,dry}) \cdot e^{-\lambda t} + K_{s,dry}, \quad (4)$$

$$\sigma_s(t) = \frac{\sigma_{s,wet} \cdot \sigma_{s,dry}}{(\sigma_{s,dry} - \sigma_{s,wet}) \cdot e^{-\lambda t} + \sigma_{s,wet}}, \quad (5)$$

where λ is the effective attenuation rate of the specular component. In the case of a given measurement, λ can be

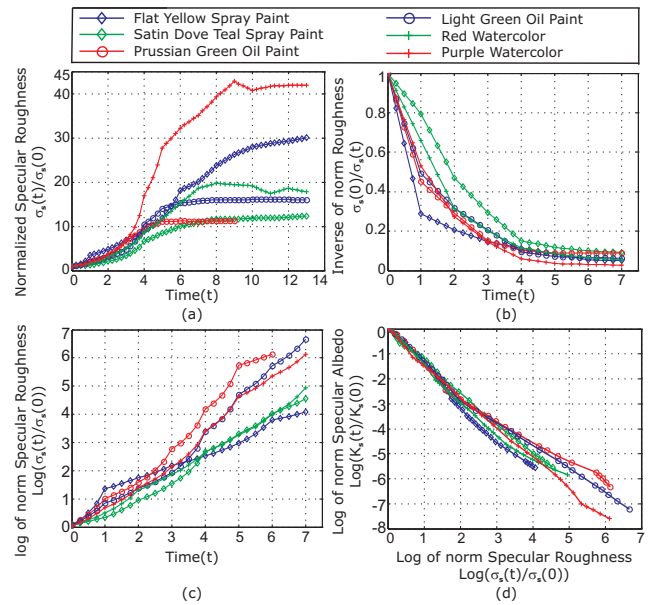


Fig. 8. (a) σ_s (normalized by the initial value $\sigma_{s,wet}$) is plotted as a function of time for several of the paint samples and can be seen to vary exponentially. (b) $1/\sigma_s$ (normalized by the initial value $\sigma_{s,wet}$) plotted as a function of time. (c) $1/\sigma_s$ plotted with a negative natural log scale after $1/\sigma_{s,dry}$ is subtracted. (d) A linear relation is observed between the log of the normalized K_s and σ_s .

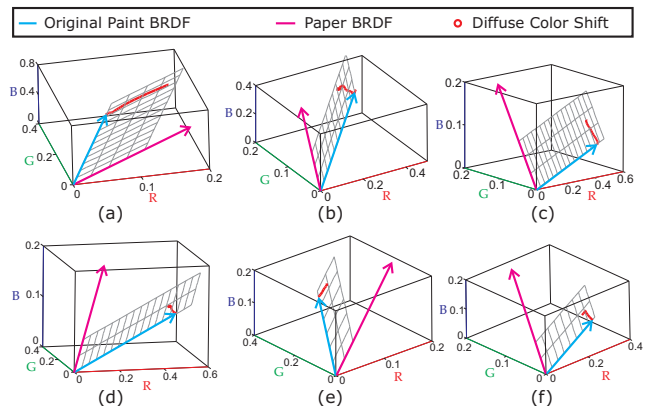


Fig. 9. The diffuse color shifts associated with drying paints lie on dichromatic planes spanned by the diffuse color vector of the surface (shown in magenta) and the diffuse color vector of the paint (shown in cyan). The first row shows the dichromatic planes for watercolors: (a) Blue watercolor, (b) purple watercolor, and (c) red watercolor. The second row shows the dichromatic planes for oil paints: (d) Cadmium yellow oil paint, (e) light green oil paint, (f) Prussian red oil paint.

estimated using the above model. Alternatively, it can be selected by a user when creating a new paint TVBRDF.

B. Temporal Diffuse Trends

In the case of paints, the diffuse color changes are more complicated as they can vary significantly with the types of pigments and solutions in the paint. For example, a watercolor can be modeled using the theory of subtractive color mixing [10] because its colorant is fully dissolved in the solution, making a wet watercolor transparent enough for light to pass through it. The color shifts associated with some of our measured watercolors are shown in Figures 9(a), (b), and (c).




Fig. 10. All diffuse color shifts of wet materials are roughly straight lines connecting the BRDFs when fully dry and wet. (a) Orange Cotton Fabric, (b) Terracotta Clay.

