

Practical Capture and Reproduction of Phosphorescent Appearance

O. Nalbach¹, H.-P. Seidel¹ and T. Ritschel²

¹Max-Planck Institut für Informatik, Germany

²University College London, United Kingdom



Figure 1: Designed phosphorescent materials before (left) and after having switched off the light for 1 s (middle) and for 80 s (right). Phosphorescent appearance changes differently for different materials over time (Render time 454 ms w., 448 ms w/o. phosp., 640 × 480 px).

Abstract

This paper proposes a pipeline to accurately acquire, efficiently reproduce and intuitively manipulate phosphorescent appearance. In contrast to common appearance models, a model of phosphorescence needs to account for temporal change (decay) and previous illumination (saturation). For reproduction, we propose a rate equation that can be efficiently solved in combination with other illumination in a mixed integro-differential equation system. We describe an acquisition system to measure spectral coefficients of this rate equation for actual materials. Our model is evaluated by comparison to photographs of actual phosphorescent objects. Finally, we propose an artist-friendly interface to control the behavior of phosphorescent materials by specifying spatio-temporal appearance constraints.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Color, shading, shadowing, and texture

1. Introduction

Phosphorescent materials, which re-emit light with a substantial temporal delay after they were exposed to illumination, are widely used in dark conditions where visibility is important but no electric power is available. Typical applications include emergency signs, door handles, light switches, user panels or automotive parts but also toys, artwork and fashion. The practical importance of phosphorescence in this list varies from entertainment to saving human lives. Despite all these use cases, no comprehensive pipeline for digital acquisition, reproduction and manipulation of phosphorescent appearance is available. In this paper, we propose such a framework.

We first devise a phenomenological model of phosphorescence based on rate equations, which allows for efficient simulation under

arbitrary changes of illumination and geometry over time (Sec. 3). Our second contribution is a practical set-up to acquire parameters for our model from physical samples (Sec. 4). The proposed reproduction method allows simulating phosphorescent appearance using the acquired data with only a small computation and storage overhead, despite phosphorescence extending the integration domain by two dimensions, making straightforward Monte Carlo integration prohibitive (Sec. 5) when aiming for real-time performance. Reproduction fidelity is evaluated by comparing rendered images to actual photographs. Our final contribution is a method to design phosphorescent materials. Instead of tweaking abstract rate equation parameters, an artist specifies appearance at keyframes and suitable parameter values for a given time-varying geometry and lighting are computed (Sec. 6).

2. Background and Previous Work

Physics A substance is called *luminescent* if it emits light in response to energy and in particular, *photoluminescent* if it does so in response to light itself. It is necessary to distinguish between photoluminescence and other forms of luminescence due to chemical reactions, other radiation (such as heat) or mechanical pressure which are similar in appearance but different in their origins. If the re-emission happens without significant time-delay (<10 ns) the material is typically called *fluorescent* while *phosphorescence* is characterized by a significant delay, up to hours or even days. Detailed introductions to various aspects of the photoluminescence of solids are available [Lev68] but the precise physical reasons for phosphorescence on the quantum mechanical level are beyond the exposition of this paper. Instead, we focus on three clearly observable phenomena of phosphorescent appearance – decay, saturation and wavelength shift – described in the following.

First, phosphorescent emission is subject to *decay*, where the material once illuminated re-emits a decreasing amount of light over time. The fall-off is often described by differential equations that, under simplified conditions, have exponential functions as closed form solutions [CHZL11]. Different from reflection, phosphorescent emission depends not only on the illumination at the current but also at previous points in time.

Second, phosphorescence is subject to *saturation* and depends on the light previously absorbed in a non-linear way: A material exposed to a certain illumination compared to the same material having received twice that illumination does not necessarily double its emission. At some point, no additional phosphorescence can be induced by additional illumination. Another consequence of this behavior is that, given a configuration of material and incident flux density, after a certain time an equilibrium is reached, where under constant lighting the re-emission is constant, too.

Third, photoluminescence entails a *wavelength-shift*. Absorbed light of one wavelength is re-radiated at a possibly different wavelength. Donaldson [Don54] introduced re-radiation matrices to describe the wavelength shift occurring during re-radiation in fluorescent pigments. For the materials we acquired, no evidence for hue shifts was observed for varying incident wavelengths and therefore we model re-emission using vectors instead of matrices.

Computer Graphics Rendering of fluorescence and phosphorescence effects was introduced by Glassner [Gla95]. However, his model does not account for saturation effects and it is unclear how to handle new excitation over time: materials just have an initial amount of emission which decays over time. No specific image generation approach was proposed, limiting the reproduction method to Monte Carlo (MC) rendering. We would argue that MC approaches, despite their generally increasing success, are not a particularly good fit for phosphorescence. They require to evaluate an integral over the additional, non-separable temporal domain, roughly resulting in a rendering cost per image similar to the cost of rendering an entire image sequence. In our same-time comparisons (Fig. 12), MC has consequently shown to produce a substantial amount of noise. Our approach handles varying excitation in dynamic scenes, includes saturation, and allows for efficient simulation of re-radiation with only moderate computational overhead (milliseconds in our GPU

implementation), independent of the animation length. The inelastic scattering framework by Gutierrez et al. [GMAS05, GSMA08] extends rendering of effects where light is re-radiated at different wavelengths to participating media. Analogously to Glassner's full radiance equation, a full radiative transfer equation is derived. However, no time-dependent effects such as phosphorescence are supported. Using our method, we can synthesize phenomenologically based, but visually convincing, re-radiating volumes (Fig. 13).

Fluorescence of materials can be measured using a Goniofluorometer [BHD*08, HMI08]. Wilkie et al. [WWLP06] and Hullin et al. [HHA*10] acquired bi-spectral, bi-directional reflection and re-radiation distribution functions and then measured for several materials. A high-resolution hyper-spectral imaging system including support for measuring fluorescence was proposed by Kim et al. [KHK*12]. Phosphorescence was not addressed in either of the works. Our additional measurements go beyond the spectral shift alone to cover temporal aspects.

In concurrent work, Alvarez-Cortes et al. [ACKM16] propose to use diffraction filters to image spectral distributions of light sources at low cost while we are using linear variable bandpass filters (LVFs) for this purpose. While pictures taken using these two types of filters both show “images” of the spectrum the spatial resolution is typically much higher for an LVF which is beneficial as we are operating at very low intensity levels and averaging pixels corresponding to the same wavelength helps increasing robustness.

Temporal aspects of appearance due to physical effects were acquired and simulated using stateless, time-dependent BRDFs [GTR*06]. In our approach, we model changes over time as differential equations that express the dependency of the current appearance on the previous illumination state. A related topic in which incident light needs to be integrated over time, too, is fading [KBC*13]. There, the spectral absorption properties of materials change due to chemical reactions. In contrast, in phosphorescence, incident light is re-emitted while the material returns to its ground state.

Finally, material appearance is subject to other advanced physical effects such as diffraction [Sta99], polarization [WTP01] or layering [HK93, JdJM14], which are orthogonal to our approach.

A broad overview on existing methods for artistic material manipulation is provided by Schmidt et al. [SPN*14]. While manipulation of time-dependent appearance, to our knowledge, has not been addressed in previous work, the exponential spatial attenuation present in volumes and tackled by Nowrouzezahrai et al. [NJS*11] bears similarities to the exponential temporal attenuation which we face in phosphorescence.

3. A Model of Phosphorescence

We will first give a motivation and then continue to introduce our phenomenological model which is formulated using rate equations.

Motivation In physics literature, various models exist to describe phosphorescent materials, in particular the decay of phosphorescent emission. Some are closed-form formulas, with or without an interpretation (such as “double exponentials” [MC97]), while others are based on differential equations modeled according to the assumed

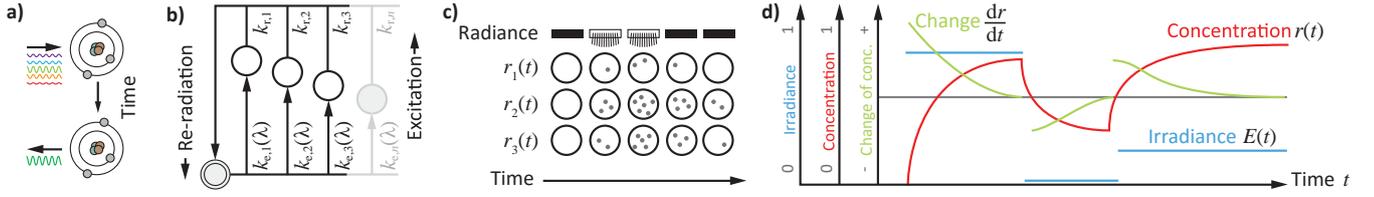


Figure 2: a): A substance is excited by light of one and re-radiates light of a different spectrum. b): From an unexcited state (lower left corner) molecules enter one out of n_c excited states. The excitation rate k_e depends on the exciting wavelength λ , the re-radiation rate k_r is independent of it. c): We model the concentration of molecules in various states. Here, for three states, concentrations are shown as varying dot populations (horizontal). d): Plots of irradiance (blue), continuous concentration (red) and change of concentration (green) over time for one state.

underlying quantum-mechanical principles [JdP84, LSW*01]. However, to our knowledge, there is no simple model describing the concurrent process of simultaneous excitation and re-radiation as found in everyday situations.

