

Time-Varying Surface Appearance: Acquisition, Modeling and Rendering

Jinwei Gu^{1*} Chien-I Tu^{1,2} Ravi Ramamoorthi¹ Peter Belhumeur¹ Wojciech Matusik² Shree Nayar¹
¹ Columbia University ² Mitsubishi Electric Research Laboratories (MERL)

Abstract

For computer graphics rendering, we generally assume that the appearance of surfaces remains static over time. Yet, there are a number of natural processes that cause surface appearance to vary dramatically, such as burning of wood, wetting and drying of rock and fabric, decay of fruit skins, or corrosion and rusting of steel and copper. In this paper, we take a significant step towards measuring, modeling and rendering time-varying surface appearance. We describe the acquisition of the first time-varying database of 26 samples, encompassing a variety of natural processes including burning, drying, decay and corrosion. Our main technical contribution is a Space-Time Appearance Factorization (STAF). This model factors space and time-varying effects. We derive an overall temporal appearance variation characteristic of the specific process, as well as space-dependent textures, rates and offsets, that control the different rates at which different spatial locations evolve, causing spatial patterns on the surface over time. We show that the model represents a variety of phenomena accurately. Moreover, it enables a number of novel rendering applications, such as transfer of the time-varying effect to a new static surface, control to accelerate time evolution in certain areas, extrapolation beyond the acquired sequence, and texture synthesis of time-varying appearance.

1 Introduction

Many interesting appearance properties of real-world surfaces are directly related to their evolution with time. Consider the charring of wood owing to heat or burning; the wetting and drying of stone, granite or fabric due to rain or spillage of water; the decay and ripening of fruit skins like apples or bananas; and the corrosion and rusting of steel or the formation of oxides on copper. Each of these natural processes forms a spatial pattern over time, often coupled with a change in reflectance, that gives rise to dramatic effects.

These processes has been studied in biology, physics and mathematics [Meinhardt 1992; Cross and Hohenberg 1993]. In computer graphics, Dorsey and collaborators have developed a number of specific models for flows, patina formation and weathering [Dorsey and Hanrahan 1996; Dorsey et al. 1996; Dorsey et al. 1999]. However, the full generality of pattern formation remains beyond the reach of any particular mathematical model or physical simulation.

In this paper, we avoid the difficulties with mathematical modelling by developing a data-driven approach, conceptually similar to recent work on data-driven static reflectance [Matusik et al. 2003] or texture [Dana et al. 1999]. We present a complete pipeline from acquisition of the first dense datasets of Time and Spatially-Varying appearance of flat samples (the TSV-BRDF) to the first data-driven models and novel renderings of time-varying appearance:

Database of Time-Varying Surface Appearance: A major contribution of our work is a database of time-varying appearance measurements, that we will release online upon publication. We have captured 26 samples, listed in figure 2, some examples of which are shown in figures 1 and 4. Because of the complexity in preparing

the samples, and developing a suitable measurement system, this database is likely to be a very relevant resource for future efforts.

Sample preparation requires careful control and significant effort—for example, we must apply a heat gun for some of the burning examples, and use special solutions to assist rusting and copper patination. Capturing the full TSV-BRDF also necessitates special measurement systems. We use a multi-light-source multi-camera dome, shown in figure 3, to simultaneously acquire time-lapse images from a variety of lighting and view directions. We then fit spatially-varying BRDF models at each time instance, capturing both spatial and temporal variation in a variety of real-world processes. Since we acquire the full TSV-BRDF, we can capture spatial patterns as well as changes in the BRDF, such as the sharp reduction in specularities over time when a surface dries.

Space-Time Appearance Factorization (STAF): Time-varying appearance is an intricate combination of many factors, such as the static surface texture, temporal variation, and spatial patterns over time. The acquired data can be used directly for rendering, but is difficult to understand, or modify for production applications (such as making wood dry faster in a wet footprint). Linear data-reduction techniques like singular-value decomposition do not easily capture the complex structures in time-varying appearance data.

We introduce a simple Space-Time Appearance Factorization (STAF), that is general and data-driven. It separates temporally varying effects from spatial variation, estimating a “temporal characteristic curve” in appearance that depends only on the physical process, as well as static spatial textures that remain constant over time. In addition, we estimate a rate and an offset at every point, that control the speed of time evolution. Spatial patterns arise because different points evolve at different rates. STAF is non-linear, with the temporal characteristic curve scaled and shifted by spatial rate and offset parameters. Our model is intuitive, accurate for the variety of time-varying phenomena in our database, and allows a user to separately modify space and time-varying effects.

Rendering Time-Varying Effects: One of the chief benefits of our data-driven STAF model is the ease with which we can generalize beyond the acquired data, to render a variety of time-varying effects. For example, we can *transfer* a time-varying process like rusting to a new static surface, like a steel plate. We can *control* the rate of time variation, such as having wet puddles or footprints on an otherwise dry wooden floor, with drying happening more at the boundaries. We can *extrapolate* to some extent beyond the acquired data. Moreover, separation of the spatial and temporal aspects allows one to use standard 2D example-based *texture synthesis*.

2 Previous Work

Physical simulation has been applied to specific weathering and corrosion effects [Dorsey et al. 1996; Dorsey and Hanrahan 1996; Dorsey et al. 1999; Merillou et al. 2001; Chen et al. 2005]. Patina formation has also been modeled based on surface accessibility [Miller 1994], while dust accumulation has been simulated based on surface geometry [Hsu and Wong 1995]. Jensen et al. [1999] render wet surfaces by combining a reflection model for surface water with subsurface scattering. Our data-driven approach generalizes and complements physical simulation of specific phenomena, much as static data-driven reflectance models complement and extend specific physically-based analytic BRDFs.

Some recent work has made a first attempt at measuring appearance changes. Koudelka [2004] considers time-varying textures imaged with fixed lighting and a single view, and extends static texture synthesis to time-varying texture synthesis. We generalize this method with images from multiple light sources and viewpoints, allowing us to fit a true TSV-BRDF model, enabling computer graphics rendering with any lighting and view. More importantly, we develop an intuitive data-driven STAF model to separate spatial and