Oil paints, on the other hand, consist of opaque particles that not only absorb but also selectively scatter light energy. Thus, the appearance of an oil paint depends on the sizes, density and shapes of the particles. Figures 9(d), (e), and (f) show the color shifts of some of our measured oil paints. Spray paints, however, cover surfaces with opaque colored spots and therefore show little color variation during drying. Moreover, irrespective of the type of paint, the diffuse color shifts may be affected by more complex factors such as the thickness of the paint coating and how absorbing the underlying surface is.

In the case of our measured paint samples, we found that the diffuse color shifts tend to lie on dichromatic planes in color space, as shown in Figure 9. For a given paint, the dichromatic plane is spanned by the color vectors of the colorant and the underlying surface. This is in line with our intuition that the appearance variation of a painted material should be a combination of the appearances of the paint and the underlying surface. Therefore, a dichromatic decomposition can be applied to separate the diffuse color into a weighted combination of the colors of the paint and the surface:

$$\rho_d(t) = \alpha(t) \cdot \rho_{d,surface} + \beta(t) \cdot \rho_{d,paint}, \quad (6)$$

where $\alpha(t)$ and $\beta(t)$ are the time-varying weights associated with the diffuse radiance $\rho_{d,surface}$ of the surface and the diffuse radiance $\rho_{d,paint}$ of the paint. These two radiances are directly measured from the bare surface and a thick layer of wet paint, respectively. In this way, we have captured all the diffuse temporal variations with two coefficients. This enables us to synthesize the effects of novel paints drying on new surfaces. Further work is needed to develop analytic models for the precise forms of $\alpha(t)$ and $\beta(t)$. For now, we simply use the measured data.

C. Analytic Time-Varying Model for Paints

We have developed an analytic time-varying BRDF model for paints which is given by Equations 4, 5 and 6. The only time-varying parameters are $\alpha(t)$ and $\beta(t)$ in Equation 6. In addition, λ in Equations 4 and 5 is a new parameter that determines the time variation of the specular component. Finally, we need the initial (wet) and final (dry) values of standard parameters such as K_s and σ_s . In practice, all of these parameters can be estimated from a measured TVBRDF.

Alternatively, some of the parameters can be selected by a user to modify the properties of the paint or the underlying

surface. For instance, by changing $\rho_{d,surface}$ and $\rho_{d,paint}$, we can synthesize the drying of a different colored paint on a new surface. We can also change the glossiness of the time-varying material by changing $K_{s,wet}$, $K_{s,dry}$, $\sigma_{s,wet}$ and $\sigma_{s,dry}$, while setting the value of λ to the one estimated from our acquired data.

D. Rendering

Our paint data can readily be used for rendering the effect of drying. The analytic TVBRDF model for paints also enables the transfer of the phenomena to novel materials. Figure 11 shows several models rendered with both acquired and synthesized materials. The dragon is rendered with our acquired blue watercolor drying on white paper. Decomposing this material into a combination of the paint color and the paper color, we can easily replace either of them to synthesize the effect of different colored paint drying on novel surfaces. In this way, we synthesize the effect of green watercolor drying on white paper to render the bunny and the effect of blue watercolor drying on red paper to render the bird (Figure 11). The specular properties of the new materials are transferred by assuming the same exponentially changing rate of K_s and σ_s as the original data, but with different initial values.

Furthermore, a heat source is suspended in the top left corner to control the drying rates of different parts of the models. The drying rate varies inversely with distance from the source. Therefore, the further away the point is, the slower it dries. Additionally, the drying rate varies linearly with the surface orientation with respect to the source so that up-facing parts of the models dry significantly faster than others. As a result, the local “time” variable t for each point ticks at different rates, causing time dependent parameters such as $\alpha(t)$, $\beta(t)$, $K_s(t)$ and $\sigma_s(t)$ to be spatially varying.

As time goes by, the two synthesized materials show changes consistent with the original paint as specular highlights diffuse out and dim and the watercolor layer thins and transmits more color from the underlying surface.

V. DRYING OF WET SURFACES

We can often tell the wetness of an object simply by observing its appearance, because wet surfaces generally lose their color contrast and exhibit significant decrease in color saturation. We apply a similar analysis to wet surfaces as we did for paints and develop their analytic TVBRDF model. Wet surfaces in our experiments refer mostly to rough and diffuse materials quenched in water and having a very thin water layer on their surfaces.