In simple terms, phosphorescence is a photoelectric effect which is caused by *excitation* of a substance by photons and later *de-excitation* resulting in *re-radiation* of new photons (Fig. 2, a) [RW45, Lev68]. Upon excitation, molecules enter excited states, from which they eventually return to the ground state, possibly by a transition coinciding with the re-emission of new photons. The re-emitted photons typically have a longer wavelength compared to the photons causing the excitation (Stokes shift). (An exception is Raman scattering whose effect is however too weak to be of practical importance in rendering.) We opt for a state-based phenomenological model that omits any considerations not relevant for plausible reproduction of the appearance of common phosphorescent materials, like effects occurring only at extreme temperatures or under special types of irradiation.

Introduction Fig. 2, b shows the state chart of our model. Upon excitation (Fig. 2, b, motion up), molecules enter one out of several states from which they return to the ground state (Fig. 2, b, motion down). This second transition results in visible light and all computations are aimed at computing this quantity. We did not observe the need to model other physically possible transitions.

Fig. 2, c shows dots for molecules in a discrete excited state which varies over time and depends on the illumination. Instead of using discrete molecules, it is more convenient to model the *concentration* of molecules in different states per unit volume of a substance. In the continuous case, the concentration takes values between 0 and 1, shown for a single state in Fig. 2, d as a thick line. Its change over time (Fig. 2, d, dotted line) depends on illumination (Fig. 2, d, thin line) as well as on its current value to model saturation effects.

The relation between the current illumination, the previous concentration and the new concentration is formalized as a *rate equation*. The stronger the illumination, the higher is the positive change of concentration. Negative changes of concentration, which also result in visible light, are due to the re-radiation. Additionally, as the concentration approaches its maximum, it cannot grow further (saturation). In the same way, when the concentration approaches zero, re-radiation does, too. For a specific illumination, there is a specific concentration where the increase and decrease cancel each other out such that the appearance reaches an equilibrium. These

effects are responsible for the non-linearities depicted in Fig. 2, d. Computing the concentrations means to solve differential equations, e.g., using numerical integration such as the forward Euler method.

In practice, the concentration is also spatially-variant and illumination itself is coupled with phosphorescence at other spatial locations described by integral equations which are solved using light transport simulation e.g., by Monte Carlo rendering. In summary, rendering phosphorescence amounts to solving a space/time/wavelength-varying system of integro-differential equations.

Rate equations Let the i -th of n_c concentrations at position \mathbf{x} and time t be denoted as $r_i(\mathbf{x}, t)$ and the (spectral) irradiance as $E(\mathbf{x}, \lambda, t)$. We define the change of concentration as

$$\frac{dr_i}{dt}(\mathbf{x}, t) = \underbrace{(1 - r_i(\mathbf{x}, t)) \int k_{e,i}(\lambda) E(\mathbf{x}, \lambda, t) d\lambda}_{\text{Excitation}} - \underbrace{k_{r,i} r_i(\mathbf{x}, t)}_{\text{Re-radiation}}, \quad (1)$$

where $k_{e,i}$ and $k_{r,i}$ are the *excitation* and *re-radiation* rates and $1 - r_i(\mathbf{x}, t)$ is the capacity remaining to full saturation. For the first-order ODE (Eq. 1) to have a unique solution, initial conditions are necessary. This amounts to specifying all initial concentrations at the start of the simulation ($t = 0$) for all spatial locations \mathbf{x} , e.g., $r_i(\mathbf{x}, 0) = 0$.

While Eq. 1 models the internal state of phosphorescent materials and its changes, the visible light caused by the re-radiation is introduced into image synthesis as an additional isotropic emissive component:

$$L_o(\mathbf{x}, \omega, \lambda, t) = \underbrace{L_e(\mathbf{x}, \omega, \lambda, t)}_{\text{Emission}} + \underbrace{L_p(\mathbf{x}, \lambda, t)}_{\text{Phosph.}} + \underbrace{\mathbf{R}L_i(\mathbf{x}, \omega, \lambda, t)}_{\text{Reflection}}, \quad (2)$$

where L_o is the outgoing radiance, L_e is the classic emissive radiance, L_p is the phosphorescent emissive radiance, L_i is the incoming radiance and \mathbf{R} the reflection operator [ATS94] to convert incoming to outgoing radiance. The phosphorescent radiance occurs at multiple wavelengths which we model as

$$L_p(\mathbf{x}, \lambda, t) = \sum_{i=1}^{n_c} \Lambda_i(\lambda) k_{r,i} r_i(\mathbf{x}, t), \quad (3)$$

where $\Lambda_i(\lambda)$ is the per-state *re-radiation* function that models how much state i re-radiates at wavelength λ . We discuss solving Eq. 2 in Sec. 5.

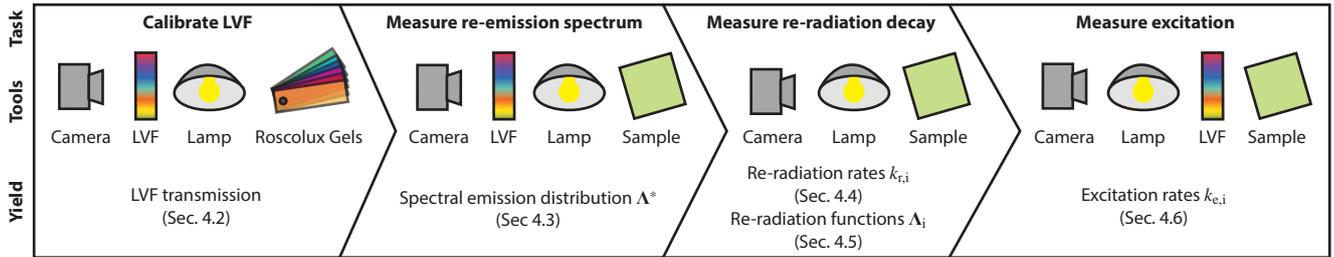


Figure 3: The acquisition workflow. After calibrating the LVF (Sec. 4.2) it is used to measure the re-emission’s spectrum (Sec. 4.3). Next, the decay process is sampled which allows to determine re-radiation rates (Sec. 4.4) and functions (Sec. 4.5). Finally, excitation by different wavelengths is determined using the LVF to produce the different wavelength bands simultaneously, and excitation rates are fit (Sec. 4.6).

Limitations Our model makes a few simplifying assumptions. First, we do not consider potential effects of the thickness of the phosphorescent substance and therefore only capture the specific material thickness used in the respective measurements in Sec. 4.

Second, all material samples used, behaved mostly isotropic in both excitation and emission. Measuring the emission from five different angles (0 to 80° in 20° steps) we found the fluctuations in intensity to be less than 5%. We consequently do not model directional dependency. Layered materials [HK93, JdJM14] where a phosphorescent substance is covered by additional layers might behave differently, which is however orthogonal and independent of the nature of phosphorescence.

Finally, we do not explicitly enforce Stokes shift as the state of the material, given by the state concentrations r_i , does not track the precise wavelength of the photons having caused excitation. Thus, technically, excitation from one wavelength may result in re-radiation at a shorter wavelength. Yet, the overlap of the non-zero regions of the excitation spectra, given by the $k_{e,i}(\lambda)$, and the re-radiation spectra, given by the $\Lambda_i(\lambda)$, is typically very small, largely excluding this possibility. We additionally measured the re-radiation of our materials after excitation by several light source types with different emission spectra and found that they were virtually the same. An example is found in the supplemental material.

4. Acquisition

To reproduce a specific material using our model, we need to acquire the constants $k_{r,i}$ and $k_{e,i}$ (Eq. 1) as well as Λ_i (Eq. 3) by measurements. We draw inspiration from existing measurements of similar parameters for models that are suitable for describing, but not for effectively simulating phosphorescence [RW45]. Typically, a sample is illuminated in a specific way, the response over time is captured and this observation is used to fit the material parameters. After describing our general setup (Sec. 4.1) and spectral acquisition method (Sec. 4.2), we discuss the measurements of the re-radiation’s spectral distribution (Sec. 4.3), re-radiation (Sec. 4.4) and excitation rates (Sec. 4.6) as well as of the re-radiation functions (Sec. 4.5). Fig. 3 gives an overview over the different acquisition steps and the parameters estimated. Our obtained parameters for various substances and a comparison to photographs are shown in Sec. 4.7. For brevity, we drop the spatial variation in this section and always consider a specific location \mathbf{x} .