*e-mail: jwgu@cs.columbia.edu

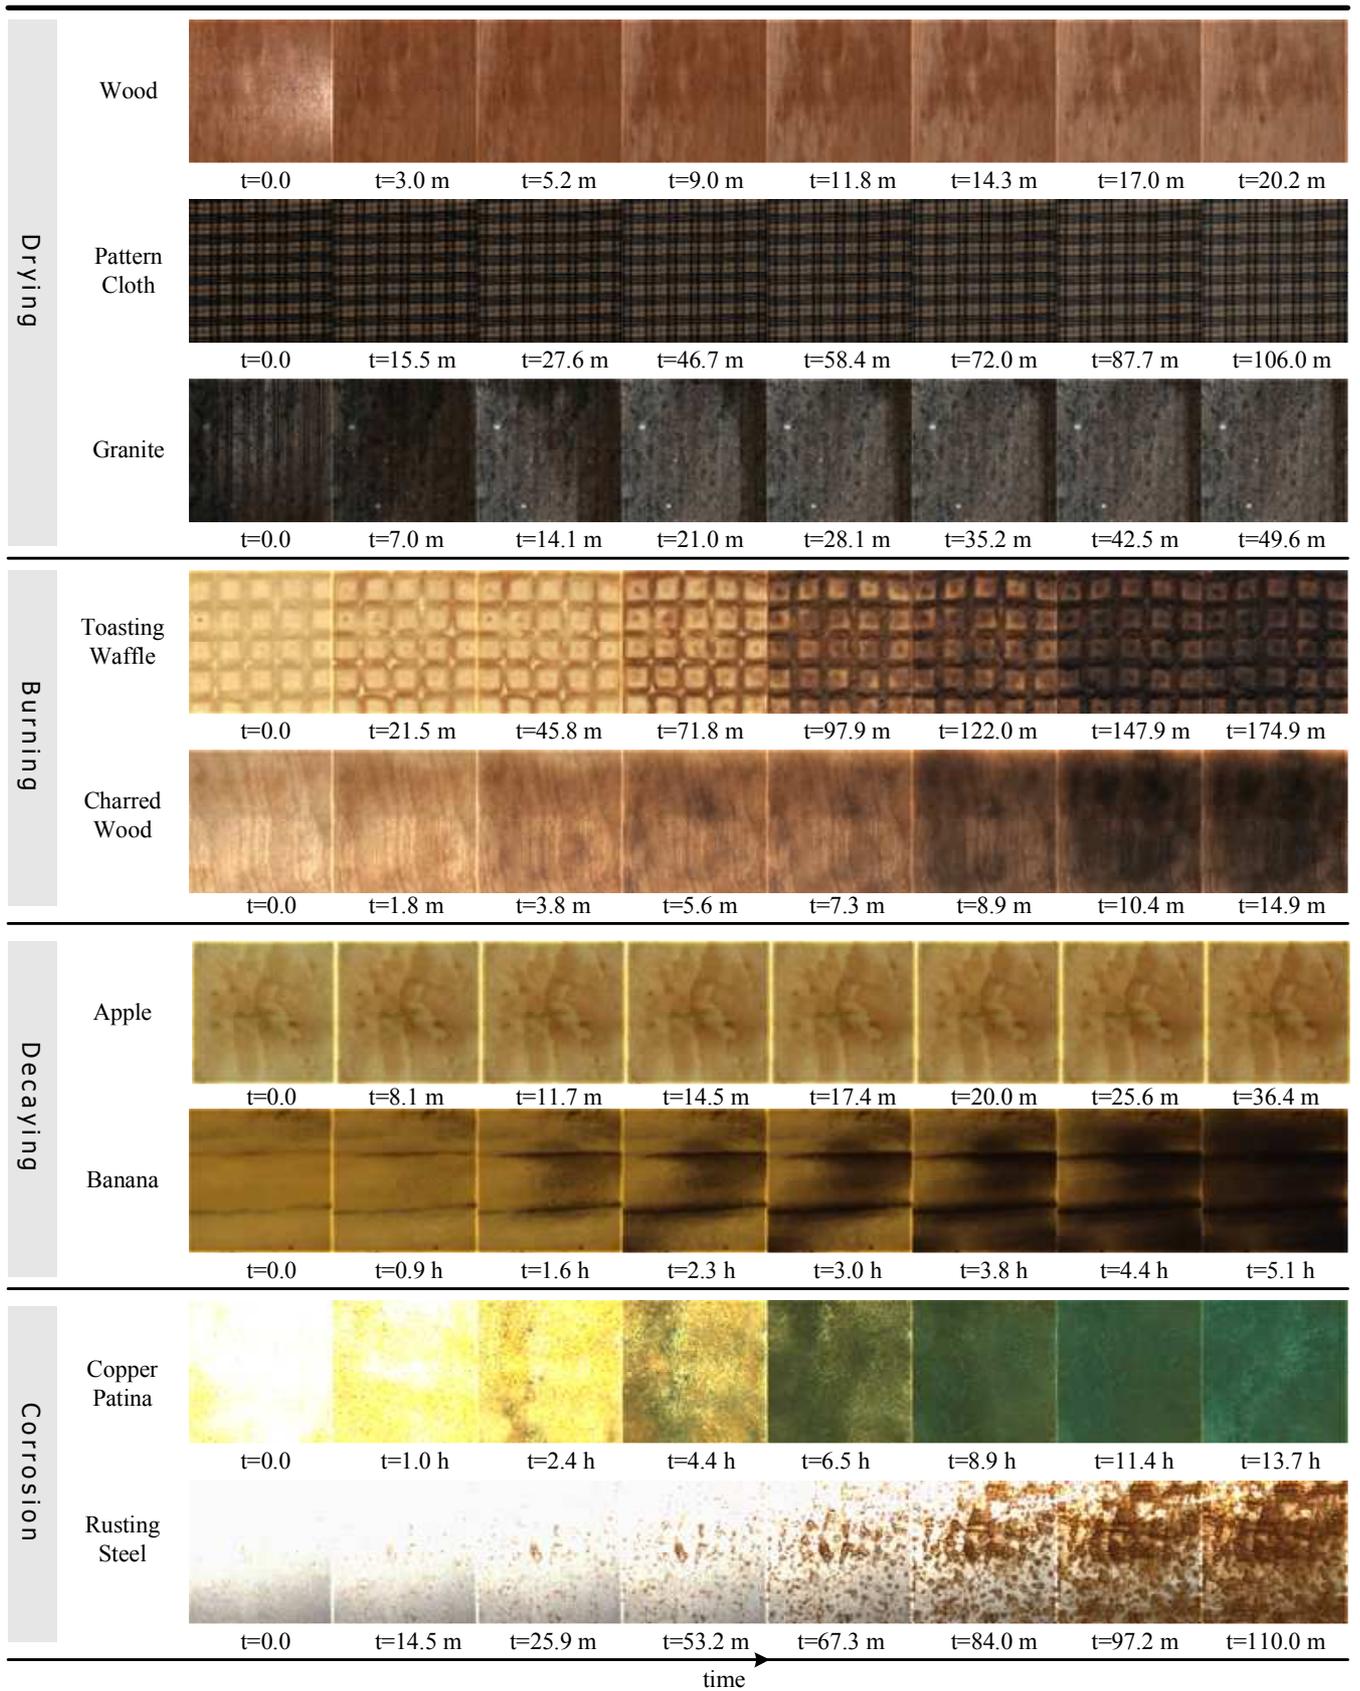


Figure 1: Some examples of the 26 samples in our database, shown here with variation across time (in minutes m or hours h) for a single light source and view. We acquire images from 1280 light and view directions at each time step—some of these images are shown for one of the samples in figure 4.

| Appearance | Time-Varying Appearance |
|--------------------------|-------------------------|
| TF (2D Texture Function) | TTF (3D) |
| BRDF (4D) | TBRDF (5D) |
| SV-BRDF (6D) | TSV-BRDF (7D) |
| BTF (6D) | TBTF (7D) |

Table 1: Extension of common appearance concepts to time-varying appearance. We also indicate the dimensionality of the function for each category. In this paper, we focus on TSV-BRDFs.

temporal effects, allowing a variety of rendering algorithms including transfer, control, extrapolation and synthesis.

For the specific case of drying on stone, [Lu et al. 2005] measure the change in diffuse appearance, and propose a sigmoid model with two spatial parameters. Similar equations can be deduced from the drying literature [Jankowsky and Santos 2004]. We generalize this work significantly, by acquiring a database of a variety of time-varying phenomena, including specular effects. Our STAF model is general and data-driven, capturing many types of time-varying processes, with intuitive rate and offset parameters at each spatial location. For specific drying scenarios, we essentially reproduce the results of [Lu et al. 2005], with our temporal characteristic curves being close to sigmoidal in those cases.

The STAF model in this paper relates to work in the statistical and speech recognition literature known as dynamic time-warping [Sakoe and Chiba 1978]. Their goal is to align time-varying curves for different subjects, in many applications like speech signals and human growth curves. Their data vary not only in amplitude, but also with respect to the time axis—different subjects experience events sooner or later. Classical linear methods like PCA cannot handle this second type of variability well [Wang and Gasser 1999]. Recently, [Kneip and Engel 1995] proposed the “shape-invariant” model, with the overall time variation known as the “structural average curve” (shape and structure are used rather differently from their traditional meaning in graphics).

In our application, we seek to align time-varying appearance curves (representing BRDF parameters like diffuse color and specular intensity) for different pixels, and we must relate this alignment to intuitive parameters like the rates and offsets at different spatial locations, as well as the static initial and final appearance. Moreover, as discussed in section 5, we develop methods to estimate the time variation of the process across the full range seen by any pixel, allowing robust extrapolation beyond the observed sequence.

3 Time-Varying Appearance

We first formalize the notion of time-varying appearance. One can imagine extending common appearance concepts, such as the BRDF or texture to include an additional time dimension, as shown in table 1. In this paper, we extend spatially-varying BRDFs (SV-BRDFs) to time and space-varying BRDFs (TSV-BRDFs). A general TSV-BRDF is a function of 7 dimensions—2 each for spatial location, incident angle and outgoing direction, and 1 for time variation. For surfaces that are rough, or have relief at a macroscopic scale, the term Bi-Directional Texture Function or BTF [Dana et al. 1999], and its time-varying extension TBTF is more appropriate, although it has the same dimensionality. While a small number of the examples in our database do have some surface relief (and may therefore not be as well modeled by the approach presented here), we focus in this paper primarily on flat surfaces or TSV-BRDFs.

4 Acquisition and Database

The first step in understanding time-varying surface appearance is to acquire datasets representing it—some examples are shown in figure 1. Figure 2 lists all of the 26 samples we have acquired and processed¹. These samples cover 5 categories—burning and charring (wood, waffles), drying of smooth surfaces (wood, fabric), drying of rough surfaces (rock, granite), corrosion and rusting (steel, copper), and decay and ripening (apples, banana).

| Type | Sample | Time Frames | Average Time Interval |
|---------------------------------|-----------------------|-------------|-----------------------|
| Burning | Charred Wood 1 | 11 | 2.1 m |
| | Charred Wood 2 | 31 | 9.9 m |
| | Waffle Toasting | 30 | 6.3 m |
| | Bread Toasting | 30 | 5.9 m |
| Drying (Smooth Surfaces) | Light Wood 1 | 14 | 3.1 m |
| | Light Wood 2 | 34 | 2.3 m |
| | Orange Cloth | 33 | 4.9 m |
| | Cotton Cloth | 30 | 11.3 m |
| | Pattern Cloth | 32 | 4.8 m |
| | White Felt | 28 | 4.4 m |
| Drying (Rough Surfaces) | Dark Wood | 32 | 3.8 m |
| | Paper Towel | 32 | 7.0 m |
| | Brick | 32 | 22.1 m |
| | Rock | 11 | 2.0 m |
| | Cardboard | 29 | 7.0 m |
| | Granite | 27 | 2.6 m |
| Corrosion | Tree Bark | 11 | 3.4 m |
| | Rusting Steel 1 | 22 | 7.3 m |
| | Rusting Steel 2 | 35 | 10.8 m |
| | Cast Iron Rusting | 30 | 13.9 m |
| Decaying | Copper with Patina | 34 | 31.6 m |
| | Apple with Core | 33 | 9.6 m |
| | Apple Slice | 13 | 3.0 m |
| | Banana | 33 | 11.3 m |
| | Potato | 31 | 8.3 m |
| | Leaf under Humid Heat | 30 | 12.6 m |

Figure 2: The 26 samples in our database, grouped into categories. For each sample, we list the number of time frames acquired, and average time interval between frames (in minutes m).



Figure 3: A photograph of the multi light-source multi-camera dome used for acquisition of our database of time-varying measurements.