The specular highlights of most wet surfaces vanish very quickly, as seen in Figure 4(d), and can be ignored for subsequent time instances (grazing angle specularities may be missed in our measurements). However, in the case of glossy underlying surfaces or a thick water layer, the specular reflectance can be important. We leave the acquisition and analysis of this effect to future exploration. Diffuse color, on the other hand, exhibits significant time variations. For most of our acquired wet materials, the diffuse color shifts in the color space are more or less straight lines, as shown in Figure




Fig. 11. Objects painted with watercolors dry over time. The dragon is rendered with our acquired blue watercolor on white paper. The bunny is rendered with a synthesized green watercolor on white paper. The bird is rendered with the blue watercolor on synthesized pink paper. The white sphere suspended in the corner represents a heat source.

10. This simple observation enables us to derive the analytic TVBRDF model for wet materials:

$$\rho_d(t) = \alpha(t) \cdot \rho_{d,dry} + (1 - \alpha(t)) \cdot \rho_{d,wet}, \quad (7)$$

where $\rho_{d,dry}$ and $\rho_{d,wet}$ are the albedos of the material when fully dry and wet, respectively, and $\alpha(t)$ can be estimated from our measured data. Over time, $\alpha(t)$ behaves as a sigmoid function as shown in Figure 12, confirming Lu et al.'s earlier results [19] for the specific case of drying stone.

VI. DUST ACCUMULATION

Dust is ubiquitous in our visual experience. Based on the temporal trends that we have observed, we develop a simple analytic TVBRDF model for dust which can generalize the dusty effect to arbitrary surfaces. This also shows that our analysis approach can be extended to a very different BRDF model from the Torrance-Sparrow + Oren-Nayar model considered in the previous two sections.

A. Dust Reflectance Model

We initially tried to use the top-lit reflectance model for dusty surfaces proposed by Blinn in [1] for fitting. A more detailed description of this model is in Appendix C. In essence, it is a weighted blending of two terms: the dust reflectance ρ_{dust} and the original surface reflectance $\rho_{surface}$.




Fig. 12. Plots of $\alpha(t)$ versus time for two sample materials.

The blending factor is a term called transparency T . The model can be briefly written as:

$$\rho = (1 - T_1) \cdot \rho_{dust} + T_2 \cdot \rho_{surface}, \quad (8)$$

$$T_1 = e^{-\tau \frac{1}{\cos \theta_r + \cos \theta_i}}, \quad (9)$$

$$T_2 = e^{-\tau \frac{1}{\cos \theta_r}}. \quad (10)$$

We can use the model directly. However, there are two important issues that need to be addressed. Firstly, though the first transparency term T_1 introduced in [1] correctly takes the dependency on both the lighting and view angles θ_i, θ_r into account, the second term T_2 , describing the transmission of the reflectance from the underlying surface, is asymmetrical and solely depends on the viewing angle θ_r . It ignores the fact that light rays have to pass through the dust as well as




Fig. 13. Plots showing results of fitting both the original and modified Blinn’s dust model to the measured data. It is clear that the modified model fits better consistently. (a) Red color channel from the third camera for Household Dust on Satin Dove Teal Paint (b) Red color channel from the top camera for Joint Compound on Almas Red Fabric.

exit it. Therefore, we reconcile the difference between the T_1 and T_2 by incorporating the lighting angle θ_i into T_2 . Both the transparency terms T_1 and T_2 can be consolidated using a single transparency T :

$$T(\tau) = e^{-\tau \left(\frac{1}{\cos \theta_i} + \frac{1}{\cos \theta_r} \right)}. \quad (11)$$

To compare the accuracies, we ran fits with both models across all our dust samples. The result clearly shows that our model fits better, as seen in Figure 13.

Second, the model defined by Equation 8 was originally developed for diffuse surfaces and does not address glossy highlights due to the difficulty in explicit modeling. Our experiments suggest that the specular component falls off at a much faster rate than the diffuse component. Therefore, we fit the specular reflectance ρ_s separately from the whole surface reflectance $\rho_{surface}$. We rewrite Equation 8 and expand some of its arguments:

$$\begin{aligned} & \rho(\omega_i, \omega_o; g, w_{r,g,b}; \sigma_d, K_d^{r,g,b}; \sigma_s(\tau), K_s(\tau)) \\ &= (1 - T(\tau)) \cdot \rho_{dust}(\omega_i, \omega_o; g, w_{r,g,b}) \\ & \quad + T(\tau) \cdot \rho_d(\omega_i, \omega_o; \sigma_d, K_d^{r,g,b}) \\ & \quad + \rho_s(\omega_i, \omega_o; K_s(\tau), \sigma_s(\tau)), \end{aligned} \quad (12)$$

where ω_i and ω_o are the incoming and outgoing directions, g is a parameter that controls the phase function of the dust particles, $w_{r,g,b}$ are the dust albedos for different color channels, K_s and σ_s are the amplitude and roughness for the specular component, $K_d^{r,g,b}$ and σ_d are the amplitudes and roughness for the diffuse component, and τ is a dimensionless quantity called *optical thickness* which describes the attenuation power of the dust layer.