4.1. General Setup

Our acquisition setup is seen in Fig. 4. In general, it consists of a strong light source with known, adjustable spectrum, such as a lamp with a smooth spectrum in combination with bandpass filters or a monochromator, as well as an instrument to capture the emission spectra of the samples like a colorimeter or a simple digital RGB or gray camera combined with additional bandpass filter(s). We take the latter approach using a Canon 5D Mark II as detailed in Sec. 4.2. The camera’s high resolution is beneficial as it allows for averaging measurements across multiple pixels reducing the effect of noise.

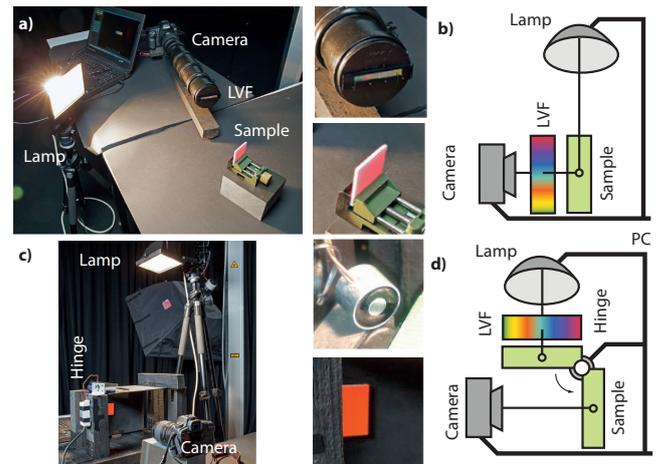


Figure 4: Our two measurement setups as described in Sec. 4.1.

We used a 400W Halogen lamp with smooth output spectrum as our light source (cf. the lamp’s spectral distribution in the supplemental material). For some measurements, we shaped the light’s spectrum using a *linear variable bandpass filter* (LVF), as seen in Fig. 4, c,d. The transmission spectrum of this type of filter corresponds to a narrow spectral band of about 25 nm, where the center of the band varies linearly across the filter. The room was strictly dark otherwise. Xenon lamps constitute an alternative type of light source to combine with an LVF, with the advantage of having stronger output in the near-UV range, but also coming with drawbacks such as being more expensive, more tenuous to operate and having spikes in the emission spectrum.

During the actual measurements, we adjusted the camera settings (ISO and aperture) depending on the emission intensity of the different materials and performed a relative calibration of the resulting values afterwards to account for this. We deemed the camera's dynamic range sufficient to capture the samples' emission.

Before obtaining a measurement series, we allowed the lamp and camera to reach a stable temperature to eliminate temperature-dependent fluctuations in the sensor's noise level or the lamp's emission. For computing the fits in Tbl. 1, we discretized all spectra into $n_\lambda = 8$ spectral bands i.e., uniform intervals over the range from 400 to 700 nm.

4.2. Spectral Acquisition using an LVF

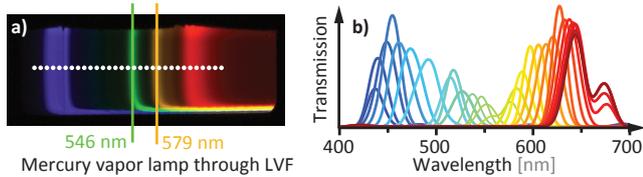


Figure 5: a): Spectrum of a mercury vapor lamp seen through an LVF. The spectral lines' wavelengths are known and hence used to select sampling points (white dots). b): Each line corresponds to the LVF's transmission at one of the sampling points highlighted in a).

Professional spectrometers are expensive. An alternative is taking RGB images of the respective sample using a camera with known spectral sensitivity, through a number of cheap color filters with different, known transmission curves, and solving the resulting system of linear equations to find the spectrum best explaining the measured values in a least-squares sense [KO11]. Unfortunately, phosphorescence is a time-varying effect requiring to take these images at the same instant, which is hard to achieve with common filters. An LVF, however, can be seen as a set of many filters with well-localized transmission curves in compact form and allows to take many differently filtered images at the same time (Fig. 4, a,b), while still being less pricey than a spectrometer.

Before being able to take spectral images, the LVF has to be calibrated by computing its own spectral transmission curves at each spatial location. For this, we first identified the locations of the 400 and 700 nm bands on the filter by localizing the spectral lines of a mercury vapor lamp and extrapolating them linearly. Then, we picked a set of 31 equidistant sample locations from 400 to 700 nm (Fig. 5, a). To determine the spectral transmission at these locations, we first combined the LVF with each of a set of 108 Roscolux color filters and imaged a neutral gray sample (with known reflectance), lit by a white LED light with known spectral distribution, before optimizing for the most likely transmission at each of the 31 locations (Fig. 5, b), analogously to Kirk and O'Brien [KO11]. During this and other LVF-based measurements, we always averaged the observations (spatially) over all pixels corresponding to the same wavelength.

Having calibrated the LVF, each picture now taken through it (Fig. 4, a,b) provides information about the RGB response for 31 different bandpass filters and to optimize for the spectrum of the

measured signal. The procedure is the same as for calibrating the LVF, just that now the signal is unknown while the filters spectral distribution is known. We validated our spectral acquisition method by imaging patches of a color checker with known reflectances and found it to perform on par with other approaches using similar tools, the results can be seen in the supplemental.

4.3. Re-radiation Spectrum

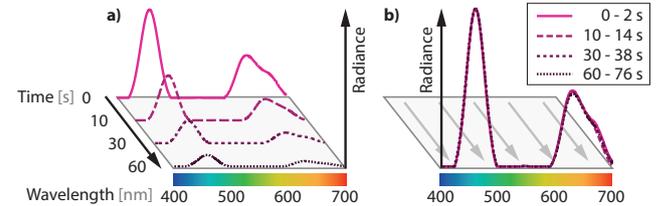


Figure 6: Constancy of the spectral distribution. a): The emission spectrum is only scaled globally over time. This can be seen even more clearly when dividing the curves by their maximum value (b). Also, note the increasing integration times (box) necessary to capture images with a similar signal to noise ratio.

To instantiate our model, the spectral emission of a material has to be captured over time. While capturing the time-dependent behavior demands for dense temporal sampling, the spectral acquisition asks for low noise images which in turn require long exposure times given the low intensity nature of phosphorescent emission. We, however, found that the relative spectral intensity of our material samples was constant over time, as can be seen for one sample in Fig. 6. This allows to decouple the measurement of this relative distribution from the measurement of the change of the absolute intensity over time, yielding a more robust fit. For these reasons, we first determined the spectral distribution $\Lambda^*(\lambda)$ of the emission separately and normalized it to have a maximum value of 1.

4.4. Re-radiation Rate

By measuring how emission decays after turning the illuminant off, we can estimate the re-radiation rates. In the absence of illumination, $E(\lambda, t) = 0$, Eq. 1 has the closed form solution

$$r_i(t) = r_i(0) \cdot e^{-k_{\epsilon,i}t}. \quad (4)$$

By measuring $L_p(\lambda, t)$ (Eq. 3) we can find the underlying re-radiation rate constant.

Procedure Initially, we illuminate the sample by a strong all-frequency light for at least 90 minutes (Fig. 4, a,b). This time is typically sufficient to achieve saturation, such that we can assume that $r_i(0) = 1, \forall i$, which is useful when fitting the re-radiation functions (Sec. 4.5). Next, we turn the lamp off and sample $L_p(\lambda, t)$ by taking direct RGB photographs, typically between $t = \epsilon$ and $t = \epsilon + 180s$, where ϵ is the delay between turning the light off and starting the measurement which we keep to a constant value of a few milliseconds by operating lamp and camera using a PC. Special care has to be taken if the lamp exhibits an afterglow; in this case, the lamp has to be shielded in addition to turning it off to

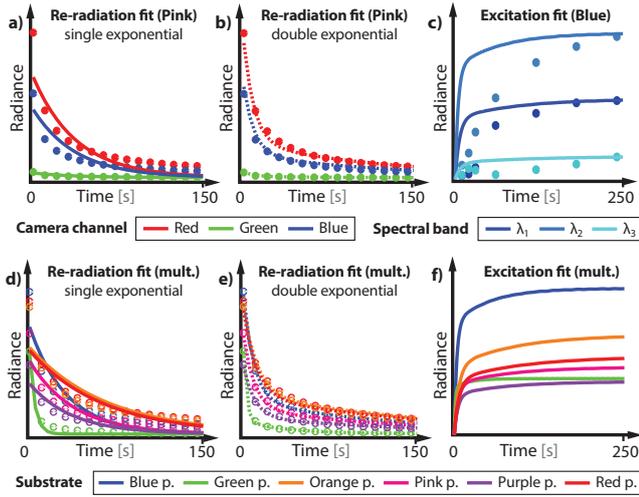


Figure 7: Measurements and fits. a): Re-radiation measurements (points) fitted by one exponential (lines) for “Pink paint”. b): A double exponential fitted to the same data. d), e): Decay fits for various materials using one or a sum of two exponential functions, respectively. c): Excitation measurements and ODE simulation (Sec. 4.6) for “Blue paint” and for others (f).

avoid a reflection by the sample. We found taking one image every 5 seconds to be sufficient. Fig. 7, a and b show the resulting curves for a sample material as points.