4.1 Acquisition

Acquisition of time-varying appearance is challenging. While some natural processes like drying occur over fairly short time scales (a few minutes), many others occur over a considerable duration under normal circumstances (several hours to days for decay of fruit skins, or a few months for corrosion of metals). In the case of burning and charring, we have used a heat gun to carefully control the process. At each time interval, we uniformly heat the sample for a fixed duration of time (typically 30 seconds). For metal corrosion, we have decided to speed up the process using specially prepared solutions [Hughes and Rowe 1991]. We spray a chemical solution before each measurement and wait a few hours. Decay of organic samples takes several hours, and is fairly difficult to speed up—we have decided to measure these processes without alteration.

A second difficulty is designing and building a measurement system that meets the following resolution requirements: 1) Dynamic range—many of the processes (e.g, drying or rusting) involve significant changes in specularity. 2) Light and view direction resolution—the sampling of the light and view directions should be sufficiently high to capture specular materials. 3) Temporal resolution—a complete acquisition at each time step, involving images with multiple lights, views and exposure settings needs to be conducted in a few seconds to avoid the sample changing sig-

¹This entire database, and our STAF model fits will be made available online. To request a copy, send e-mail to staf@cs.columbia.edu.

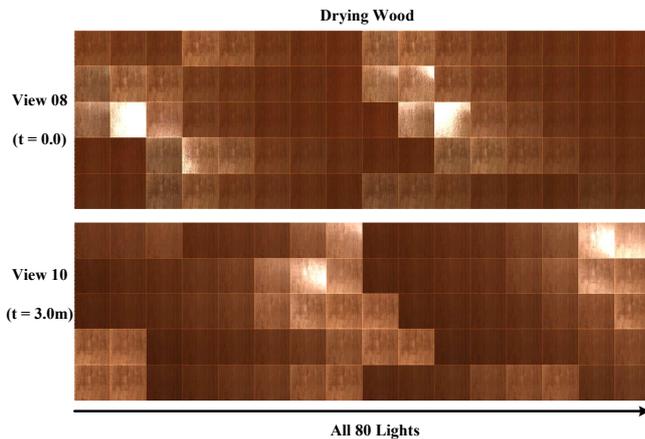


Figure 4: Acquired images of wood drying. We show two separate views/time instances, and all of the useful lighting directions for those.

nificantly over this time. This rules out gantry-based systems, that typically take a few seconds to acquire even a single image.

We have decided to use a multi-light-source multi-camera dome, shown in figure 3. The dome skeleton is based on an icosahedron. We use 16 Basler cameras (resolution 1300×1030 pixels) placed on the icosahedron vertices and 150 white LED light sources spaced evenly on the edges (approximately 80 of these lights lie in the visible hemisphere with respect to the flat sample, and therefore give useful images). This design is similar to the light stage [Debevec et al. 2002], but includes multiple cameras as well. The cameras and light sources are synchronized using a custom-built controller.

The cameras are geometrically calibrated by moving a small LED diode in the working volume and detecting its 2D location in all cameras. A bundle adjustment is performed to obtain the precise geometric location and projection matrices for all cameras. Since we know the dome’s design specifications, this allows us to register all light and camera positions to a common coordinate frame. We also perform a photometric calibration of the system, by capturing images of a perfectly white diffuse standard (spectralon) from all camera viewpoints under all light combinations. To obtain normalized BRDF values for each surface point, we divide by the corresponding observation of the white diffuse standard.

For acquisition, we place a prepared sample in the center of the dome. At each time step, we capture a high dynamic range data set—we take images at two different exposures (typically 2 and 82 msec) for each light-camera pair. This results in 4,800 photographs captured in 22 seconds. It takes about 90 seconds to save the data to the hard disk (therefore, the minimum time between two consecutive measurements is about 2 minutes). We typically capture appearance data sets at 30 time frames.

Once a complete time-varying appearance data set is captured, we resample the data on a uniform grid (typically 400×400 pixels) for each light and view direction. Some of our data, showing time variation for a single light source and view has already been seen in figure 1. Figure 4 shows all of the 80 useful images (lighting directions in the visible hemisphere) for two time instances/viewpoints.

4.2 Time and Spatially-Varying Parametric Reflectance

Initially we attempted to take a straightforward non-parametric approach to represent the BRDF at every point directly by the acquired raw data. For rendering (i.e. to create images under novel view and lighting), we used the algorithm in [Vlasic et al. 2003] and did barycentric interpolation twice, once over view and then over lighting. A similar algorithm is used in [Vasilescu and Terzopoulos 2004]. However, as shown in figure 5, since the light-view sampling of our samples is not dense enough, direct interpolation produces artifacts. In figure 5, we have “texture-mapped”² the TSV-BRDF

²When we refer to “texture mapping” throughout this paper, we mean mapping the complete TSV-BRDF, i.e. all 5 BRDF parameters, including diffuse RGB color and specular K_s and σ , and including time variation. These BRDF parameters at each point in space and time can then be used with any lighting model and rendering computation.

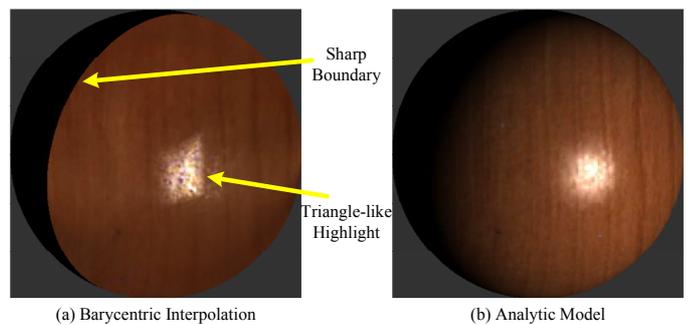


Figure 5: Comparison of (a) barycentric interpolation and (b) parametric spatially-varying reflectance fits, texture-mapped onto a sphere. The parametric reflectance model is quite accurate, preserving the fine details of the wood grain, while eliminating artifacts in the highlights and boundaries.

onto a 3D sphere to better make these comparisons.

Fortunately, we have enough measurements to effectively fit parametric reflectance models, including specular lobes, to each spatial location. We use a simple combination of diffuse Lambertian and simplified Torrance-Sparrow reflectance, with the BRDF given by

$$\rho(x, y, \vec{\omega}_i, \vec{\omega}_o, t) = K_d(x, y, t) + \frac{K_s(x, y, t)}{4(\vec{\omega}_i \cdot \vec{n})(\vec{\omega}_o \cdot \vec{n})} \exp \left[- \left(\frac{\cos^{-1}(\vec{\omega}_h \cdot \vec{n})}{\sigma(x, y, t)} \right)^2 \right], \quad (1)$$

where $\vec{\omega}_i$ and $\vec{\omega}_o$ are incident and outgoing directions, \vec{n} is the surface normal and $\vec{\omega}_h$ is the half-angle vector. The BRDF parameters are the diffuse intensity K_d , the specular intensity K_s and the surface roughness σ . Since K_d is an RGB color, we have a total of 5 parameters for each spatial location (x, y) and time t .

Note that the BRDF model used to fit the raw data is independent of the STAF model in the remaining sections. Other kinds of parametric BRDF models (e.g. Lafortune model) could also be used.

The diffuse and specular parameters are estimated separately in two steps, since for some materials there are only a few samples in the specular lobe. To fit the diffuse color K_d , we consider a frontal view, which gives the highest-resolution image. At each spatial location, we optimize over only those light source directions where specular highlights are not present (conservatively, we require the light source and the reflected view direction to be separated by at least 30° which works well for most of the samples in the database). We consider each time instance separately for the fits.

We fit the specular intensity K_s and roughness σ separately for each spatial location. To do so, we consider all light source directions and views. Since σ is the only non-linear parameter, we have found it most robust to do a linear exhaustive search to determine it. For a given σ , we solve a linear system for K_d and K_s , choosing the σ (and K_s) that has minimum error. Although we do estimate the diffuse K_d in this process again, we prefer to use the K_d described earlier, which is determined from the highest-resolution frontal view, and with specularly completely absent. To make the two estimates of K_d consistent, we scale the earlier estimate of K_d by the average value of the latter estimate of K_d over all spatial locations. As seen in figures 5 and 6, we capture the important qualitative aspects of the specularly, without artifacts. Quantitative analysis is difficult, since some spatial locations have only a sparse set of BRDF samples in the specular lobe.

4.3 Summary and Results

From now on, we will use the notation $p(x, y, t)$ for the parametric fits to the TSV-BRDF. p can be thought of as a vector of 5 space and time-varying parameters, the diffuse RGB color $K_d(x, y, t)$ and the specular $K_s(x, y, t)$ and $\sigma(x, y, t)$. The angular dependence is implicit in the form of the specular term controlled by K_s and σ . Note that although the BRDF representation is parametric, the estimated parameters $p(x, y, t)$ capture the space and time-variation of surface appearance in a non-parametric way (i.e., directly from the acquired raw data).