The dust reflectance ρ_{dust} depends on a few other static parameters such as the dust albedo $w_{r,g,b}$ and g in the Henyey-Greenstein phase function. These parameters need to be computed only once from pure dust. Similarly, the static BRDF parameters for the underlying surfaces such as σ_d and K_d can be estimated from the materials when they are completely dust free.

Levenberg-Marquardt method [20] is used to fit our modified model to the acquired data. The fitting result for the combined BRDF contains only two time-varying parameters: $[K_s(\tau)]$ and $[\sigma_s(\tau)]$.




Fig. 14. (a) The natural log of K_s normalized by its initial value K_{s0} versus optical thickness τ . It clearly shows that K_s decays exponentially with the optical thickness τ and the slopes of the lines depend on the types of dust. (b) Absolute values of σ_s versus optical thickness τ .

B. Temporal Parameter Trends

The change in diffuse colors due to dust has been addressed by the modified transparency T in Equation 11. Therefore, we focus on the behavior of the specular highlights subject to dust.

Due to complicated interactions such as scattering and inter-reflection inside the dust layer, specular highlights should not be attenuated by exactly $\exp(-\tau)$. However, considering that dust is primarily a single scattering medium and the net effect of dust on a surface is observed from a distance, we estimate that the specularities should still fall off at an exponential rate.

After normalizing the specular amplitude K_s by its initial value, we plot it in log scale versus the optical thickness τ . As Figure 14(a) shows, the log value of specular parameter K_s decreases essentially linearly with the optical thickness τ , which confirms our intuition about the exponential decay. Moreover, the slopes of these curves actually relate to the scattering properties of the dust particles and are dust dependent. This can be modeled by the effective specular optical thickness λ . On the other hand, most changes of σ_s are rather small as shown in Figure 14(b), and thus can be treated as negligible.

C. Analytic Time Varying BRDF Model for Dust

Based on the temporal trends mentioned above, we have developed an analytic TVBRDF model for dust:

$$\rho(\tau) = (1 - T(\tau)) \cdot \rho_{dust} + T(\tau) \cdot \rho_d + e^{-\lambda\tau} \cdot \rho_s, \quad (13)$$

where $T(\tau)$ is the modified transparency term as defined in Equation 11 and blends the dust color with the diffuse reflectance of the surface. The specular reflectance of the surface is attenuated at an exponential rate described by the effective specular optical thickness λ . If the specular component behaved the same way as the diffuse component, λ would equal 1. However since specular highlights fall off faster, in practice, λ is greater than 1 and depends on the type of the dust particles. In our database, λ is about 11 for joint compound dust and $\frac{10}{3}$ for household dust. In Figure 15, the teapot and teacup are rendered directly using our acquired data while the material of the table is synthesized with a low specular exponent (still using the λ for household dust from our data).




Fig. 15. A tea set scene accumulating dust across time and its close ups (bottom). A sequence of this scene across time is shown in Figure 1. Please note effects such as intricate dust shadows under the teacup and the teapot and its knob (white arrows), fine dust gradient on the teapot body and diffusing specularities on the saucer and teapot (yellow haloed arrows).

D. Rendering and Physics Controls

Dust accumulation is affected by many factors, including wind, the position and orientation of the dust source, the inclination, stickiness and exposure of a surface and its contact with other objects, as discussed in Hsu et al. [13]. With our analytic time-varying BRDF model for dust, different physics control mechanisms only need to modify τ spatially to generate compelling spatially varying effects. In Figure 15, a tea set scene is accumulating dust cast from a circular dust source above. The effect of gravity and surface inclination on the rate of dust accumulation is modeled by the cosine of the angle between the surface normal and the vertical axis. As a result, the dust is not evenly distributed and less accumulates on steeper surfaces – for instance, on the side of the teapot and teacup. Wherever dust is present, the surface becomes dull and specular highlights are attenuated, as shown by the saucer and the teapot in the insets.

Further, due to occlusion, the surface exposure at all points is computed as the solid angle subtended by the dust source and is used to linearly control the rate of dust accumulation. Certain areas exhibit a “dust shadow” effect and remain shiny and in high contrast across time – for example, the areas just under the teapot and its knob, under the teacup and on the saucer, and on the bottom side of the table, as shown in the insets.