Fitting As the spectral distribution is constant over time (Sec. 4.3), the measured RGB values are multiples of the camera’s RGB response to the spectrum given by $\Lambda^*(\lambda)$, which we denote by $\mathbf{\Lambda}_3^* \in \mathcal{R}^3$. We (simultaneously) fit a sum of n_c exponential decay functions of the form

$$\sum_{i=0}^{n_c} a_i \cdot \mathbf{\Lambda}_3^* \cdot e^{-k_{r,i}t}$$

to the observed RGB values of $L_p(\lambda_j, t)$ for $\lambda_j \in \{1, 2, 3\}$ using Matlab’s constrained non-linear optimization functionality, restricting the parameters to the positive range. Constancy of the spectral distribution implies the single scaling factor a_i and a single decay rate $k_{r,i}$ per exponential term which are shared among all wavelengths. Note that, while fitting a single exponential function corresponds to fitting a line in log-space, this cannot be generalized to sums of exponential functions. An example fit for one material is shown in Fig. 7, a and b. Additional individual plots are included in the supplemental material. Fits for multiple other materials are shown in Fig. 7, d and e as well as in Tbl. 1.

Discussion We found that $n_c = 2$ produces a good fit for all measured materials (cf. Fig. 7, b). As outlined in Sec. 3, such a “double exponential” fit for decay curves is common in physics [CHZL11], with explanations including multiple trap states [MC97], (Eq. 138). It is worth noting that stretched exponentials [LSW*01] or power law decays [JdP84] provide different interpretations leading to different ODEs which however contain the time since illumination as a term, defying efficient simulation in general conditions. A fit

using $n_c = 1$ provides an approximation with larger numerical error but still resembles typical phosphorescence (Fig. 8). The different shapes of Fig. 7, d and e show that different substances have different decay behavior that can be substantially different from a single exponential.

Numerically, using two (three) exponentials compared to only one, the median approximation error of the fit over all materials is reduced to 8% (4%) of its original magnitude. The main improvement of the fit concerns the long-term error: The error for the last time sample is reduced to 18% (3%). We however found $n_c = 2$ to be a good compromise as adding one exponential more would also roughly increase evaluation time by 50% (Sec. 5).

By measuring the decay from different excitation levels, we verified that the shape of the curve is independent of the excitation state i.e., that the $k_{r,i}$ are the same for all curves (cf. the supplemental material for a corresponding plot). Also, constancy of the spectrum (Fig. 6) indicates that a single $k_{r,i}$ shared by all wavelengths, as used in our model, is appropriate.

4.5. Re-radiation Function

The re-radiation function $\Lambda_i(\lambda)$ models to which extent excitation in state i results in re-radiation at wavelength λ . For n_λ discrete wavelength bands, the operator becomes a vector $\mathbf{\Lambda}_i \in \mathcal{R}^{n_\lambda}$. As outlined in the following, we can re-use the data acquired as in Sec. 4.4 without conducting further measurements.

Fitting First consider $L_p(\lambda_j, t)$ during the experiment conducted in Sec. 4.4 and denote the re-radiation vectors in this case as $\mathbf{\Lambda}_{3i} \in \mathcal{R}^3$. With $n_\lambda = 3$, for a given position, Eq. 3 becomes

$$L_p(\lambda_j, t) = \sum_{i=1}^{n_c} \mathbf{\Lambda}_{3i,j} k_{r,i} r_i(t) \quad (5)$$

$$= \sum_{i=1}^{n_c} \mathbf{\Lambda}_{3i,j} k_{r,i} e^{-k_{r,i}t} = \sum_{i=0}^{n_c} a_i \cdot \mathbf{\Lambda}_3^* \cdot e^{-k_{r,i}t}. \quad (6)$$

As in the second line, we have an equality between two sums of exponential functions, we can compare their coefficients, which have to be identical for the equality to hold i.e., $\mathbf{\Lambda}_{3i,j} k_{r,i} \stackrel{!}{=} a_i \cdot \mathbf{\Lambda}_3^*$. This yields $\mathbf{\Lambda}_{3i,j} = a_i \cdot \mathbf{\Lambda}_3^* / k_{r,i}$. We obtain the spectral $\mathbf{\Lambda}_{i,j}$ by applying the scaling found using the RGB case to the discretized version $\mathbf{\Lambda}^*$ of $\Lambda^*(\lambda)$ i.e., $\mathbf{\Lambda}_{i,j} = a_i \cdot \mathbf{\Lambda}_j^* / k_{r,i}$.

Discussion The second equality holds since $r_i(0) = 1$ and the illumination is turned off (i.e., under the conditions of the experiment in Sec. 4.4). As in general $a_i \neq a_k$ for $i \neq k$, one $\mathbf{\Lambda}_i$ per state is necessary to reproduce the measured emission behavior, where all $\mathbf{\Lambda}_i$ are however scaled versions of the per-material $\mathbf{\Lambda}^*$.

4.6. Excitation Rate

After the previous experiments, re-radiation rate and function are known. They describe how fast a material loses excitation and the color of the resulting emission. To complete the parameter set, what remains is the excitation rate, i.e., how fast excitation is gained. Conceptually, we would like to observe the change of the phosphorescent intensity of an initially unexcited sample being exposed

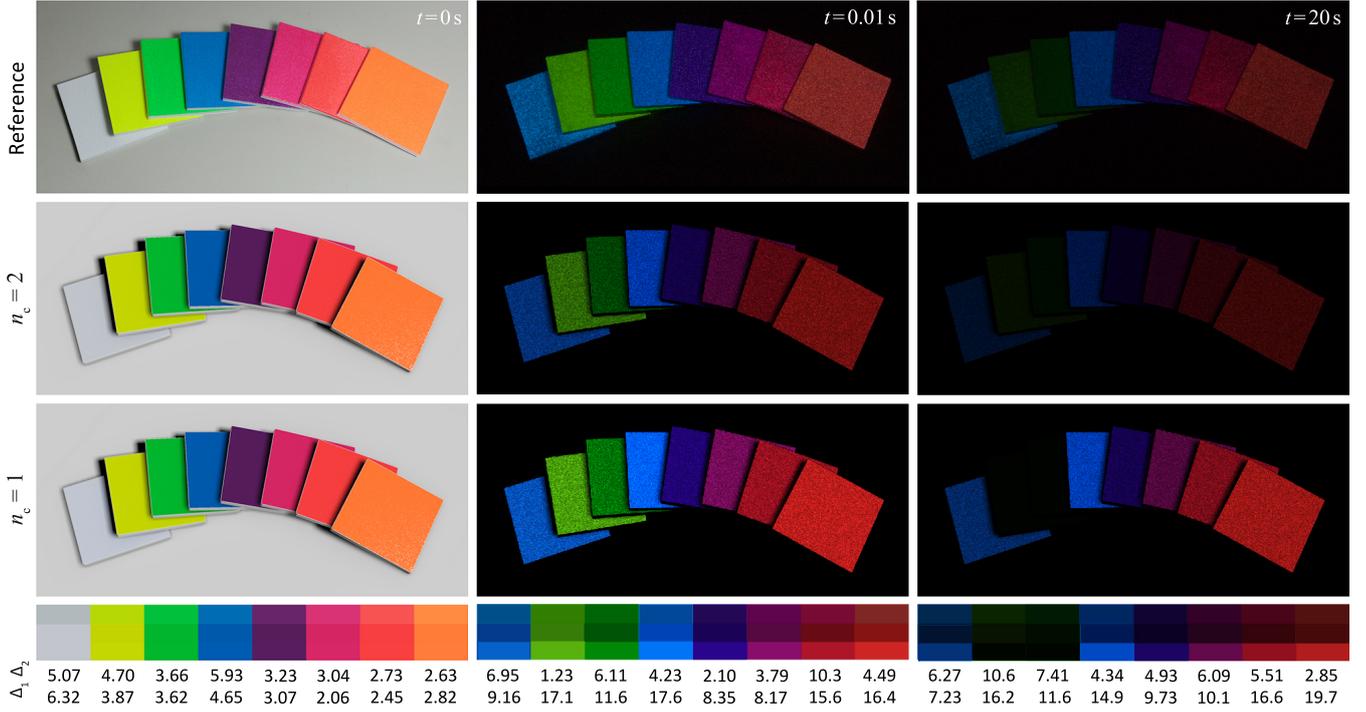


Figure 8: Materials acquired with $n_\lambda = 8$ and using single ($n_c = 1$) and double exponential ($n_c = 2$) fits compared to photographs. For display and to enable comparisons, spectral re-emission is mapped back to RGB using the camera sensitivity curves. The actual emission spectra are displayed in Fig. 9. The CIEDE2000 differences [SWD05] to the reference for $n_c = i$ are labeled with Δ_i . The avg./min./max. error over all samples and points in time amounts to 5.46/1.23/10.64 for the double- and 13.16/7.24/19.75 for the single-exponential fit.

