

Time-Varying BRDFs

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

Department of Computer Science, Columbia University, New York, NY 10027, U.S.A.

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.

1. 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 [JdVL95]. Dust simulation is very popular in driving simulators, games and visualization of interacting galaxies [HW95, CF99].

While there has been a good deal of work on physics-based techniques for simulating time-varying effects due to weathering and aging [DH96, DPH96, DEJ*99], 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.

Depending on the specific properties of the medium and its particles, light can not only be refracted by the liquid-air interface and reflected by the underlying surfaces, but can also be scattered by dust particles, attenuated/reflected by pigments, or forward scattered by water droplets. Exact simulation of the light transport in these cases, based on the properties of the scattering particles, is too complex in terms of computations, even for a single time instance. Yet, the material changes with time.

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 [War92, DVGNK99, MWL*99, MPBM03], 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

[†] bosun@cs.columbia.edu



Figure 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 11 for enlarged insets and Section 6 for more details.

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 2) 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. 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 will be made available at the time of publication.)

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 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, 8 and 11.

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.

2. 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.

Time-Varying Texture Patterns: Time-varying texture patterns have been studied at various levels over the past two decades. [WB90] modeled surface imperfections through texture specification and generation. [Kou04] and [EKB*05] 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 [DH96], aging of stone [DEJ*99], and appearance changes [DPH96]. Most recently, [LGR*05] 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. Moreover, [GTR*06, WTL*06] focus on the change in the texture of a surface with time and a simple BRDF model is used to model changes with lighting and viewpoint. We study temporal BRDF variations more carefully and explore a different set of phenomena, using more complicated BRDF models, e.g., the dust model in Section 6.

Existing BRDF Models and Databases: Models for surface reflection date as far back as Lambert [Lam60], 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) [DVGNK99], Ward [War92], Marschner’s skin measurements [MWL*99], and MIT/MERL [MPBM03]. 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. [HM92] applies




Figure 2: 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.

the Kubelka-Munk theory of pigment mixing to computer graphics to improve image synthesis. [CAS*97] 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 [LD88] works best for rough solid surfaces, such as blackboards and asphalt. In computer vision, [JdVL95] applied the L&D model to the problem of wet surface identification. In computer graphics, [JLD99] 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 [NKON90]. However, none of these techniques address "partially wet" surfaces or how drying influences surface appearance.

Dust on diffuse surfaces has also been studied. [Bl82] introduced a reflectance model for dusty surfaces to the graphics community. [HW95] studied dust accumulation using cosine functions and "dust maps" to simulate dust adherence and scratching effects. [CF99] 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 tell the optical properties of different dust particles.

3. Acquisition and TVBRDF Database

3.1. 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

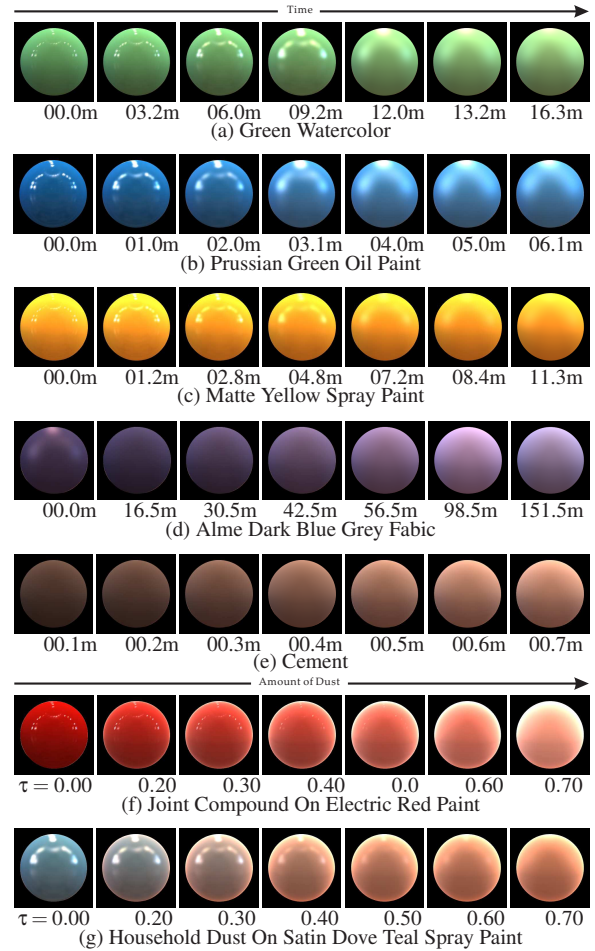


Figure 3: Rendered spheres with time-varying BRDF data captured using our system. To fully illustrate the time-varying phenomena, the renderings use complex natural lighting from the St. Peters environment map.

lighting direction 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 [MWLT00, MPBM03].