VII. VALIDATION

We validate both the acquisition and the analytic TVBRDF models presented in earlier sections.

A. Validation of Acquisition

In Section III, we quantitatively showed the accuracy of our fits to the original measurements in Figure 5. To further test the fidelity of our data to the real materials, we take photographs




Fig. 16. Photos of a real sphere (top row) and rendered images of a virtual sphere covered with Satin Dove Teal Spray Paint (bottom row). Across time, both spheres show similar appearance changes, such as the diffusing out of specular highlights. Color differences between the two spheres are due to distinct lighting environments.

of actual painted spheres and compare them with rendered images using the parameters in our database. As Figure 16 shows, despite different external factors such as the lighting, our data qualitatively agrees with real observed material appearance. Across time, the specular highlights broaden and diffuse out on both the real and rendered spheres. Slight color differences can be observed because of distinct lighting conditions.

B. Validation of analytic TVBRDF Models

In this section, we test the effectiveness of our TVBRDF models in capturing the temporal variations of the original measurements. We focus on the TVBRDF models for paints and dusts. For qualitative comparisons, we render teapots with both the original measurements and the analytic TVBRDF models and compare their visual differences. For quantitative comparisons, we plot the parameter curves of the analytic TVBRDF models against the the original measured data.

For the TVBRDF model of paints, Figure 17 shows two teapots rendered respectively with the captured data of a light green oil paint and the analytic model at different time instances. Though the specular highlights predicted by the analytic model grows slightly stronger in the middle column, the overall temporal variations of the visual appearances of the teapots are very similar. For quantitative comparison, we compare the specular component of the measurements and the analytic TVBRDF model (since the diffuse component is accurately decomposed by $\alpha(t)$ and $\beta(t)$ in the model). Figure 18 shows K_s and σ_s predicted by our analytic TVBRDF model against the measured data in log scale. Different line slopes are described by λ , the effective attenuation rate of the specular component. Despite slight deviations of the data from the model, such as that of the Satin Dove Teal Spray Paint in Figure 18, they largely agree on their overall variations across time. The misalignment of the data with the model could be due to many reasons such as acquisition noise, surface properties, airflows and even some micro-level chemical reactions.




Fig. 17. Two teapots rendered with the measured data (top row) and the analytic TVBRDF model (bottom row) for light green oil paint at different time instances.




Fig. 18. The temporal variation of K_s (left) and σ_s (right) of both the paint measurements and their analytic TVBRDF model. Curves of the same material have the same color, but with markers to distinguish the measurements and the TVBRDF model. All curves are normalized and then plotted in the natural log space.

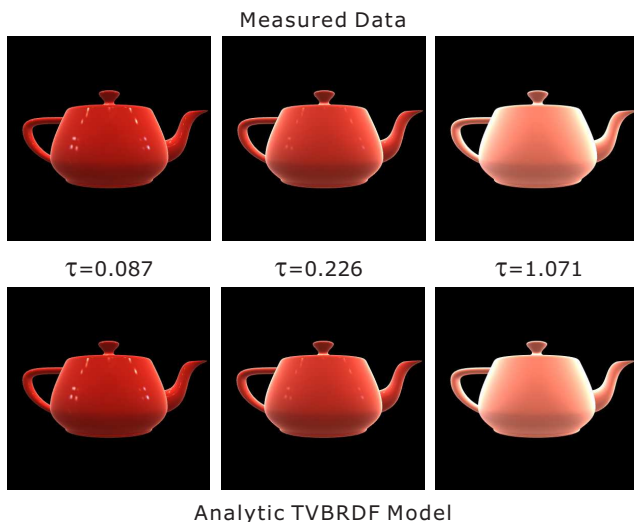


Fig. 19. Two teapots rendered with the measured data (top row) and the analytic TVBRDF model (bottom row) for joint compound powder dust on electric red paint at different time instances.




Fig. 20. The attenuation of K_s with time for both the dust measurements and its analytic TVBRDF model. Black lines are the TVBRDF models, while the marked color curves are the measurements. All curves are normalized and then plotted in the natural log space.

To validate the TVBRDF model for dust, we first show two teapots rendered with the acquired joint compound powder on electric red paint and the model in Figure 19. The two teapots look visually similar across time. The quantitative validation tests the model's specular component. In Figure 20, we plot K_s best predicted by our model in comparison with the measurements in log scale. Though the curves of different materials do not exactly follow the straight lines of the model, their temporal behavior can be well approximated by the TVBRDF model within a large range of the optical thickness.