Table 1: Re-radiation and excitation rates in units s^{-1} . The first column shows the daylight color of the samples, the second the decay gradients (scale of 60 s). The response for some wavelengths was too low to be measured with our setup, thus excitation rates are only given for the first three wavelength bands. For the last three materials, only their decay behavior was measured.

| Decay | Material | k_r | λ_1 | k_e λ_2 | λ_3 | λ_1 | λ_2 | λ_3 | λ_4 | Λ λ_5 | λ_6 | λ_7 | λ_8 |
|-------|--------------|-------|-------------|----------------------|-------------|-------------|-------------|-------------|-------------|--------------------------|-------------|-------------|-------------|
| | Blue paint | 0.032 | 0.0170 | 0.0500 | 0.0040 | 0.34 | 4.37 | 3.65 | 0.55 | 0.0 | 0.01 | 0.0 | 0.0 |
| | Green paint | 0.227 | 0.1000 | 0.0350 | 0.0060 | 0.0 | 0.0 | 0.07 | 0.64 | 0.05 | 0.0 | 0.0 | 0.0 |
| | Orange paint | 0.015 | 0.0150 | 0.0060 | 0.0003 | 0.04 | 1.85 | 0.47 | 0.0 | 0.45 | 7.70 | 4.88 | 2.49 |
| | Pink paint | 0.027 | 0.0060 | 0.0150 | 0.0033 | 0.59 | 4.07 | 0.49 | 0.0 | 0.0 | 1.22 | 1.77 | 0.09 |
| | Purple paint | 0.028 | 0.0120 | 0.0180 | 0.0008 | 0.46 | 3.08 | 0.36 | 0.0 | 0.0 | 0.19 | 0.08 | 0.0 |
| | Red paint | 0.016 | 0.0120 | 0.0038 | 0.0000 | 0.13 | 3.03 | 0.44 | 0.0 | 0.02 | 5.49 | 6.21 | 0.63 |
| | White paint | 0.047 | 0.0180 | 0.0400 | 0.0058 | 0.26 | 2.55 | 2.49 | 0.61 | 0.0 | 0.01 | 0.0 | 0.10 |
| | Yellow paint | 0.251 | 0.1450 | 0.0380 | 0.0070 | 0.0 | 0.0 | 0.08 | 0.80 | 0.15 | 0.12 | 0.02 | 0.0 |
| | Bottle | 0.185 | - | - | - | 0.0 | 0.04 | 0.41 | 0.58 | 0.25 | 0.06 | 0.01 | 0.0 |
| | Gloves | 0.067 | - | - | - | 0.10 | 0.88 | 5.29 | 5.28 | 2.01 | 0.66 | 0.12 | 0.01 |
| | Stickers | 0.039 | - | - | - | 0.01 | 6.02 | 40.75 | 40.66 | 14.86 | 5.08 | 1.11 | 0.14 |

to a constant flux of a certain wavelength band until equilibrium is reached. Using the other, already known, parameters, we could then fit the excitation rate in Eq. 1 so that the shape of the curve and, above all, the equilibrium level, are matched best. However, in practice the phosphorescent emission can only be measured when irradiance is very small compared to the former as it is otherwise dwarfed by the effects of reflection. As an approximation, we mea-

sure the phosphorescent emission directly after the illuminant is switched off.

Procedure Before starting measurements, we keep the sample in darkness for several hours until no emission can be detected anymore. Next, we expose it to an illuminant of a specific wavelength band for t seconds (Fig. 4, b), then switch the latter off and measure the emission immediately to obtain an estimate of $L_p(t)$. We execute

this process for all n_λ wavelength bands and different excitation times. As the time for the excitation to reach an equilibrium is typically unknown, we recommend to initially double the excitation time with each sample until emission does not change anymore before selecting further intermediate sampling times as desired.

This procedure is more challenging to perform than the previous ones. While re-radiation can be easily measured in a single run (i.e., with linear effort in the number of samples), to obtain emission after time t , an uncharged sample has to be previously charged for precisely this amount of time (quadratic effort). To accelerate the process, we perform the experiment for all wavelengths in parallel by illuminating the sample through an LVF placed directly on top of the sample, so that different spatial locations correspond to different excitation wavelengths. To quickly remove the LVF from the sample in order to take the actual measurement, we built an automated trap door-like mechanism in which the sample is swung away from the LVF and directed towards the camera, before releasing the shutter, immediately when the light is turned off (cf. Fig. 4, c,d). In our setup, the light takes about 320 ms to fade, then the door begins to open and stops 280 ms later.

Fitting As the excitation rate is wavelength-dependent, we perform the fit for the n_λ different wavelength bands independently, using the respective excitation-curve measured for light of that band. We chose the excitation rate $k_{e,i}(\lambda_j)$ for band j using Matlab’s non-linear least-squares regression such that, when simulated using the measured spectrum E of the light source and the already determined $k_{r,i}$ and $\Lambda_i(\lambda_j)$, it explains the sample curve best. Note, that in our case the sample curve consists of (linear) RGB values measured by the camera and, in order to be able to compare them to the spectral simulation results, the latter have to be mapped to RGB, too, using the camera’s spectral sensitivity curves. We used [KHTI13] as a source for the curves for our camera model. For numerical reasons, we normalized the length of E to 1 in the simulations instead of scaling it correctly relative to the samples of L_p (i.e., performing an absolute calibration) to avoid tiny values of $k_{e,i}$ which would only complicate reproduction. The resulting samples and a fit using $n_c = 1$ is seen in Fig. 7, c. The spectrum of the light source is contained in the supplemental material.

Discussion Our excitation measurement process bears two potential sources of imprecision: the short delay when moving the sample and camera noise. Furthermore, slight fluctuations of the aforementioned delay and vibrations of the sample when our trap door stops may lead to inconsistencies. Comparing Fig. 7, a and Fig. 7, c it can be seen that the re-radiation fit is precise, while the excitation fit reaches the correct equilibrium but at a faster rate than measured. Despite the imprecision in the measurement, we attribute this mostly to the model itself. One reason for this is that we fit the re-radiation first to match it as well as possible, as differences in the re-radiation are very noticeable when observed in darkness. The speed at which the (correct) equilibrium is reached during excitation, however, is of lesser importance as, while the illuminant is switched on, the phosphorescent emission is dwarfed by reflection and the process of excitation cannot be observed by the naked eye.

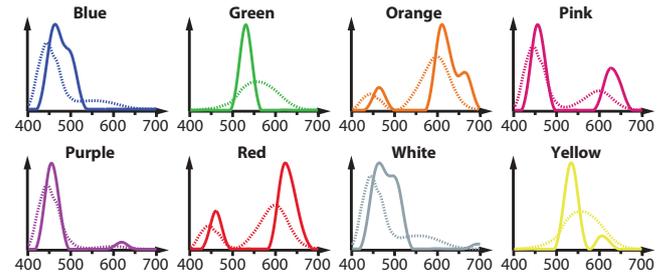


Figure 9: Spectral intensity distributions (solid lines) of the samples shown in Fig. 8, measured according to Sec. 4.2. Relative spectral radiance is plotted for varying wavelength [nm]. We also show the corresponding best approximations as a weighted sum of the CIE2006 color matching functions (dashed) as an indication for how well the spectra can be represented in a three channel color space. The spectral acquisition captures the localized peaks better.

4.7. Results

We acquired parameters for 11 phosphorescent materials, from special paints to everyday objects. As previously described, after initially measuring the emission spectra at a 10nm resolution using the calibrated LVF (Sec. 4.2), we used $n_\lambda = 8$ wavelength “bands” of 37.5 nm width, from 400 to 700 nm for the remaining fit. The outcome is shown in Tbl. 1 and Fig. 9. For the parameters to a double exponential model, please see the supplemental material.