Even without the analysis and modeling in the rest of this paper, our database of TSV-BRDFs can be texture-mapped onto arbitrary 3D objects and used directly for rendering with general lighting

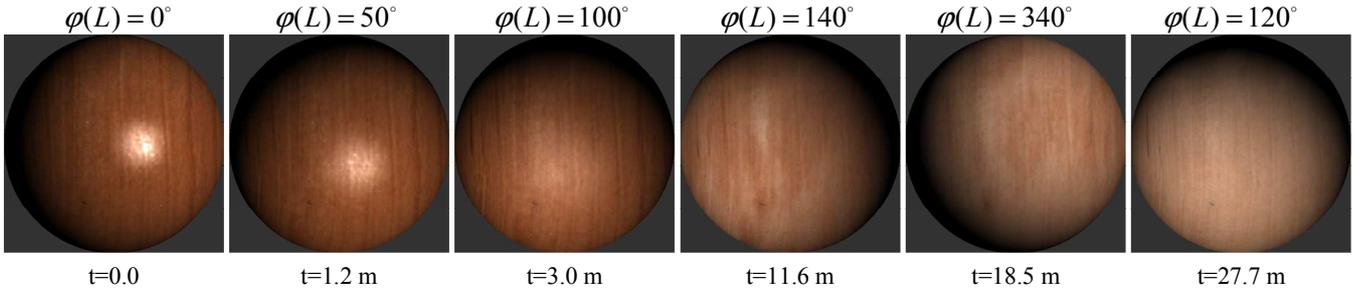


Figure 6: Drying wood TSV-BRDF, texture-mapped onto a sphere. This example demonstrates the power of our database, which enables us to render with simultaneous changes in lighting and evolution with time. Note the diffuse spatial drying patterns, and the early dimming and diffusing of specularities. The elevation angle of the light with respect to the center is fixed at $\theta(L) = 30^\circ$, while the azimuthal lighting angle varies as the sample dries.

direction, viewing angle, and time variation. Indeed, our use of standard parametric models allows time-varying effects to be easily incorporated in almost any interactive or off-line rendering system. As one example, figure 6 shows drying wood texture-mapped onto a sphere. We show a sequence of frames, where we simultaneously change the lighting, and evolve the sample over time. Note the spatial drying patterns, as well as BRDF changes, wherein the initial sharp specularities quickly diffuses and dims over time.

5 Modeling and Analysis of Time Variation

While our TSV-BRDF database can often be used directly, there are many rendering applications where the user desires more control. For example, he may want to control the spatial drying patterns on a wooden floor to dry slower near recent wet footprints. Or she may want to remove the spatial drying patterns altogether allowing the surface to dry uniformly. The user might also want to change the underlying spatial texture, to create a different appearance for the wood grain. These effects are difficult to create, because space and time variation are deeply coupled in the TSV-BRDF, while we seek to separately modify or edit intuitive spatial or temporal functions (like overall spatial texture or rate of variation).

We present a Space-Time Appearance Factorization (STAF) that separates effects because of space and time-variation, showing how they interact. In this section, we introduce the STAF model, and show how to estimate it from the TSV-BRDF. We present results indicating its accuracy for the large variety of time-varying phenomena in our database. In section 6, we will show the power and flexibility of the STAF model in creating novel rendering effects.

5.1 Space-Time Appearance Factorization (STAF)

Our approach is based on the idea that most physical processes have an overall temporal behavior associated with them. For example, drying wood may get lighter over time. For a given parameter of the BRDF, like the diffuse red channel, this time variation can be expressed by a curve $p(x, y, t)$ for each spatial location. However, different points dry at different rates. Moreover, in situations like a puddle, some parts start out wetter than others. Intuitively, we seek to align the time variation for different spatial locations, by forming a single “temporal characteristic curve” $\phi(t)$ according to spatially-varying parameters for “rate” $R(x, y)$ and “offset” $O(x, y)$,

$$\begin{aligned} p(x, y, t) &= A(x, y)\phi(t') + D(x, y) \\ t' &= R(x, y)t - O(x, y). \end{aligned} \quad (2)$$

In this equation, we consider each of the 5 parameters of the TSV-BRDF separately. For example, for the diffuse component, one can think of all quantities as being RGB colors. The model is data-driven, since the factors or terms A , D , R , O and ϕ are estimated directly from the acquired data, and represented in a purely data-driven way. We now describe the meanings of the various terms.

$\phi(t')$ – Temporal Characteristic Curve: The overall time variation characteristic of the physical process is captured by the curve $\phi(t')$. The form of ϕ will vary with the specific phenomenon. It can be exponential for some decays, sigmoidal for drying and burning, a more complex polynomial form for rusting, or any other type of curve. Since our representation is fully data-driven, we can handle a variety of effects. ϕ is a function of t' , that we call the *effective time*, as described below.

$R(x, y)$ and $O(x, y)$ – Spatial Rate and Offset : Different spatial locations evolve differently. We capture these effects with spatially varying rate $R(x, y)$ and offset $O(x, y)$ parameters. If R is large, the rate of change will be rapid. If O is positive, the point will start from an earlier state. The *effective time* t' for a given point is given by $t' = R(x, y)t - O(x, y)$, where we refer to t as the *global time*.

$A(x, y)$ and $D(x, y)$ – Static SV-BRDFs : $A(x, y)$ and $D(x, y)$ are static over time. The diffuse components correspond to standard spatial textures like wood-grain that remain fixed throughout the time variation. Consider the special case when $R(x, y) = 1$ and $O(x, y) = 0$ so all points evolve in the same way. Equation 2 simply becomes $A(x, y)\phi(t) + D(x, y)$. In this case, we simply interpolate from one texture (or more generally, SV-BRDF) to another. The initial and final appearance are simply $A\phi(0) + D$ and $A\phi(1) + D$.

5.2 Discussion

Separating Spatial and Temporal Variation: The STAF model in equation 2 has factored spatial and temporal variation in a compact representation. We now have quantities (A, D, R, O) , that depend only on spatial location (x, y) , and a temporal characteristic curve $\phi(t)$ that controls time variation. Unlike linear decompositions, the STAF model is *non-linear*, because $\phi(t)$ is stretched and offset by the spatial rate and offset $R(x, y)$ and $O(x, y)$. A similar separation of spatial and temporal effects could not be accurately achieved by linear data reduction methods like PCA, nor would the terms in a linear model correspond to physically intuitive and editable factors.

Extrapolation: Another interesting aspect of the model is its power to extrapolate beyond the acquired sequence. Let us normalize the global time t in the range of $[0 \dots 1]$. Now, consider the effective time $t' = R(x, y)t - O(x, y)$, which lies in the range $J(x, y) = [-O(x, y), R(x, y) - O(x, y)]$. If either R and/or O is large, this range can extend considerably beyond the global $[0 \dots 1]$ time. The valid domain of effective times for the full curve $\phi(t')$ is now

$$J = \bigcup_{(x, y)} J(x, y) = \left[\min_{(x, y)} (-O(x, y)), \max_{(x, y)} (R(x, y) - O(x, y)) \right], \quad (3)$$

which considers the minimum and maximum effective time t' over all points (x, y) . By definition, the overall range of J is a superset of that for each point, enabling individual pixels to be backed up or extended beyond the sequence captured, allowing time extrapolation. This is reasonable because early-starting points can provide information for other similar points which start later by some offset.

5.3 Estimating the STAF model

We use a simple iterative optimization to estimate the factors in equation 2. Each iteration consists of two steps. In the first step, we fix the spatial parameters A , D , R and O to update our estimate $\phi(t')$. If the other terms are fixed, we can solve directly for ϕ in equation 2. The second step of the iteration fixes $\phi(t')$ and solves for the spatial parameters A , D , R and O . This requires non-linear optimization, but can be carried out separately for each spatial location (x, y) . We have found that only 5 iterations are needed to obtain accurate estimates of all parameters. This algorithm is very easy to implement, requiring fewer than 50 lines of Matlab code, while being robust and effective for the entire variety of samples in our database. We describe the technical details below. Some readers may wish to skip to our results in section 5.4.