VIII. COMPARISON WITH CONTEMPORANEOUS WORK

We compare our work with two most recent publications, Wang et al. [28] and Gu et al. [11], in terms of research goals and results, acquisition setups and BRDF models used for fitting.

Wang et al. [28] studied the appearance manifolds that model the time-variant surface appearance of a material. They observed that concurrent variations in appearance over a surface represent different degrees of weathering. They started from a weathered material sample at a single time instance, and inferred its spatial and temporal variations during the weathering process.

Gu et al. [11] explored a number of natural processes that cause the surface appearance to vary, such as burning, decay, corrosion, drying and rusting. A multi-light system with 16 Basler cameras and 150 white LED light sources was used to capture the time-varying surface appearance. A Space-Time Appearance Factorization (STAF) model was also proposed to factor the space and time-varying effects.

Our work differs from these contemporaneous works primarily in three aspects. (1) They focus primarily on spatial patterns due to spatially different rates of weathering for general time-varying process, but do not discover any trends between different materials (each process/sample is treated completely independently). On the other hand, we do not consider spatial variation, but instead study temporal BRDF variations

more carefully and discover trends and analytic models for TVBRDFs among different specific types of processes (paints drying, dust) (2) In terms of acquisition, Wang et al. [28] use the data of a weathered material at a single time instance, and use simpler linear light-source reflectometry. By contrast, Gu et al. [11] and we consider time-varying acquisition, and both methods use rapid acquisition setups. They use a dome-type system with a sparse sampling of the full light-view space, while we use a robotic rig to densely sample the BRDF along a single incidence plane and for a small number of viewpoints to capture high-frequency specularities accurately (while they also fit parametric models, the sparse sampling means they cannot make quantitative claims regarding the accuracy of their specular fits). (3) The isotropic Ward model [29] was used in [28] to fit all materials. Gu et al. [11] used a combination of diffuse Lambertian and simplified Torrance-Sparrow model for fitting. In our work, more complicated BRDF models such as the Blinn's dust model are also used to study the appearance of different types of materials.

IX. CONCLUSIONS AND FUTURE WORK

We have introduced the notion of time-varying BRDFs, and have for the first time captured, modeled and rendered such temporal changes in appearance. A major result of our work is a comprehensive database which is currently accessible from www1.cs.columbia.edu/CAVE/databases/tvbrdf/tvbrdf.php. Our data can be directly used for many important qualitative time-varying effects such as drying, dusting and melting, in any standard rendering package. Moreover, we have analyzed the temporal trends of the model parameters, and developed analytic TVBRDF models which are useful in extending these time-varying phenomena to novel materials.

We are interested in exploring many related aspects of time-varying BRDFs. One avenue would be to incorporate time-varying BRDFs into existing Precomputed Radiance Transfer methods for real-time rendering. An alternative direction can be to couple important appearance changes such as burning and melting with physical simulation of those processes.

ACKNOWLEDGMENT

This research was funded in part by a Sloan Research Fellowship and NSF grants #0305322 and #0446916.

REFERENCES

- [1] J. Blinn. Light reflection functions for simulation of clouds and dusty surfaces. In *SIGGRAPH 82*, pages 21–29, 1982.
- [2] J. X. Chen and X. Fu. Integrating physics-based computing and visualization: Modeling dust behavior. *Computing in Science and Engg.*, 1(1):12–16, 1999.
- [3] J. X. Chen, X. Fu, and J. Wegman. Real-time simulation of dust behavior generated by a fast traveling vehicle. *ACM Trans. Model. Comput. Simul.*, 9(2):81–104, 1999.
- [4] C. J. Curtis, S. E. Anderson, J. E. Seims, K. W. Fleischer, and D. H. Salesin. Computer-generated watercolor. In *SIGGRAPH '97*, pages 421–430, 1997.
- [5] K.J. Dana, B. Van-Ginneken, S.K. Nayar, and J.J. Koenderink. Reflectance and Texture of Real World Surfaces. *ACM Transactions on Graphics (TOG)*, 18(1):1–34, Jan 1999.
- [6] J. Dorsey, A. Edelman, H. Jensen, J. Legakis, and H. Pedersen. Modeling and rendering of weathered stone. In *SIGGRAPH 99*, pages 225–234, 1999.