Fig. 8 compares synthetic images, rendered using the acquired data according to Sec. 5, to actual photographs. The most prominent differences appear for the yellow and green paints which fade particularly fast before continuing to glow at low brightness. Especially the fit with only one exponential fails to capture this dualistic behavior, while the double exponential fit can cover both short and long-term behavior. For some materials their albedo is close to their re-radiation spectrum, which is however no general rule as can be seen e.g., for “White paint” and “Stickers” (Tbl. 1).

The acquisition process might appear to be very time consuming, however this is only true if only one physical material sample is available. Determining the emission spectrum (Sec. 4.2) merely requires a single photo at a sufficient excitation level. For the re-radiation acquisition, exciting the sample to saturation usually only takes half an hour. The following measurement only takes about 3 minutes in which the decay curve is sampled. For the excitation rates, if a sufficient number of samples is at hand (about 10 per material), it is not necessary to wait for de-excitation of the samples and the total time is reduced to not more than two hours per material.

5. Reproduction

To reproduce phosphorescence, Eq. 1, an ordinary differential equation, has to be solved. As the irradiance E appears in Eq. 1, this also requires solving the time-stationary rendering equation which itself depends on the emission from previous points in time.

Representation To solve Eq. 1 we need a spatial discretization of the concentration r_i . For this, the scene surfaces are sampled into



Figure 10: Photographs of phosphorescent objects compared to results rendered using our approach with designed materials.

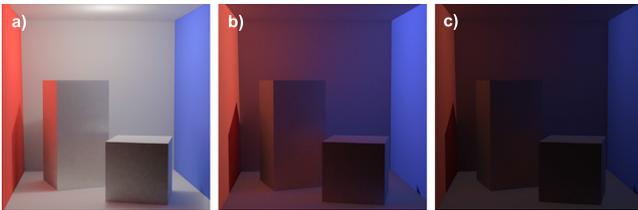


Figure 11: Path-traced Cornell box with phosphorescent red and blue walls and non-phosphorescent boxes: Initial illumination (a) and after the light has been turned off for 0.1 (b) and 60 seconds (c), respectively. Non-phosphorescent areas are lit by bounced phosphorescent emission. The materials used are artificial.

a point cloud. This sampling pattern is held constant but deformed in every frame, assuming isometric deformations of objects. A particularly efficient sampling of this form avoiding explicit sample locations is produced by parameterizing the phosphorescent surfaces and associating one sample with each texel of a texture. We used the UV mapping functions of Blender for this. Re-radiation and excitation parameters are stored as per-material constants or, if spatial variation is desired, in textures.

Solver We perform incremental Euler integration, starting at frame $j = 1$. At each sample position, we store n_c floats in the range $[0, 1]$ to maintain the r_i which are initialized to 0. Let $\mathbf{r}_i^{(j)}$ denote the value of r_i at frame j . Frame $j > 1$ is produced as follows. First the irradiance $\mathbf{E}^{(j)}$ is computed using spectral rendering to solve Eq. 2 and stored as an n_λ -component vector. This may include the phosphorescent emission from previous time steps. Next, the change of ratio $\dot{\mathbf{r}}_i^{(j)}$ is computed according to Eq. 1 using $\mathbf{E}^{(j)}$ and the per-material or per-texel parameters. Finally, $\mathbf{r}^{(j)}$ is updated by $\mathbf{r}_i^{(j+1)} = \mathbf{r}_i^{(j)} + \Delta_t \dot{\mathbf{r}}_i^{(j)}$, where Δ_t is the time step size.

Discussion We use simple Euler integration to solve Eq. 1 as we found no need to use a more advanced solver after comparing to other solvers e.g., Runge-Kutta (cf. supplemental material). Using

those would be straightforward, though. Our discretization in time is synced to the rendering speed. We observed this to be sufficient for the materials and lighting conditions we used. Even the highly space-time discontinuous functions in Fig. 12 are captured well. The spatial discretization of \mathbf{r}_i has to be fine enough to capture the spatial variation of illumination and phosphorescence parameters.

Results We implemented the solver using GPU shaders. As an example, for Fig. 1, 4.7 m sampling points were used, for which it required 3.4 ms on an Nvidia Geforce GTX 770 to update their values in parallel, including the time for computing the irradiance from direct light using shadow maps. The results in Fig. 1, Fig. 10, Fig. 12 and Fig. 14 have been post-processed to emulate scotopic vision by the method of Thompson et al. [TSF02], i.e., applying slight amounts of blur and noise to enhance their nocturnal appearance.

A comparison of real and reproduced phosphorescence for complex object is shown in Fig. 10. Fig. 12 demonstrates the approach for animated scenes in real time. Fig. 11 shows global phosphorescent light transport. Please see the supplemental video for animations. Finally, in Fig. 13, we demonstrate an extension to the volume case based on a smoothed-particle hydrodynamics (SPH) simulation where each particle represents one “sample” and tracks its own current level of excitation.

Comparison to Monte Carlo rendering The only existing framework to render phosphorescence is due to Glassner [Gla95]. Since it is given as a closed-form formula, sampling irradiance over previous points in time is necessary to determine the current phosphorescent emission. This either requires an irradiance cache that grows linearly with simulation length or re-sampling of the irradiance in each time step leading to a sampling effort for the phosphorescence component which grows quadratically with simulation length. Furthermore, the evaluation of the decay cannot be cached and also leads to quadratic effort in time. Either of these options defies use in interactive applications (Fig. 12), while for our method memory consumption and effort per frame are constant.

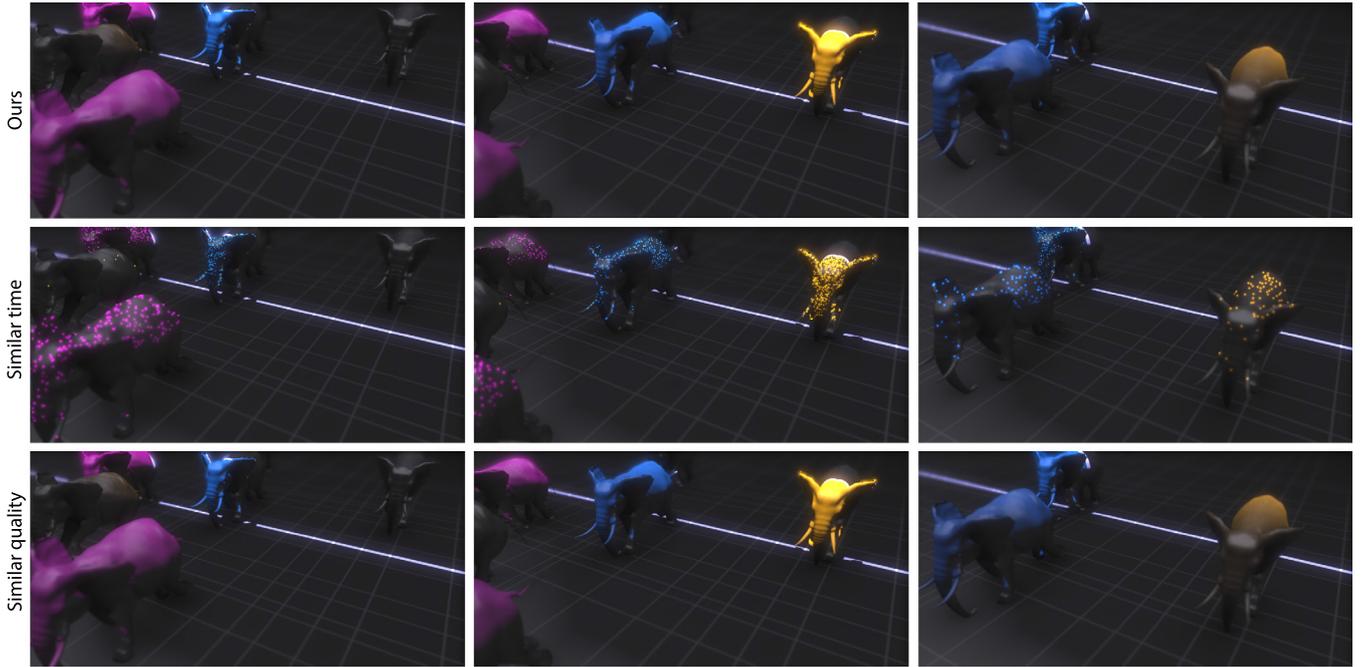


Figure 12: Frames of an image sequence (1024×512 px) computed using our approach (top, 142 fps) and a MC method based on [Gla95], using temporal importance sampling according to the decay function, running in the same time (middle, 1 time sample per pixel) and the same quality (bottom, 1024 tspp.).