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. 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.

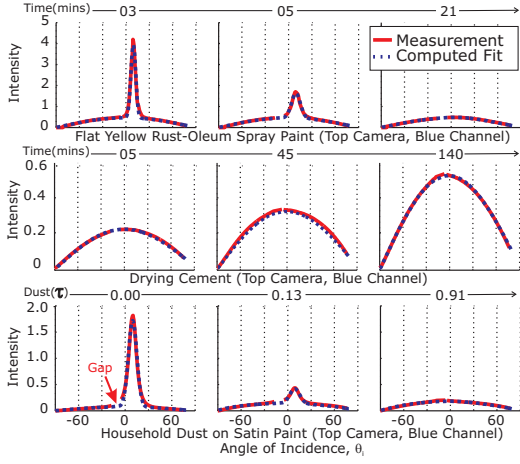


Figure 4: Example fits for three of our acquired samples. Each row shows the measurement (red solid lines) obtained in the blue color channel of the top camera for three different time instances, and the results of fitting appropriate analytic BRDF functions (blue dotted lines). Though the changes in 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 shown in Figure 2, our system is composed of four key components: Four 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 2. 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 [DVGK99]. 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.

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 2. 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 4.

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).

3.2. Data Fitting

In this section, we focus on the fitting of analytic BRDF functions to our acquired data.

Drying Paints and Rough Surfaces: We fit a combination of the Oren-Nayar diffuse model [ON94], denoted as ρ_d , and the Torrance-Sparrow specular model [TS67], denoted as ρ_s , to the BRDF measurement obtained from drying paints and drying wet surfaces. This combined BRDF model can be written as:

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

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.

Dust Accumulation: We have modified Blinn's reflectance model for dusty surfaces [Bli82] 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

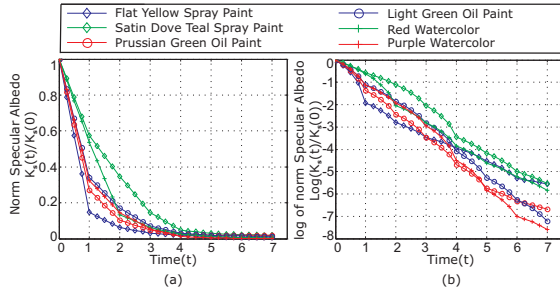


Figure 5: (a) The fall-off with time of K_s (normalized by the initial value $K_{s,wet}$) of various paint samples. (b) This plot is similar to the one in (a) except that the dry value is first subtracted and then the natural log is applied. Note that K_s in the case of paints attenuates exponentially with time.

roughness for the diffuse component, and τ is a dimensionless quantity called *optical thickness* which represents the attenuating power of the dust layer.

Fitting Algorithm: The Levenberg-Marquardt non-linear least-squares optimization algorithm [Mar63] 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 1.

3.3. Database

As shown in Figure 3, 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 1. 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.

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). 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.

4. 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.

4.1. Temporal Specular Trends

Materials with wet paint applied are highly specular due to strong surface reflection at the liquid-air interface. As the

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

Table 1: 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 %.

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 5(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 5(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 6(a). We plot $1/\sigma_s$ in linear scale in Figure 6(b) and see that it falls off exponentially with time. In Figure 6(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 6(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 estimated using the above model. Alternatively, it can be selected by a user when creating a new paint TVBRDF.

4.2. 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 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 7(a), (b), and (c). 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 7(d), (e), and (f) show the color shifts of some of our measured oil paints. Spray paints, however, cover surfaces with opaque colored

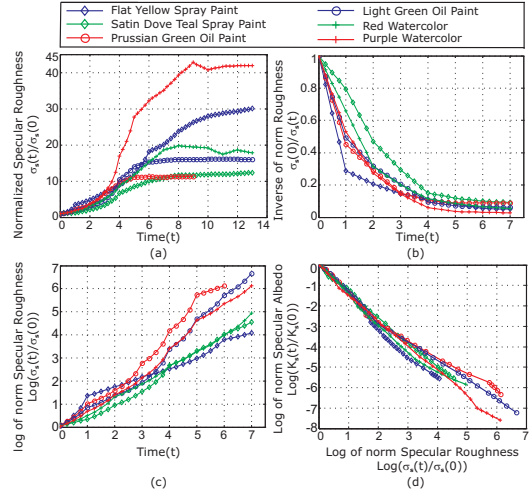


Figure 6: (a) σ_s (normalized by the initial value $\sigma_{s,wet}$) is plotted as a function of time for several of the paint samples and can be seen to vary exponentially. (b) $1/\sigma_s$ (normalized by the initial value $\sigma_{s,wet}$) plotted as a function of time. (c) $1/\sigma_s$ plotted with 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 .