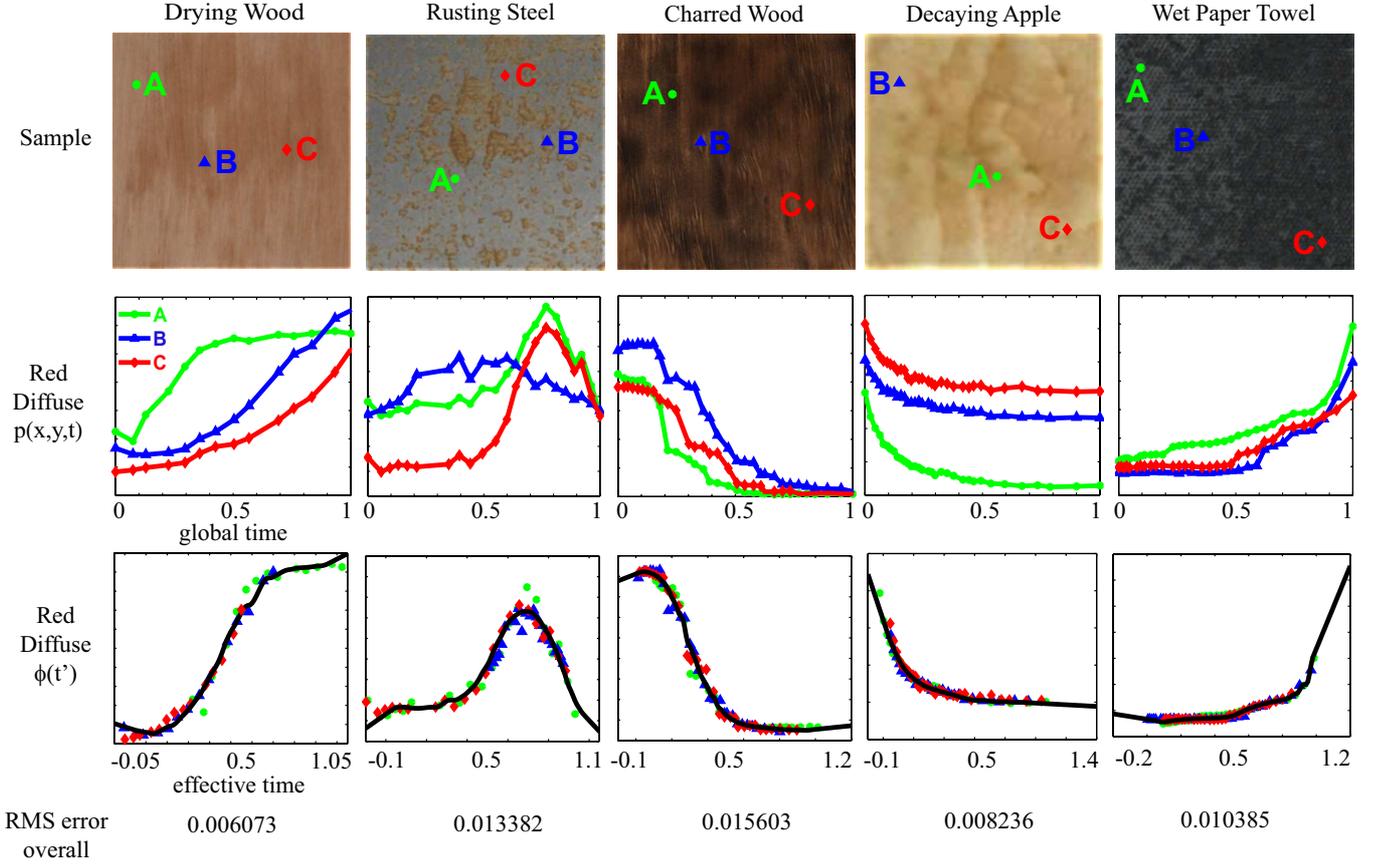


Figure 7: Estimating the factored representation. **Top:** A range of different phenomena, with 3 spatial locations marked on each sample. **Middle:** Time-varying curves $p(x, y, t)$ (for the red diffuse component) for spatial locations A, B and C. The curves are quite different for different points A, B and C. **Bottom:** We align these time-varying curves using our model. The data accurately matches the temporal characteristic curve $\phi(t')$ computed from all the points on the sample. The overall RMS image reconstruction error (across all temporal frames and spatial locations) is very low. This indicates the generality of our model.

Preprocessing: Our inputs are discrete fits of the parameters p at pixels i and times j , which we denote $p_i(t_j)$. The pixel i corresponds to spatial location (x_i, y_i) . It is simpler to work with continuous functions of time. For each pixel, we construct a continuous curve $p_i(t)$ using the kernel-based method [Gasser et al. 1985]. Splines or local polynomial fitting can also be used. We are now ready to begin our iterative optimization. To initialize, we set $A_i = 1$, $D_i = 0$ and $R_i = 1$, $O_i = 0$ for all pixels i .

Step 1 – Estimating $\phi(t')$: We first fix the spatial parameters A , D , R , O in order to estimate $\phi(t')$. For estimation, we re-arrange equation 2, writing $t = (t' + O_i)/R_i$ to derive for point i ,

$$\begin{aligned} A_i \phi(t') + D_i &= p_i \left(\frac{t' + O_i}{R_i} \right) \\ \phi(t') &= \frac{p_i((t' + O_i)/R_i) - D_i}{A_i}, \end{aligned} \quad (4)$$

for $t' \in J_i$, where J_i is the range of effective times. For robustness, and to consider the full effective time range, we add multiple points,

$$\phi(t') = \frac{\sum_{i:t' \in J_i} p_i((t' + O_i)/R_i) - \sum_{i:t' \in J_i} D_i}{\sum_{i:t' \in J_i} A_i}. \quad (5)$$

Step 2 – Estimating A , D , R , O : We now keep our value for the overall time curve $\phi(t')$ fixed, and estimate the spatial parameters. This is a separate optimization problem for each spatial location i ,

$$\min \sum_{j=1}^N [p_i(t_j) - A_i \phi(R_i t_j - O_i) - D_i]^2. \quad (6)$$

Note that this expression uses the discrete observations $p_i(t_j)$, finding spatial parameters that best match our input data. This is a nonlinear least-squares optimization problem, and we use the *lsqnonlin* function in Matlab, with Levenberg-Marquardt minimization.

Normalization: We are now almost ready to start a new iteration in the optimization, returning to step 1. One final detail is that the STAF model involves a product, and requires normalization of the factors for uniqueness. We use the following normalization,

$$\langle A_i \rangle = 1, \quad \langle D_i \rangle = 0 \quad \langle R_i \rangle = 1, \quad \langle O_i \rangle = 0, \quad (7)$$

where $\langle \cdot \rangle$ stands for the average over all spatial locations. This simply says that the overall spatial textures are normalized, and that the average rate is 1, while the average offset is 0.

Let us call the un-normalized results at the end of step 2, \hat{A}_i , \hat{D}_i , \hat{R}_i and \hat{O}_i . To normalize A_i and R_i , we simply divide by the average values for \hat{A}_i and \hat{R}_i . Then, we normalize D_i and O_i as follows,

$$D_i = \hat{D}_i - A_i \langle \hat{D}_j \rangle, \quad O_i = \hat{O}_i - R_i \langle \hat{O}_j \rangle. \quad (8)$$

We can now start the next iteration of the optimization, returning to step 1. In general, we find five iterations enough for convergence.

Efficiency and Robustness: For efficiency, instead of using all the points on the sample, we randomly select 400 points as input to the algorithm. Therefore, the iterative optimization itself takes only a few minutes. Once the final $\phi(t')$ is known, we use step 2 (equation 6) to directly estimate A, D, R, O for all points on the original sample. Since the sample resolution is large (400×400), and we must solve a nonlinear optimization for each pixel, the total time can be a few hours, but the process is completely automated.

One final issue is that we want to estimate $\phi(t')$ in the full range J , while the iterative optimization uses only part of the data. The kernel-based curve $p_i(t)$ cannot extrapolate well, and therefore neither can $\phi(t')$ from step 1 of the optimization. Therefore, in step 2 of the algorithm, instead of using $\phi(t')$ in equation 6, we fit $\phi(t')$ by a smooth polynomial $q(t')$ and use $q(t')$ to estimate A, D, R, O .

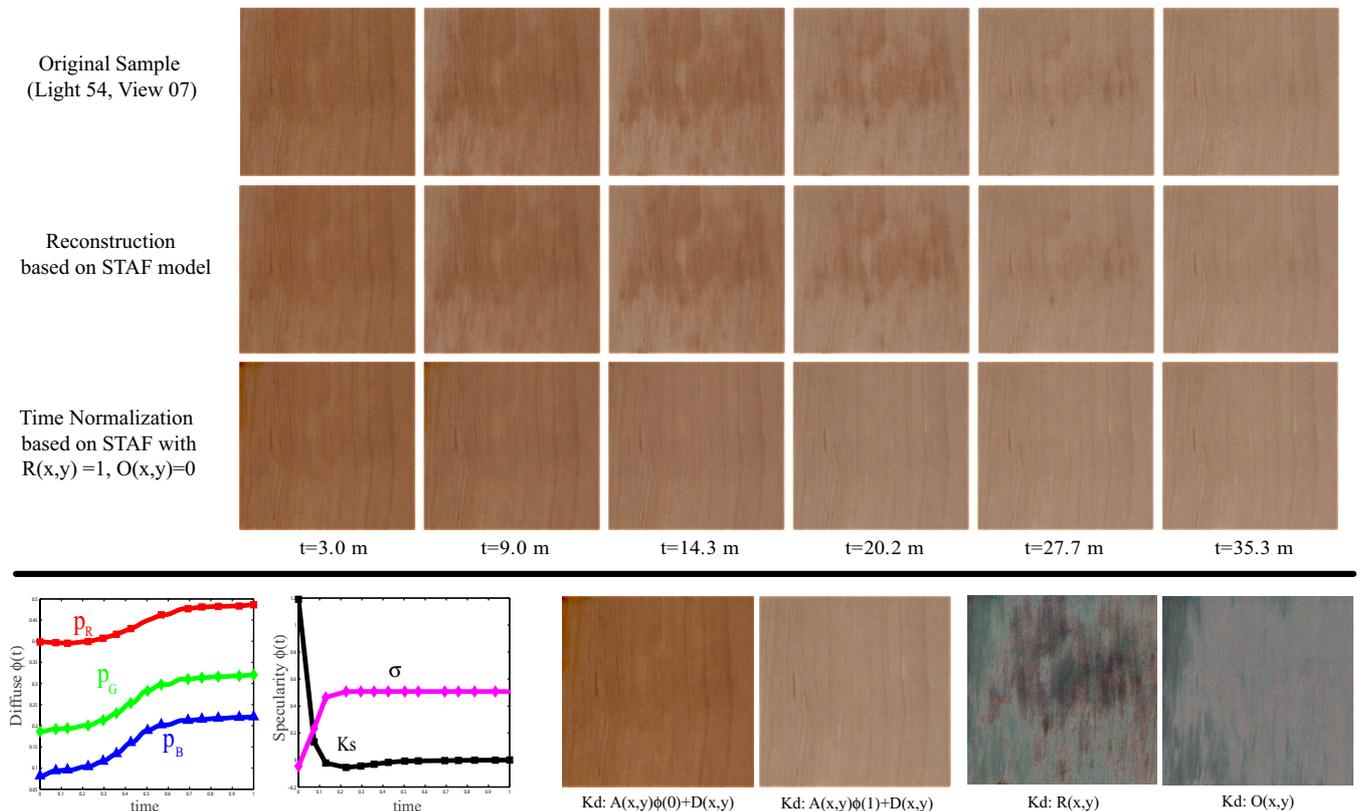


Figure 8: Factored representation for drying wood. The panel above compares STAF to the original sample for one light and view in the acquired dataset. Our model is seen to be quite accurate, and can also be used for time normalization, wherein we keep the overall appearance changes but eliminate the spatial drying patterns. The panel below shows the estimated $\phi(t')$ for both diffuse and specular parameters. We also show visualizations of the spatial diffuse “textures” A, D, R, O . In particular, we show the normalized initial frame $A\phi(0) + D$, and final frame $A\phi(1) + D$. We show $R(x, y)$ that controls the drying rate, and is responsible for the spatial patterns. Finally, the offset $O(x, y)$ is mostly uniform, but indicates a slightly faster start to drying in the left region.