- [7] J. Dorsey and P. Hanrahan. Modeling and rendering of metallic patinas. In *SIGGRAPH 96*, pages 387–396, 1996.
- [8] J. Dorsey, H. Kohling Pedersen, and P. Hanrahan. Flow and changes in appearance. In *SIGGRAPH 96*, pages 411–420, 1996.
- [9] S. Enrique, M. Koudelka, P. Belhumeur, J. Dorsey, S. Nayar, and R. Ramamoorthi. Time-varying textures: Definition, acquisition, and synthesis. Technical Report CUCS-023-05, Columbia University.
- [10] R. M. Evans. An introduction to color. Wiley, 1948.
- [11] J. Gu, C. Tu, R. Ramamoorthi, P. Belhumeur, W. Matusik, and S. Nayar. Time-varying surface appearance: Acquisition, modeling, and rendering. In *SIGGRAPH '06*, 2006.
- [12] C. S. Haase and G. W. Meyer. Modeling pigmented materials for realistic image synthesis. *ACM Trans. Graph.*, 11(4):305–335, 1992.
- [13] S. Hsu and T. Wong. Simulating dust accumulation. *IEEE Comput. Graph. Appl.*, 15(1):18–22, 1995.
- [14] H. W. Jensen, J. Legakis, and J. Dorsey. Rendering of wet materials. In *Rendering Techniques 99*, pages 273–282, 1999.
- [15] H. B. Mall Jr. and N. da Vitoria Lobo. Determining wet surfaces from dry. In *ICCV '95: Proceedings of the Fifth International Conference on Computer Vision*, pages 963–968, 1995.
- [16] M. L. Koudelka. *Capture, Analysis and Synthesis of Textured Surfaces With Variation in Illumination, Viewpoint, and Time*. PhD thesis, Yale University, 2004.
- [17] J. H. Lambert. *Photometria sive de mensura de gratibus lumi-nis. colorum umbrae*, 1760.
- [18] J. Lekner and M.C. Dorf. Why some things are darker when wet. *Applied Optics*, 27:1278–1280, 1988.
- [19] J. Lu, A. S. Georghiadis, H. Rushmeier, J. Dorsey, and C. Xu. Synthesis of material drying history: Phenomenon modeling, transferring and rendering. In *Eurographics Workshop on Natural Phenomena*, 2005.
- [20] D. Marquardt. An algorithm for least-squares estimation of nonlinear parameters. *SIAM J. Appl. Math.*, 11:431–441, 1963.
- [21] S. R. Marschner, S. H. Westin, E. P. F. Lafortune, and K. E. Torrance. Image-based bidirectional reflectance distribution function measurement. *Applied Optics*, 39:2592–2600, 2000.
- [22] S. R. Marschner, S. H. Westin, E. P. F. Lafortune, K. E. Torrance, and D. P. Greenberg. Image-based BRDF measurement including human skin. In *In Proceedings of 10th Eurographics Workshop on Rendering*, pages 139–152, 1999.
- [23] W. Matusik, H. Pfister, M. Brand, and L. McMillan. A data-driven reflectance model. *ACM Transactions on Graphics*, 22(3):759–769, July 2003.
- [24] E. Nakamae, K. Kaneda, T. Okamoto, and T. Nishita. A lighting model aiming at drive simulators. In *SIGGRAPH '90*, pages 395–404, 1990.
- [25] M. Oren and S.K. Nayar. Generalization of Lambert's reflectance model. In *SIGGRAPH 94*, pages 239–246, Jul 1994.
- [26] B. Sun, K. Sunkavalli, R. Ramamoorthi, P. Belhumeur, and S. Nayar. Time-varying BRDFs. In *Eurographics Workshop on Natural Phenomena*, pages 15–24, 2006.
- [27] K. Torrance and E. Sparrow. Theory for off-specular reflection from rough surfaces. *Journal of the Optical Society of America*, 57:1105–1114, Sep 1967.
- [28] J. Wang, X. Tong, S. Lin, M. Pan, H. Bao, B. Guo, and H. Shum. Appearance manifolds for modeling time-variant appearance of materials. In *SIGGRAPH '06*, 2006.
- [29] G. Ward. Measuring and modeling anisotropic reflection. In *SIGGRAPH '92*, pages 265–272, 1992.
- [30] W.Becket and N. Badler. Imperfection for realistic image synthesis. *The Journal of Visualization and Computer Animation*, 1:26–32, 1990.