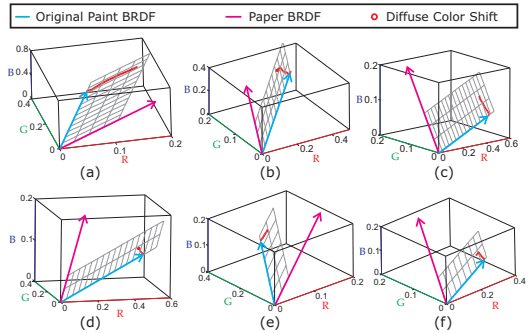


Figure 7: The diffuse color shifts associated with drying paints lie on dichromatic planes spanned by the diffuse color vector of the surface (shown in magenta) and the diffuse color vector of the paint (shown in cyan). The first row shows the dichromatic planes for watercolors: (a) Blue watercolor, (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.

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 7. 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 material and the underlying surface. Therefore, a dichromatic decomposition can be applied to separate the diffuse color into a




Figure 8: 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 white paper. The white sphere suspended in the corner represents a heat source.

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.

4.3. 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. Some other 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.

4.4. Rendering

The analytic TVBRDF model for paints also enables the transfer of the phenomena to novel materials. Figure 8 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 effect of green watercolor drying on white paper (bunny) and the effect of blue watercolor drying on red paper (bird). 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.




Figure 9: 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.

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.

5. Drying of Wet Surfaces

We apply a similar analysis to wet surfaces as we did for paints and develop their analytic TVBRDF model.

The specular highlights of most wet surfaces vanish very quickly, as seen in Figure 3(d), and can be ignored for subsequent time instances (grazing angle specularity may be missed in our measurements). 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 9. 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.

6. 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. 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.

6.1. Temporal Parameter Trends

The change in diffuse colors due to dust has been addressed by transparency T in Equation 2. Therefore, we focus on the behavior of the specular highlights subject to dust.

After normalizing the specular amplitude K_s by its initial value, we plot it in log scale versus the optical thickness τ . As Figure 10(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. 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 10(b), and thus can be treated as negligible.




Figure 10: (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 .

6.2. Analytic Time Varying BRDF Model for Dust

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

$$\rho(\tau) = (1 - T(\tau)) \cdot \rho_{dust} + T(\tau) \cdot \rho_d + e^{-\lambda\tau} \cdot \rho_s, \quad (8)$$

where $T(\tau)$ is the modified transparency term as defined in Equation 2 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 . 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 $\frac{10}{3}$ for household dust and 11 for joint compound dust.

6.3. 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 [HW95]. 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 11, a tea set scene is accumulating dust cast from a circular dust




Figure 11: 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).

source above. The teapot and teacup are rendered directly using our acquired data while the material of the table is synthesized with a low specular exponent. 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.

Further, due to occlusion, the surface exposure at all points 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.

7. Conclusions and Future Work

We have for the first time captured, modeled and rendered time-varying BRDFs. A major result of our work is a comprehensive database which will be made accessible upon publication. 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 physics simulations.

Acknowledgment

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

References

- [Bli82] BLINN J.: Light reflection functions for simulation of clouds and dusty surfaces. In *SIGGRAPH 82* (1982), pp. 21–29.
- [CAS*97] CURTIS C. J., ANDERSON S. E., SEIMS J. E., FLEISCHER K. W., SALESIN D. H.: Computer-generated watercolor. In *SIGGRAPH '97* (1997), pp. 421–430.
- [CF99] CHEN J. X., FU X.: Integrating physics-based computing and visualization: Modeling dust behavior. *Computing in Science and Engg. 1*, 1 (1999), 12–16.
- [DEJ*99] DORSEY J., EDELMAN A., JENSEN H., LEGAKIS J., PEDERSEN H.: Modeling and rendering of weathered stone. In *SIGGRAPH 99* (1999), pp. 225–234.
- [DH96] DORSEY J., HANRAHAN P.: Modeling and rendering of metallic patinas. In *SIGGRAPH 96* (1996), pp. 387–396.
- [DPH96] DORSEY J., PEDERSEN H. K., HANRAHAN P.: Flow and changes in appearance. In *SIGGRAPH 96* (1996), pp. 411–420.