5.4 Results

The top row of figure 7 shows five samples, with three spatial locations marked on each. The middle row shows curves for the red diffuse component over time (similar results are obtained for other parameters). As can be seen, the curves from different points on the sample are quite different. In the bottom row, we show alignment of these separate time-varying curves by estimating our factored representation (the $\phi(t')$ curve is plotted in black). Specifically, the x-axis is the *effective time*, t' while the y-axis is the normalized function value $(p(x, y, t) - D(x, y)) / A(x, y)$. The green/red/blue dots overlaid on the black curves show which portions of the black curves $\phi(t')$ correspond to each of the three original curves in the second row. Note that the $\phi(t')$ curves extrapolate beyond the data, having a larger range of effective times than $[0 \dots 1]$.

If the model in equation 2 is perfect, the curves from different spatial locations should now all be aligned, exactly fitting $\phi(t')$. Indeed, the time-aligned data in the bottom row of figure 7 matches very well to the canonical curve. The overall RMS image reconstruction errors are computed across all temporal frames and spatial locations. The range of the image intensity is generally in $[0, 1]$, except for samples with strong specular components, such as the steel for which the intensity of the specular pixels is in $[0, 30]$. Note that figure 7 shows a variety of phenomena, with a number of different data-driven forms and mathematical curve-types for the canonical $\phi(t')$.

The accuracy of our factored model is evaluated in figure 8. We accurately capture drying patterns over time. We also show the estimated diffuse “textures”. Instead of A and D , that are somewhat harder to interpret, we show the normalized initial frame, $A(x, y)\phi(0) + D(x, y)$ and final frame $A(x, y)\phi(1) + D(x, y)$. We also show $R(x, y)$ that controls the rate at which different points dry. It corresponds closely with the spatial patterns observed at later frames. Finally, we show the offset $O(x, y)$. It is mostly close to 0, since we wet our sample uniformly before starting acquisition.

However, it does indicate small non-uniformities and the slightly faster start to drying in the top left region. We also show the canonical diffuse and specular $\phi(t')$ curves. The specular K_s decreases exponentially, changing more rapidly than diffuse color.

One of the principal benefits of our factored representation is that it enables a variety of rendering applications, as discussed in the next section. Figure 8 indicates one way in which we can separate space and time-varying effects by “time-normalization”, making all points on the surface evolve at the same rate. For this purpose, we leave $A(x, y)$, $D(x, y)$ and $\phi(t)$ unchanged. However, we set $O(x, y) = 0$ to eliminate offsets and $R(x, y) = 1$ to eliminate differences in rates. The third row of figure 8 compares the time-normalized results with the original, showing that all pixels now change at the same rate, removing the spatial patterns. For rendering in the next section, we can now modify R and O , to create the spatial patterns and variations we desire, while still preserving the essence of the acquired time-varying phenomenon.

Figure 9 uses the rusting steel sample to compare renderings from the STAF model with the original TSV-BRDF. This example is particularly challenging, because the $\phi(t)$ red curve is not even monotonic (since rust forms, reddening the material, but also darkening its base color). Nevertheless, our factored data-driven model is accurate. We capture the dimming of the specular highlight, and the intricate spreading of the spatial rust patterns over time.

In terms of compression, the average size of the raw data (high dynamic range images) of one sample is about 30 GB. Fitting parametric BRDF models for each time step reduces the size to about 80 MB. The STAF model can further reduce the size of one sample to about 6 MB on average – we only need to store four texture images A, D, R, O and the curve $\phi(t)$ for each of the five parameters in the BRDF model. Using other image compression techniques (e.g., JPEG), we can reduce the size even further to about 1-2 MB without producing noticeable artifacts.

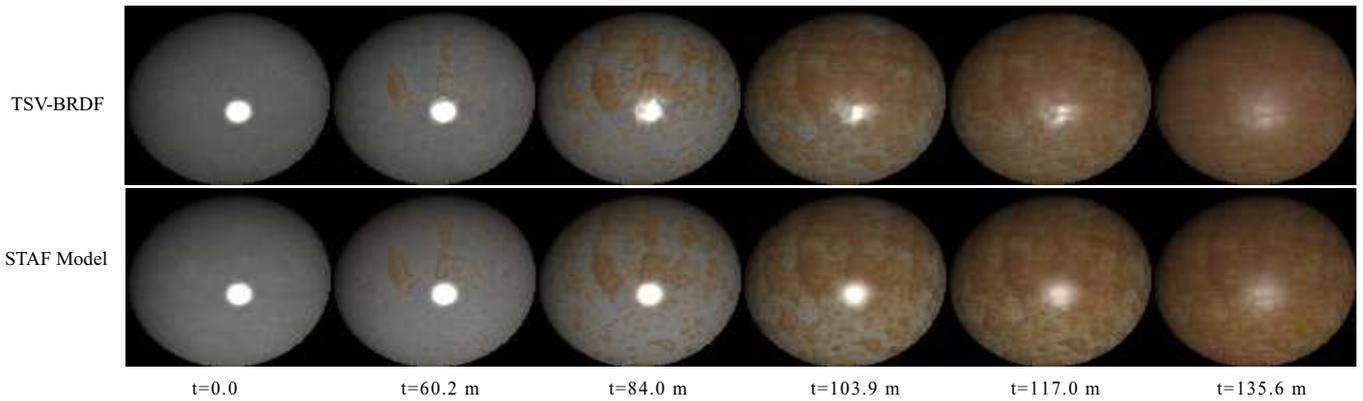


Figure 9: Comparison of original rust TSV-BRDF (texture mapped onto a sphere and rendered) with our STAF model for several time frames.

6 Rendering

It is possible to use the TSV-BRDF database directly for rendering, even without the STAF model. However, the types of appearance we could create are limited, since one cannot control or modify the TSV-BRDF. On the other hand, the STAF model completely factors space and time-varying effects, allowing either or both to be manipulated and edited separately.

Extrapolation: The temporal characteristic curve $\phi(t')$ extends beyond the actual global time range over which the data is acquired, allowing us to back up or extend the process beyond the acquired data for many pixels.

Control: By changing rate and offset parameters $R(x, y)$ and $O(x, y)$, we can control the rate at which different points on the surface change, while still preserving the characteristic features of the time-varying process. We could set R and O according to physical principles like the amount of light or humidity. In practice, we use simple procedural ideas—for example, a wet puddle dries faster near the edges, so we increase the rates in those regions.

Transfer: By changing the textures $A(x, y)$ and $D(x, y)$, to those obtained from a new static photograph, we can transfer the time-varying effect, such as burning or rusting, to a new object, while still preserving the essence of the data-driven appearance change.

Time-Varying Texture Synthesis: Our database is acquired on small flat samples. Of course, we can texture map these onto arbitrary 3D objects, but we also seek to use texture synthesis to create larger spatial patterns. With our factored form, we simply synthesize the spatial textures using standard 2D methods.

We now use these ideas to render a variety of examples, that showcase the full power of our method. The 3D renderings were done using the PBRT package [Pharr and Humphreys 2004].

Figure 10 shows a texture synthesis of the drying rock example, to create a much larger spatial pattern. To maintain temporal coherence from initial to final frame, we treat the spatial textures A and D together. We first synthesize $I_0 = A\phi(0) + D$ using image quilting [Efros and Freeman 2001], and then use the same patches to synthesize $I_1 = A\phi(1) + D$. Given the synthesized “initial” and “final” textures I_0 and I_1 , it is easy to find the new A and D . It is possible to also apply texture synthesis to the rate and offset independently, in a similar fashion. However, in this case $R(x, y)$ and $O(x, y)$ are not textures in the conventional sense, but encode an overall variation over the surface, where the rock dries from left to right. In this example, we choose to preserve this overall effect, simply enlarging R and O with standard image processing.

Figure 11 shows how standard static texture-mapping may be combined with TSV-BRDFs. In this case, we use a photograph of an Apple Records logo from a Beatles album, to modulate the TSV-BRDF in the mapping region (with an alpha blend near the edges). Thus, we create the effect of the cut apple logo decaying.

Figure 11 also demonstrates extrapolation on the apple slice dataset, to obtain virtual frames even before the actual start of acquisition. For extrapolation, we simply use our factored representation, evaluating $\phi(t')$, and clamping t' at its overall minimum and maximum value as per equation 3. In this dataset, most of the decay actually happens in the first 30 minutes, and we use input from the

corresponding 10 initial frames only for this figure. We show a significant backing up of the process for many pixels up to $t = -20m$, to the point where the apple is much greener. We can also decay the apple beyond the end-point of the acquisition.