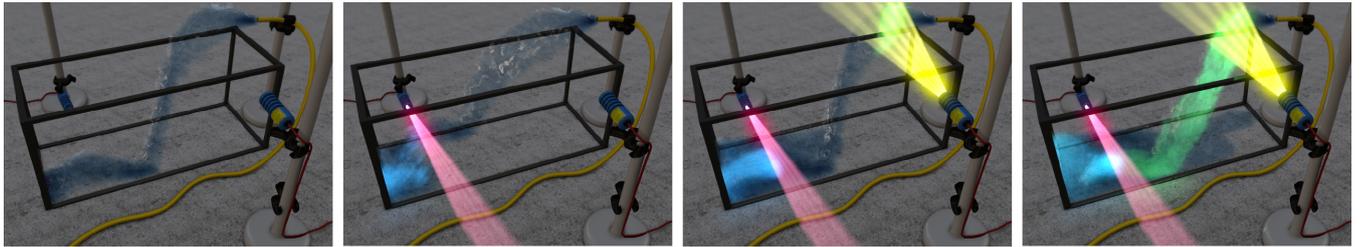


Figure 13: Our model also covers hypothetical liquids with re-radiation depending on the exciting wavelengths. This sequence was rendered using a particle-based simulation, tracking per-particle concentrations r_i . An animated version is part of the supplemental video.

6. Manipulation

Tweaking the parameters of our model to achieve a desired appearance is tedious for novice users. While we provide presets for a selection of typical phosphorescent materials (Tbl. 1), users may wish to create novel or non-physical materials. To this end, we propose to automatically adjust [SDS*93] the parameters of our model to match user-provided constraints. Different from classic inverse lighting, constraints need to be provided in both space and time.

The parameter optimization is performed for each spatial location independently and effects of global light transport of phosphorescent emission are ignored for simplicity. In our manipulation tool, we use a discretization into three channels ($n_\lambda = 3$) and a single concentration ($n_c = 1$) to reduce the number of constraints a user needs to provide. Formally, the constraints are outgoing radiance values $\hat{L}_{o,i}$ at a single spatial location $\hat{\mathbf{x}}$ for n_u different points in

time $\hat{t}_1, \dots, \hat{t}_{n_u}$ (key frames). We now would like to find an ordered set of parameters $\theta = \{\hat{\Lambda}, \hat{k}_e, \hat{k}_r\}$:

$$\theta = \arg \min_{\theta' \in \mathbb{R}^7} f(\theta'), \quad f(\theta') = \sum_{i=1}^{n_u} \|L_o^{\theta'}(\hat{\mathbf{x}}, \hat{t}_i) - \hat{L}_{o,i}\|^2, \quad (7)$$

where $L_o^{\theta'}$ is the outgoing radiance for phosphorescence parameters chosen according to θ' . The parameter set has seven degrees of freedom: two RGB values $\hat{\Lambda}$ and \hat{k}_e and a scalar \hat{k}_r . Note, that this situation is a more general one than for the acquisition as now the data we fit to are provided for an arbitrary illuminant, not the result of measurements under controlled illumination. This entails that we cannot simply apply the procedures as described in Sec. 4. After finding the best parameters θ they are propagated from the spatial positions of the constraints to the rest of the scene using edit

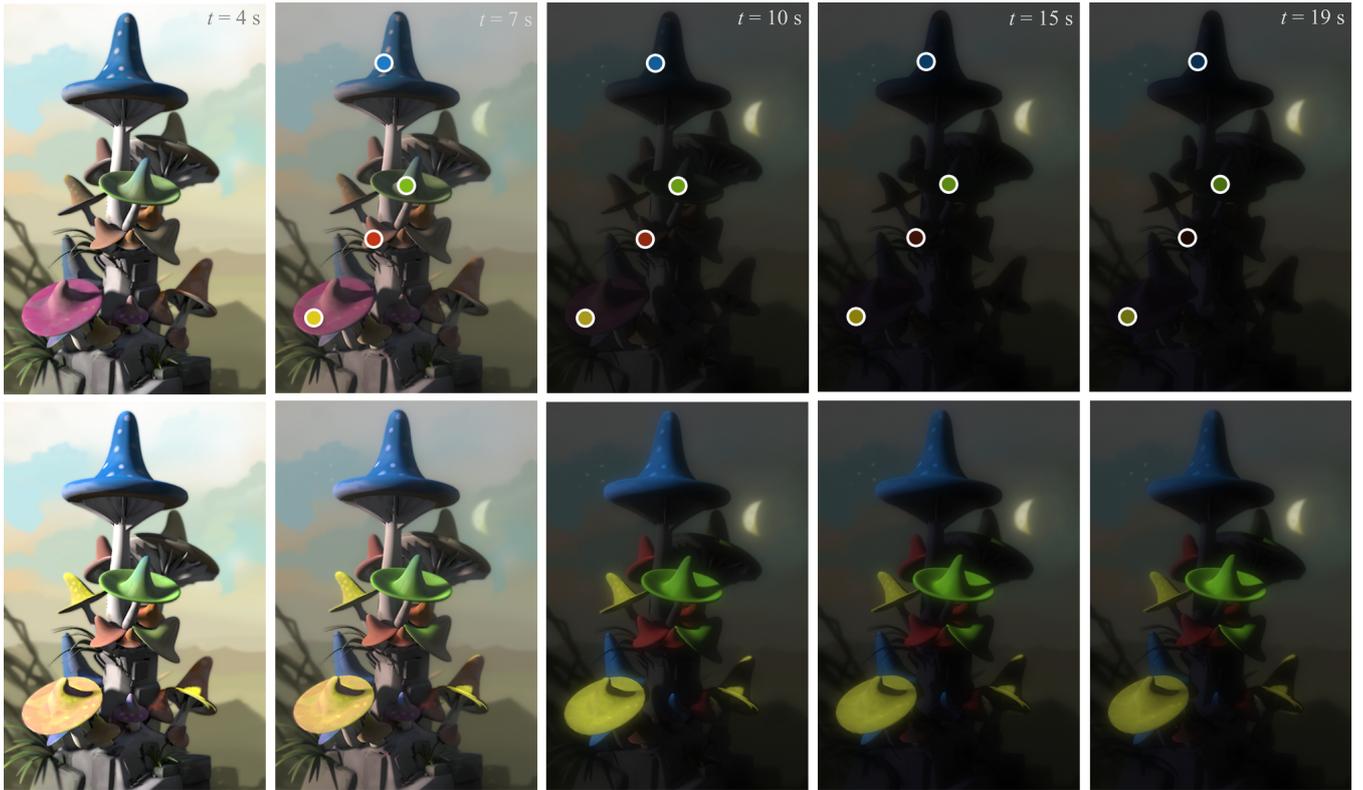


Figure 14: Manipulation: A user puts spatio-temporal appearance constraints (dots, top). The system finds parameter values to best reproduce this appearance while conforming to the model. Rendering the animation again yields an appearance consistent with the constraints (bottom). The manipulation is performed interactively and modifying one constraint with 4 keyframes out of a 200 frame sequence requires 80ms.

propagation [PL07]. Examples of manipulation of phosphorescent appearance are shown in Fig. 14.

Eq. 7 is minimized in three steps. First, the irradiance $\hat{E} = E(\hat{\mathbf{x}})$ is computed and stored for all frames up to the time \hat{t}_{n_u} of the last constraint. Now, for each parameter set θ' the outgoing radiance $L_o^{\theta'}$, consisting of both reflected light and phosphorescent re-emission, is computed using the pre-computed \hat{E} by numerical integration of Eq. 1 up to the last key frame. The cost $f(\theta')$ is evaluated by summing over all key frames. As these computations are independent for all θ' they are well-suited for a parallel implementation on a GPU. Finally, the parameter set with the lowest total cost is selected using parallel minima reduction. Our search range typically covers 10 different values in each dimension which are chosen as $a \cdot 2^b$, where a is a parameter-dependent scale and the exponent b is sampled equidistantly from a parameter-dependent range $[b_{\min}, b_{\max}]$. We chose these ranges to cover the materials we measured. Please see the supplemental for typical search ranges. The time to pre-compute E depends on how light transport is simulated and on the length of the sequence; the optimization itself is performed interactively.

7. Conclusion

We presented a holistic pipeline which, based on a simple model, allows for acquisition of actual materials and an efficient imple-

mentation suitable for interactive manipulation. We believe that phosphorescence, enabling appearance changes in reaction to illumination, could find uses as a special effect or element of storytelling in movies or interactive applications such as games, e.g., to convey the nocturnal appearance of a scene.

Our model is phenomenological and not strictly motivated by physics. While it is physically-plausible, the exact nature of phosphorescence is beyond what might be required in computer graphics. We opt to choose the most simple model that describes the perceivable effects - decay, saturation and specific re-radiation spectrum - and model phosphorescence as the result of changes in multiple concentrations. This allows a good fit to data which we measured and is in agreement with the material science literature [MC97].