APPENDIX

a) Appendix A: Oren-Nayar Diffuse Reflectance Model:

The Oren-Nayar reflectance model was designed for rough surfaces. The model is composed of two parts: the direct illumination component and the inter-reflection component. The direct illumination component in the radiance for this model is given by

$$\begin{aligned}
 & \rho_d^1(\theta_r, \theta_i, \phi_r - \phi_i; \sigma_d, K_d) \\
 = & \frac{K_d}{\pi} \left[C_1(\sigma_d) + \cos(\phi_r - \phi_i) C_2(\alpha, \beta, \phi_r - \phi_i, \sigma_d) \tan \beta \right. \\
 & \left. + (1 - |\cos(\phi_r - \phi_i)|) C_3(\alpha, \beta, \sigma_d) \tan\left(\frac{\alpha + \beta}{2}\right) \right] \quad (14)
 \end{aligned}$$

where,

$$\alpha = \max(\theta_r, \theta_i) \quad (15)$$

$$\beta = \min(\theta_r, \theta_i) \quad (16)$$

$$C_1 = 1 - 0.5 \frac{\sigma_d^2}{\sigma_d^2 + 0.33} \quad (17)$$

$$C_2 = \begin{cases} 0.45 \frac{\sigma_d^2}{\sigma_d^2 + 0.09} \sin \alpha & \cos(\phi_r - \phi_i) \geq 0 \\ 0.45 \frac{\sigma_d^2}{\sigma_d^2 + 0.09} (\sin \alpha - (\frac{2\beta}{\pi})^3) & \text{otherwise} \end{cases} \quad (18)$$

$$C_3 = 0.125 \left(\frac{\sigma_d^2}{\sigma_d^2 + 0.09} \right) \left(\frac{4\alpha\beta}{\pi^2} \right)^2 \quad (19)$$

The inter-reflection component is given by

$$\begin{aligned} & \rho_d^2(\theta_r, \theta_i, \phi_r - \phi_i; \sigma_d, K_d) \\ &= 0.17 \frac{K_d^2}{\pi} \frac{\sigma_d^2}{\sigma_d^2 + 0.13} \left[1 - \cos(\phi_r - \phi_i) \left(\frac{2\beta}{\pi} \right)^2 \right] \end{aligned} \quad (20)$$

These two components combine to give the total diffuse surface radiance.

$$\begin{aligned} & \rho_d(\theta_r, \theta_i, \phi_r - \phi_i; \sigma_d, K_d) \\ &= \rho_d^1(\theta_r, \theta_i, \phi_r - \phi_i; \sigma_d, K_d) + \rho_d^2(\theta_r, \theta_i, \phi_r - \phi_i; \sigma_d, K_d) \end{aligned} \quad (21)$$

b) Appendix B: Torrance-Sparrow Specular Reflectance Model: The Torrance-Sparrow specular model is expressed by specular component amplitude, facet normal distribution, geometrical attenuation and fresnel reflection terms as

$$\rho_s = \frac{K_s \cdot D \cdot G \cdot F}{4 \cos \theta_i \cos \theta_r} \quad (22)$$

where K_s is the specular component amplitude, D describes the distribution of facet normals over the surface and G is a geometrical attenuation factor.

$$D = e^{-(\theta_h/\sigma_s)^2} \quad (23)$$

$$G = \max(0, \min(1, \frac{2 \cos \theta_i \cos \theta_h}{\cos \theta_i \cos \theta_h + \sin \theta_i \sin \theta_h \cos(\phi_i - \phi_h)}, \frac{2 \cos \theta_r \cos \theta_h}{\cos \theta_r \cos \theta_h + \sin \theta_r \sin \theta_h \cos(\phi_r - \phi_h)})) \quad (24)$$

F is the fresnel reflection term and depends on the refractive index n of the material. We have set the Fresnel term to 1 for convenience of measurement and fitting.

c) Appendix C: Top-lit Dust Reflectance Model: The dust reflectance ρ_{dust} is from the top lit brightness function in [1] which is given by

$$\rho_{dust}(\theta_r, \theta_i, \phi_r - \phi_i; g, w_{r,g,b}) = w_{r,g,b} \Phi(\gamma) \frac{\cos \theta_i}{(\cos \theta_i + \cos \theta_r)} \quad (25)$$

where γ is computed as the angle between the light and viewing ray. Φ is the popular Henyey-Greenstein phase function which describes the dependence of scattering on deviation angle γ .

$$\Phi(\gamma, g) = \frac{1 - g^2}{(1 + g^2 - 2g \cos \gamma)^{3/2}} \quad (26)$$

This is the equation of an ellipse in polar coordinates, centered at one focus. The parameter g is the eccentricity of the ellipse and is a property of the material. When g equals 0, scattering is isotropic. When g is greater than 0, it is predominantly forward scattering.