Figure 12 shows how the drying wood can be controlled to create the appearance of drying footprints on a wooden floor. The offsets $O(x, y)$ ensure the floor starts out dry ($O = -1$), while the lower left footprint dries earlier (has a smaller offset $O = 0$, compared to $O = 0.3$ for the upper right footprint). We set $R(x, y)$ to control the rate of drying, depending on the distance from the edge of the footprint. Motivated by observation, the rate is set higher towards the edges and decreases towards the center. We compute a distance transform $d(x, y)$ for points inside the footprint, and set $R(x, y) \sim d^{-1}(x, y)$. Finally, we use a 7×7 Gaussian filter on the resulting maps $R(x, y)$ and $O(x, y)$ to ensure smooth transitions, especially at the edges.

In figure 13, we transfer the rusting steel time-varying process to a new (unrusted) steel plate, using only a single image of its initial condition. The ratio of the new photograph I_{new} to frame 0 of the original sample I_0 is used to modulate both static textures $A_{new}(x, y) = A(x, y) * I_{new}/I_0$ and $D_{new}(x, y) = D(x, y) * I_{new}/I_0$. We then texture-map the time-varying pattern onto a 3D teapot. Note that both diffuse and specular effects, and their time variations are preserved. We also use control to increase the rate of rusting in high curvature regions. In addition, we do edge detection on our static 2D image of the steel plate, to increase the rate near edges. The net rate $R(x, y) = \kappa(x, y)\mu(x, y)$ where μ is an edge map and κ is the average curvature. The insets in the bottom row clearly show that different parts of the object rust at different rates. We have full 3D rendering capabilities, and can see the teapot from different viewpoints while the appearance is evolving.

Figure 14 shows how user-specified patterns can be created in the otherwise natural time-varying processes, with implications for special effects and animations. We texture-mapped the burning wood onto a bowl model; the table cover is from our drying orange cloth dataset. Control is effected through a virtual heat source, for both burning and drying. In addition, we manually modify the rate $R(x, y)$ to resemble the Siggraph logo, for both the bowl and the cloth. For the initial frame, the samples have their normal static appearance. As time progresses, the patterns gradually appear on the bowl and table. With further progression, charring on the bowl and drying of cloth is complete, and the patterns disappear.

7 Conclusions and Future Work

We have presented a complete pipeline from acquisition to rendering for time and space-varying appearance or TSV-BRDFs. This leads to a new capability for computer graphics imagery, to include the dynamic evolution of surfaces and scenes. Our contributions include a newly acquired dataset of time-lapse images for many natural processes from multiple light source and viewing directions, along with estimated parametric TSV-BRDFs. Our main technical contribution is a compact intuitive factored representation that separates spatially varying aspects from temporal variation, being accurate for a variety of natural phenomena. With this representa-

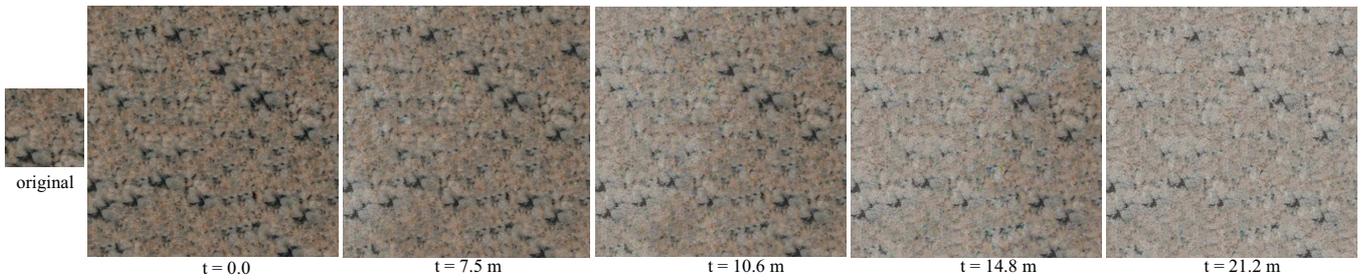


Figure 10: Time-varying texture synthesis can be reduced to 2D synthesis of static spatial textures A and D with our model. We choose to preserve the overall drying pattern from left to right in the original sample.

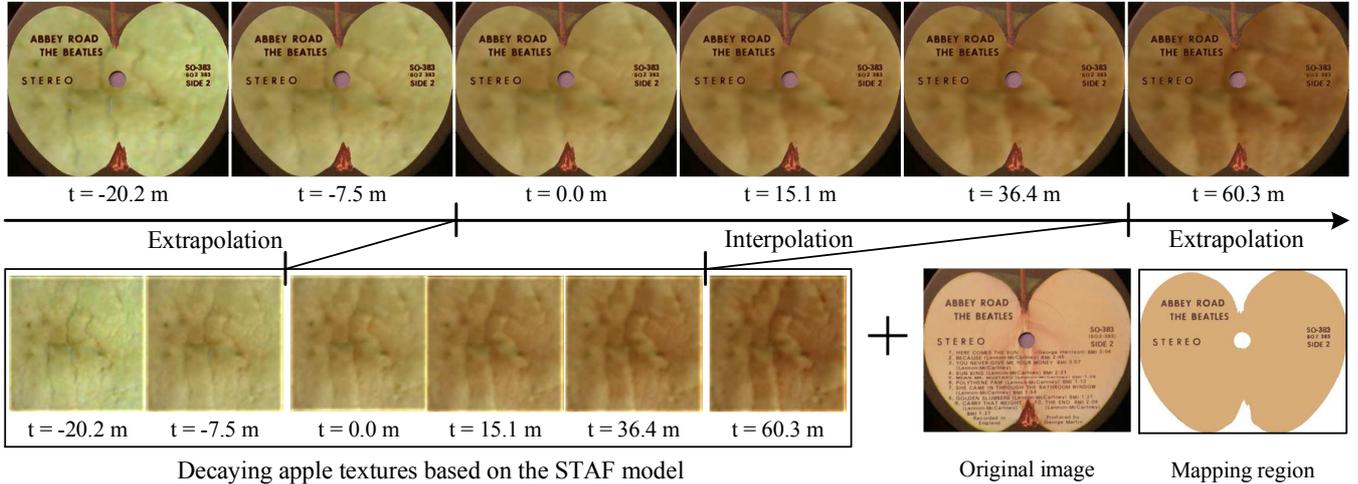


Figure 11: Decaying Apple Records logo, using our apple slice dataset, and modulating by a static texture map of the logo from a Beatles record. This example demonstrates extrapolation, wherein we back up the decay process to considerably before actual start of acquisition—the decay is mostly complete at +30 minutes, and we back up to -20 minutes, getting a much greener look on the apple (we are also able to extrapolate beyond the final time frame).

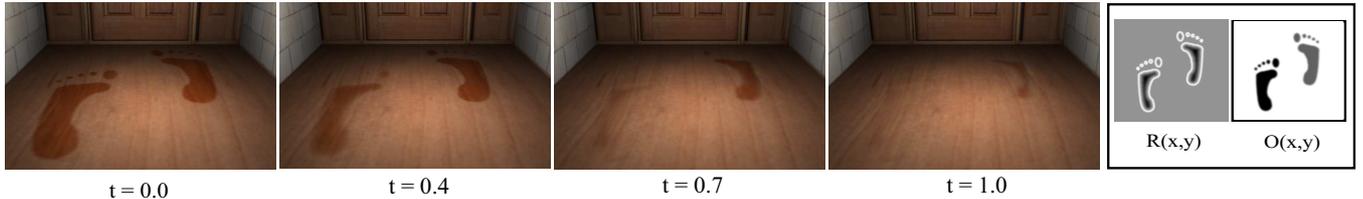


Figure 12: Footprints drying on a wooden floor. We use the drying wood dataset, controlling the rate and offset of drying as shown in the maps on the far right. Specifically, the prints dry faster towards the edges, and the left footprint has a lower offset (higher effective time) and so dries earlier.

tion, we can generalize to a number of novel rendering tasks such as transfer, control, extrapolation and texture synthesis.

In future work, we seek to address current limitations of the STAF model. One example is the decaying apple slice with core, where there are multiple types of time-varying processes occurring so that the “single temporal characteristic curve” assumption does not hold. Extending the STAF model to consist of multiple temporal characteristic curves is one possible solution. Another example is the drying tree bark which is essentially a TBTF instead of TSV-BRDF. Explicitly modeling the relations between neighboring spatial locations may be needed to handle these types of datasets.

The idea of time-varying surface appearance extends beyond the datasets and models reported here. We currently represent the temporal characteristic curve $\phi(t')$, which is effectively the overall time-varying BRDF, in a purely data-driven way, without further analysis. In future work, we seek to understand the time evolution of $\phi(t')$, to develop TBRDF models that would form the time-varying counterpart of common static BRDF models. In general, we believe that measuring, modeling and rendering time-varying surface appearance is an important problem, and this paper is a significant step towards a comprehensive study of this new area.

Acknowledgements: We thank Jason Lawrence for acquiring the first test datasets of time-varying appearance at MERL, and for some early discussions with him and Szymon Rusinkiewicz regard-

ing non-linear factorizations of BRDFs. This work was supported in part by NSF grants 0085864, 0325867, 0305322, 0446916 and a Sloan research fellowship.