Orthogonal extensions include spatially variant acquisition or layered materials [HK93] with phosphorescent layers and anisotropic emission behavior. To achieve this, existing frameworks like the one by Jakob et al. [JdJM14] could be easily adopted. Investigating effects of the thickness of the phosphorescent substances is another direction of future work.

We produced images in a physically-based real-time renderer and a path tracer. Phosphorescence is only perceived at scotopic conditions and reproduction of these has its own challenges [TSF02]. Images need to be processed to best match the scotopic visual ex-

perience. We have complemented the appearance differences by numerical results in Fig. 8 for inspection. A more precise comparison would consider the additional effects of scotopic vision. A formulation as a differential equation also requires to render all frames before frame t in order to actually produce render frame t . We would argue, that this is inevitable except in simplified conditions (e.g., point lights with a step function intensity).

Our manipulation supports time-varying constraints to manipulate discrete materials but is limited to a single exponential that keeps the number of constraints required for a unique solution low. More advanced manipulation should account for the joint effects of global illumination and phosphorescence and use advanced propagation.

Acknowledgements We would like to thank Tobias Bertel and Dushyant Mehta for discussions and help with carrying out measurements. The voice-over features Mathieu Flinders. Furthermore, we thank Bernhard Reinert, Michael Wessely, Klaus Hildebrandt and Tim Weyrich for proof reading and Heike Dobicki for providing us with pieces of equipment. The moon scene in Fig. 10 was contributed by Elena Arabadzhiyska. Two objects of Fig. 1 are CC Zero objects from BlendSwap: the fire extinguisher by user Bastable and the mask by YellowPanda. The running elephant sequence in Fig. 12 is courtesy of Bob Sumner. Finally, we are grateful for the helpful suggestions from people involved in the reviewing process of this paper.

References

- [ACKM16] ALVAREZ-CORTES S., KUNKEL T., MASIA B.: Practical low-cost recovery of spectral power distributions. *Comp. Graph. Forum* 35, 1 (2016), 166–178. 2
- [ATS94] ARVO J., TORRANCE K., SMITS B.: A framework for the analysis of error in global illumination algorithms. In *Proc. SIGGRAPH* (1994), pp. 75–84. 3
- [BHD*08] BENDIG M., HANIKA J., DAMMERTZ H., GOLDSCHMIDT J. C., PETERS M., WEBER M.: Simulation of fluorescent concentrators. In *Proc. IEEE Ray Tracing* (2008), pp. 93–8. 2
- [CHZL11] CHEN B., HAO H.-C., ZHU J., LU M.: A Phenomenological Model for Decay Process of Long-Persistent Phosphorescence. *Chinese Physics Letters* 28, 5 (May 2011). 2, 6
- [Don54] DONALDSON R.: Spectrophotometry of fluorescent pigments. *British J App. Phy.* 5, 6 (1954), 210. 2
- [Gla95] GLASSNER A. S.: A model for fluorescence and phosphorescence. In *Photorealistic Rendering Techniques*. Springer, 1995, pp. 60–70. 2, 9, 10
- [GMAS05] GUTIERREZ D., MUNOZ A., ANSON O., SERON F. J.: Non-linear volume photon mapping. In *Proc. EGSR* (2005), pp. 291–300. 2
- [GSMA08] GUTIERREZ D., SERON F., MUNOZ A., ANSON O.: Visualizing underwater ocean optics. *Comp. Graph. Forum (Proc. EUROGRAPHICS)* 27, 2 (2008), 547–56. 2
- [GTR*06] GU J., TU C.-I., RAMAMOORTHY R., BELHUMEUR P., MATUSIK W., NAYAR S.: Time-varying surface appearance: Acquisition, modeling and rendering. *ACM Trans. Graph (Proc. SIGGRAPH)* 25, 3 (2006), 762–71. 2
- [HHA*10] HULLIN M. B., HANIKA J., AJDIN B., SEIDEL H.-P., KAUTZ J., LENSCH H. P. A.: Acquisition and analysis of bispectral bidirectional reflectance and reradiation distribution functions. *ACM Trans. Graph (Proc. SIGGRAPH)* (2010), 97:1–97:7. 2
- [HK93] HANRAHAN P., KRUEGER W.: Reflection from layered surfaces due to subsurface scattering. In *SIGGRAPH* (1993), pp. 165–74. 2, 4, 11
- [HMI08] HOLOPAINEN S., MANOCHERI F., IKONEN E.: Goniofluorometer for characterization of fluorescent materials. *Appl. Opt.* 47, 6 (Feb 2008), 835–842. 2
- [JdJM14] JAKOB W., D'EON E., JAKOB O., MARSCHNER S.: A comprehensive framework for rendering layered materials. *ACM Trans. Graph.* 33, 4 (July 2014), 118:1–118:14. 2, 4, 11
- [JdP84] JONSCHE A. K., DE POLIGNAC A.: The time dependence of luminescence in solids. *Journal of Physics C: Solid State Physics* 17, 35 (1984). 3, 6
- [KBC*13] KIMMEL B. W., BARANOSKI G. V. G., CHEN T. F., YIM D., MIRANDA E.: Spectral appearance changes induced by light exposure. *ACM Trans. Graph.* 32, 1 (Feb. 2013), 10:1–10:13. 2
- [KHK*12] KIM M. H., HARVEY T. A., KITTLE D. S., RUSHMEIER H., DORSEY J., PRUM R. O., BRADY D. J.: 3d imaging spectroscopy for measuring hyperspectral patterns on solid objects. *ACM Trans. Graph.* 31, 4 (2012). 2
- [KHTI13] KAWAKAMI R., HONGXUN Z., TAN R. T., IKEUCHI K.: Camera spectral sensitivity and white balance estimation from sky images. *International Journal of Computer Vision* (June 2013). 8
- [KO11] KIRK A. G., O'BRIEN J. F.: Perceptually based tone mapping for low-light conditions. *ACM Trans. Graph (Proc. SIGGRAPH)* 30, 4 (2011), 42:1–42:10. 5
- [Lev68] LEVERENZ H. W.: *An introduction to luminescence of solids*. Dover Publications, New York, 1968. 2, 3
- [LSW*01] LEE K., SIEGEL J., WEBB S., LEVEQUE-FORT S., COLE M., JONES R., DOWLING K., LEVER M., FRENCH P.: Application of the stretched exponential function to fluorescence lifetime imaging. *Biophysical J* 81, 3 (2001), 65–1274. 3, 6
- [MC97] MCKEEVER S., CHEN R.: Luminescence models. *Radiation Measurements* 27, 5 (1997), 625–61. 2, 6, 11
- [NJS*11] NOWROUZEZAHRAI D., JOHNSON J., SELLE A., LACEWELL D., KASCHALK M., JAROSZ W.: A programmable system for artistic volumetric lighting. *ACM Trans. Graph. (Proc. SIGGRAPH)* 30, 4 (2011), 29:1–29:8. 2
- [PL07] PELLACINI F., LAWRENCE J.: AppWand: Editing measured materials using appearance-driven optimization. In *ACM Trans. Graph. (Proc. SIGGRAPH)* (2007), vol. 26, p. 54. 11
- [RW45] RANDALL J., WILKINS M.: The phosphorescence of various solids. *Proc. Roy. Soc. London. Math and Phys.* (1945), 347–64. 3, 4
- [SDS*93] SCHOENEMAN C., DORSEY J., SMITS B., ARVO J., GREENBERG D.: Painting with light. In *Proc. SIGGRAPH* (1993), pp. 143–146. 10
- [SPN*14] SCHMIDT T.-W., PELLACINI F., NOWROUZEZAHRAI D., JAROSZ W., DACHSBACHER C.: State of the art in artistic editing of appearance, lighting, and material. In *EUROGRAPHICS 2014 STAR* (2014). 2
- [Sta99] STAM J.: Diffraction shaders. In *Proc. SIGGRAPH* (1999), pp. 101–110. 2
- [SWD05] SHARMA G., WU W., DALAL E. N.: The CIEDE2000 color-difference formula. *Color Res. & App.* 30, 1 (2005), 21–30. 7
- [TSF02] THOMPSON W. B., SHIRLEY P., FERWERDA J. A.: A spatial post-processing algorithm for images of night scenes. *J Graphics Tools* 7, 1 (2002), 1–12. 9, 11
- [WTP01] WILKIE A., TOBLER R. F., PURGATHOFER W.: Combined rendering of polarization and fluorescence effects. In *Proc. EGSR* (2001), pp. 197–204. 2
- [WWLP06] WILKIE A., WEIDLICH A., LARBOULETTE C., PURGATHOFER W.: A reflectance model for diffuse fluorescent surfaces. In *Proc. GRAPHITE* (New York, NY, USA, 2006), pp. 321–31. 2