- [DVGNK99] DANA K., VAN-GINNEKEN B., NAYAR S., KOENDERINK J.: Reflectance and Texture of Real World Surfaces. *ACM Transactions on Graphics (TOG)* 18, 1 (Jan 1999), 1–34.
- [EKB*05] ENRIQUE S., KOUDELKA M., BELHUMEUR P., DORSEY J., NAYAR S., RAMAMOORTHI R.: Time-varying textures: Definition, acquisition, and synthesis. *SIGGRAPH 2005 Sketch* (2005).
- [GTR*06] GU J., TU C., RAMAMOORTHI R., BELHUMEUR P., MATUSIK W., NAYAR S.: Time-varying surface appearance: Acquisition, modeling, and rendering. In *SIGGRAPH '06* (2006).
- [HM92] HAASE C. S., MEYER G. W.: Modeling pigmented materials for realistic image synthesis. *ACM Trans. Graph.* 11, 4 (1992), 305–335.
- [HW95] HSU S., WONG T.: Simulating dust accumulation. *IEEE Comput. Graph. Appl.* 15, 1 (1995), 18–22.
- [JdVL95] JR. H. B. M., DA VITORIA LOBO N.: Determining wet surfaces from dry. In *ICCV '95: Proceedings of the Fifth International Conference on Computer Vision* (1995), p. 963.
- [JLD99] JENSEN H. W., LEGAKIS J., DORSEY J.: Rendering of wet materials. In *Rendering Techniques 99* (1999), pp. 273–282.
- [Kou04] KOUDELKA M. L.: *Capture, Analysis and Synthesis of Textured Surfaces With Variation in Illumination, Viewpoint, and Time*. PhD thesis, Yale University, 2004.
- [Lam60] LAMBERT J. H.: Photometria sive de mensura de gratibus lumi-nis. *colorum umbrae* (1760).
- [LD88] LEKNER J., DORF M.: Why some things are darker when wet. *Applied Optics* 27 (1988), 1278–1280.
- [LGR*05] LU J., GEORGHIADES A. S., RUSHMEIER H., DORSEY J., XU C.: Synthesis of material drying history: Phenomenon modeling, transferring and rendering. In *Eurographics Workshop on Natural Phenomena* (2005).
- [Mar63] MARQUARDT D.: An algorithm for least-squares estimation of nonlinear parameters. *SIAM J. Appl. Math* 11 (1963), 431–441.
- [MPBM03] MATUSIK W., PFISTER H., BRAND M., MCMILLAN L.: A data-driven reflectance model. *ACM Transactions on Graphics* 22, 3 (July 2003), 759–769.
- [MWL*99] MARSCHNER S. R., WESTIN S. H., LAFORTUNE E. P. F., TORRANCE K. E., GREENBERG D. P.: Image-based brdf measurement including human skin. In *In Proceedings of 10th Eurographics Workshop on Rendering* (1999), pp. 139–152.
- [MWLT00] MARSCHNER S. R., WESTIN S. H., LAFORTUNE E. P. F., TORRANCE K. E.: Image-based bidirectional reflectance distribution function measurement. *Applied Optics* 39 (2000), 2592–2600.
- [NKON90] NAKAMAE E., KANEDA K., OKAMOTO T., NISHITA T.: A lighting model aiming at drive simulators. In *SIGGRAPH '90* (1990), pp. 395–404.
- [ON94] OREN M., NAYAR S.: Generalization of Lambert's reflectance model. In *SIGGRAPH 94* (Jul 1994), pp. 239–246.
- [TS67] TORRANCE K., SPARROW E.: Theory for off-specular reflection from rough surfaces. *Journal of the Optical Society of America* 57 (Sep 1967), 1105–1114.
- [War92] WARD G.: Measuring and modeling anisotropic reflection. In *SIGGRAPH '92* (1992), pp. 265–272.
- [WB90] W.BECKET, BADLER N.: Imperfection for realistic image synthesis. *The Journal of Visualization and Computer Animation* 1 (1990), 26–32.
- [WTL*06] WANG J., TONG X., LIN S., PAN M., BAO H., GUO B., SHUM H.: Appearance manifolds for modeling time-variant appearance of materials. In *SIGGRAPH '06* (2006).