References

- CHEN, Y., XIA, L., WONG, T., TONG, X., BAO, H., GUO, B., AND SHUM, H. 2005. Visual simulation of weathering by gammaton tracing. *ACM Transactions on Graphics, SIGGRAPH 2005*.
- CROSS, M., AND HOHENBERG, P. 1993. Pattern formation out of equilibrium. *Reviews of Modern Physics* 65, 851–1112.
- DANA, K., VAN GINNEKEN, B., NAYAR, S., AND KOENDERINK, J. 1999. Reflectance and texture of real-world surfaces. *ACM Transactions on Graphics* 18, 1 (January), 1–34.
- DEBEVEC, P., WENGER, A., TCHOU, C., GARDNER, A., WAESE, J., AND HAWKINS, T. 2002. A lighting reproduction approach to live-action compositing. *ACM TOG (SIGGRAPH 02)* 21, 3, 547–556.
- DORSEY, J., AND HANRAHAN, P. 1996. Modeling and rendering of metallic patinas. In *SIGGRAPH*, 387–396.
- DORSEY, J., PEDERSON, H., AND HANRAHAN, P. 1996. Flow and changes in appearance. In *SIGGRAPH*, 411–420.
- DORSEY, J., EDELMAN, A., WANN JENSEN, H., LEGAKIS, J., AND PEDERSON, H. 1999. Modeling and rendering of weathered stone. In *SIGGRAPH*.

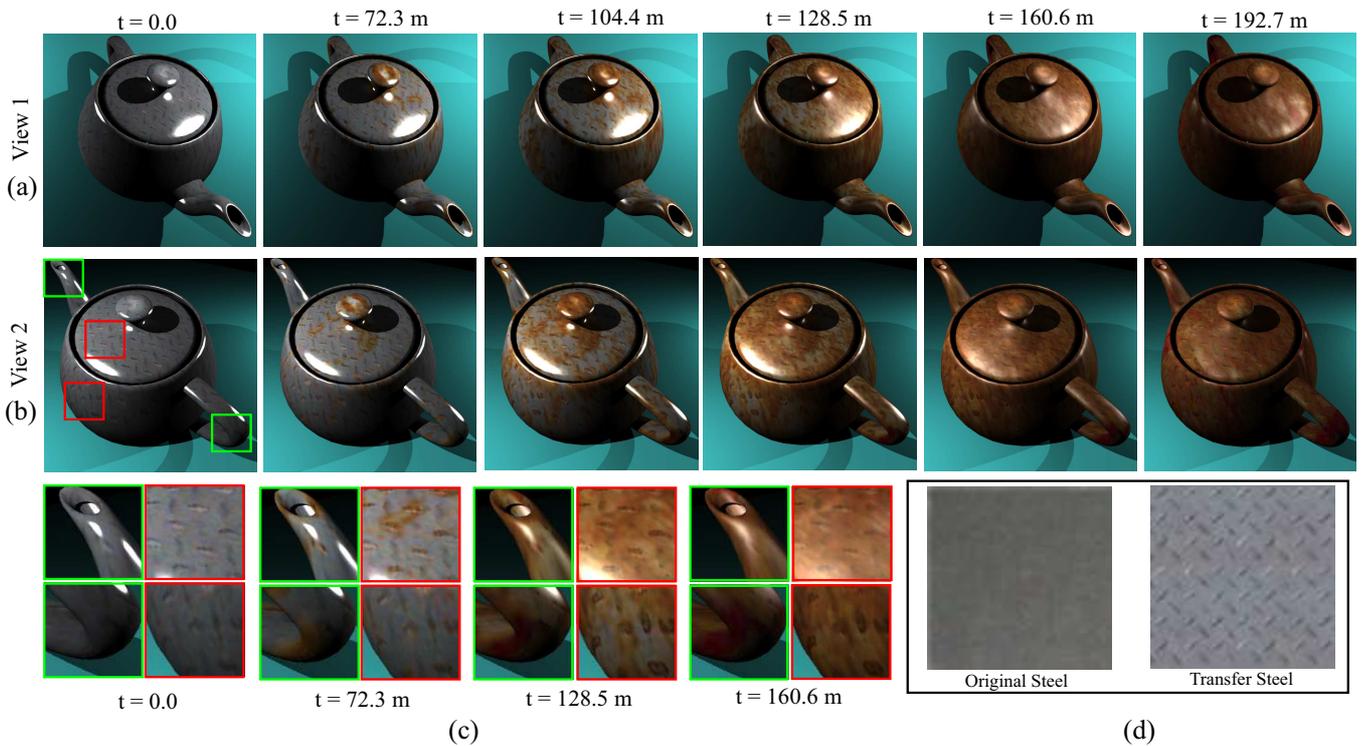


Figure 13: Rusting teapot example. We transfer the rust TSV-BRDF to a new static texture (shown at bottom right (d)). We show two views (a) and (b), with specular and diffuse effects. The insets (c) in the bottom left show the effects of control, where edges and high-curvature areas rust faster.

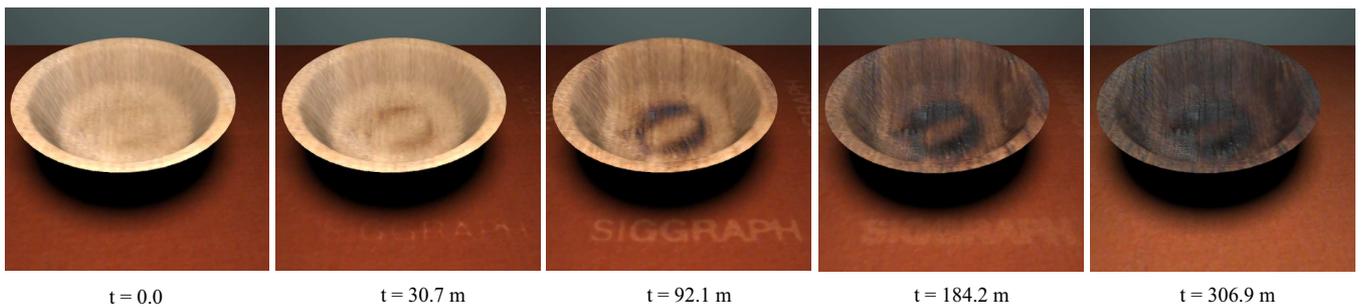


Figure 14: Creating ephemeral patterns by adjusting rates. The bowl uses the burning wood dataset, and the table is drying orange cloth. Control is provided by a virtual heat source. We start with the static appearance, gradually evolving into the Siggraph logo, and then into a fully charred and dry state.

- EFROS, A., AND FREEMAN, W. 2001. Image quilting for texture synthesis and transfer. In *SIGGRAPH*, 341–346.
- GASSER, T., MULLER, H. G., AND MAMMITZSCH, V. 1985. Kernels for nonparametric curve estimation. *Journal of the Royal Statistical Society. Series B (Methodological)* 47, 2, 238–252.
- HSU, S., AND WONG, T. 1995. Simulating dust accumulation. *IEEE Computer Graphics & Applications* 15, 1 (January), 18–22.
- HUGHES, R., AND ROWE, M. 1991. *The colouring, bronzing and patination of metals*. Watson-Guption Publications, New York.
- JANKOWSKY, I., AND SANTOS, G. 2004. Drying behavior and permeability of eucalyptus grandis lumber. In *International Drying Symposium*, 1385–1389.
- JENSEN, H., LEGAKIS, J., AND DORSEY, J. 1999. Rendering of wet materials. In *Rendering Techniques'99*, Springer-Verlag, D. Lischinski and G. Larson, Eds., 273–282.
- KNEIP, A., AND ENGEL, J. 1995. Model estimation in nonlinear regression under shape invariance. *The Annals of Statistics* 23, 2 (April), 551–570.
- KOUDELKA, M. 2004. *Capture, Analysis and Synthesis of Textured Surfaces With Variation in Illumination, Viewpoint and Time*. PhD thesis, Yale University.
- LU, J., GEORGIADIS, A., RUSHMEIER, H., DORSEY, J., AND XU, C. 2005. Synthesis of material drying history: Phenomenon modeling, transferring and rendering. In *Eurographics Workshop on Natural Phenomena*.
- MATUSIK, W., PFISTER, H., BRAND, M., AND MCMILLAN, L. 2003. A data-driven reflectance model. *ACM Transactions on Graphics, SIGGRAPH 2003* (July), 759–769.
- MEINHARDT, H. 1992. Pattern formation in biology: a comparison of models and experiments. *Reports on Progress in Physics* 55, 797–849.
- MERILLOU, S., DISCHLER, J.-M., AND GHAZANFARPOUR, D. 2001. Corrosion: simulating and rendering. In *Graphics Interface 01*.
- MILLER, G. 1994. Efficient algorithms for local and global accessibility shading. In *SIGGRAPH*, 319–326.
- PHARR, M., AND HUMPHREYS, G. 2004. *Physically Based Rendering: From Theory to Implementation*. Morgan Kaufmann.
- SAKOE, H., AND CHIBA, S. 1978. Dynamic programming algorithm optimization for spoken word recognition. *IEEE Trans Acoustic Speech and Signal Processing* 26, 43–49.
- VASILESCU, M., AND TERZOPOULOS, D. 2004. Tensortextures: Multilinear image-based rendering. In *SIGGRAPH*, 336–342.
- VLASIC, D., PFISTER, H., MOLINOV, S., GRZESZCZUK, R., AND MATUSIK, W. 2003. Opacity light fields: Interactive rendering of surface light fields with view-dependent opacity. In *Symposium on Interactive 3D graphics*, 65–74.
- WANG, K., AND GASSER, T. 1999. Synchronizing sample curves non-parametrically. *The Annals of Statistics* 27, 2, 439–460.