




Navigating the Manifold of Translucent Appearance

Dario Lanza , Belen Masia , and Adrian Jarabo 

Universidad de Zaragoza - I3A, Spain




Figure 1: We propose a continuous, perceptually-motivated manifold for translucent appearance (illustrated on the right of the right panel, color-coding is for visualization purposes only). The manifold covers a wide range of materials, as shown by the objects on the right panel, whose appearance has been selected from the corresponding points selected in the manifold. We also propose an interactive editing interface that leverages this manifold, and allows users to navigate the space of translucent appearance more effectively than when using traditional slider-based interfaces. We evaluate this by means of a comprehensive user study including two different editing tasks. An example of the results obtained in this study with the two interfaces (ours and a slider-based one) for a certain target image is shown on the left panel. The user study shows that participants were able to better match the target image using our interface, with higher level of satisfaction.

Abstract

We present a perceptually-motivated manifold for translucent appearance, designed for intuitive editing of translucent materials by navigating through the manifold. Classic tools for editing translucent appearance, based on the use of sliders to tune a number of parameters, are challenging for non-expert users: These parameters have a highly non-linear effect on appearance, and exhibit complex interplay and similarity relations between them. Instead, we pose editing as a navigation task in a low-dimensional space of appearances, which abstracts the user from the underlying optical parameters. To achieve this, we build a low-dimensional continuous manifold of translucent appearance that correlates with how humans perceive this type of materials. We first analyze the correlation of different distance metrics in image space with human perception. We select the best-performing metric to build a low-dimensional manifold, which can be used to navigate the space of translucent appearance. To evaluate the validity of our proposed manifold within its intended application scenario, we build an editing interface that leverages the manifold, and relies on image navigation plus a fine-tuning step to edit appearance. We compare our intuitive interface to a traditional, slider-based one in a user study, demonstrating its effectiveness and superior performance when editing translucent objects.

Keywords: material appearance, translucency, perception

CCS Concepts

• **Computing methodologies** → Perception; Graphics systems and interfaces; Appearance and texture representations;

1. Introduction

Materials that exhibit some degree of translucency are ubiquitous in the real world, ranging from organic materials, such as milk or wax, to inorganic materials like jade or glass. Light transport algorithms capable of simulating them with great accuracy, and in a computationally efficient manner, are well-developed [NGHJ18]. When it comes to *editing* these material models, however, existing techniques are considerably less developed.

Translucent materials are typically modeled using the bidirectional scattering surface reflectance distribution function (BSSRDF), and manipulating their appearance through tuning of the optical parameters of the BSSRDF is a daunting challenge. These parameters span a high-dimensional space, with dimensions that correlate poorly with human perception, and complex, non-linear interactions between them, leading to distant optical parameters resulting in similar appearances [ZRB14]. This problem is aggravated in the case of novice users, who are not acquainted with the individual impact of each of these parameters on the final appearance, not to mention the intricate manner in which they interact.

Bridging the gap between physically-based parameters, used in analytical or measured models of appearance, and high-level attributes that humans can better understand and control is a long-standing problem in the field of material appearance modeling. Typically, this involves finding low-dimensional perceptual manifolds, which enable a more intuitive navigation of the targeted range of appearance [SGM*16; TGG*20; SWSR21], and in some cases even provide a relationship between the dimensions of these manifolds and the parameters of reflectance models [PFG00; WAKB09]. Unfortunately, all these approaches to building perceptually-based material appearance manifolds target opaque BRDFs.

Despite –or owing to– their higher complexity, translucent appearance models have received less attention in this area. Most efforts have focused on optically-thick materials, where editing can be done by manipulating the diffusion profile [WCW*08; STPP09; Bur15], thus limiting material editing to a particular set of translucent materials. In addition, the perception of translucency, and its interaction with factors like lighting and geometry, is not fully understood [FJB04; FB05; GWA*15; GTPH21; KNYK22; LJM22], further hindering the task.

Our goal is to explore intuitive methods for navigating and editing translucent appearance, allowing users to abstract themselves from the low-level optical properties defining appearance and focus only on the appearance itself. We do this by building a perceptually-meaningful continuous manifold of translucent appearance (Figure 1). This manifold is built under the premise that perceptually similar appearances should be closer together, so global exploration of our manifold will allow drastic appearance changes, and local navigation will result in fine-tuning of appearance. We focus on homogeneous, achromatic translucent appearance, with the goal of evaluating whether an editing paradigm based on image navigation of such a manifold can be practical for users and outperform more traditional, slider-based editing approaches, thus justifying further exploration of its applicability to the vast space of translucent appearance.

We build our manifold leveraging an objective distance measure, which we select by conducting a perceptual study. Since our ultimate objective is appearance editing, we validate the usefulness of such space by proposing an interface that builds on top of it; we evaluate this interface with a user study, comparing it to a standard, slider-based interface. We show that novice users have better objective performance in editing tasks using our prototype interface, and at the same time have a more satisfactory editing experience.

Specifically, our contributions are:

- A perceptual study to determine an adequate objective distance measure for translucent appearance that correlates with human perception of translucency over a wide range of optical parameters.
- A perceptually-meaningful continuous manifold of the space of translucent appearance, suitable for editing homogeneous translucent materials.
- An editing interface that leverages this manifold, and which we validate through a user study, showing its effectiveness and superior performance in comparison to a standard, slider-based approach.

2. Background and Related Work

In this section, we first briefly introduce the terminology used in this paper for modeling translucent appearance, and then summarize the related work on translucency editing, perceptual spaces of appearance, and material design and editing interfaces.

2.1. Light Transport in Translucent Materials

Translucent materials can be modeled as a dielectric material embedding a scattering medium. When a light beam interacts with the dielectric boundary, it is both reflected and transmitted inside the object following the Fresnel equations. Light transmitted inside the medium is then scattered multiple times until it re-emerges back to the surface, potentially far away from the entrance point.

Light scattering at the boundary is determined by the change of index of refraction of the medium, η , and the microscopic geometry modeled statistically via the microfacet model [WMLT07], where the microgeometry is represented by a statistical distribution of microscopic smooth planes, parametrized by the surface roughness ρ .

Inside the medium, light interacts with it following the radiative transfer equation [Cha60]:

$$\omega \cdot \nabla L(\mathbf{x}, \omega) = -\sigma_t L(\mathbf{x}, \omega) + \alpha \sigma_t \int_{\mathcal{S}^2} f_p(\mu) L(\mathbf{x}, \omega') d\omega', \quad (1)$$

where \mathbf{x} is a point inside the medium, $\omega \in \mathcal{S}^2$ is a direction, and $L(\mathbf{x}, \omega)$ is the radiance at \mathbf{x} in direction ω . The medium is characterized by the extinction coefficient σ_t , the single scattering albedo α , and the phase function describing the angular scattering $f_p(\mu)$, with $\mu = \omega \cdot \omega'$. We restrict our work to homogeneous achromatic media, avoiding spatial and spectral dependence of these parameters. For the phase function we use the common Henyey-Greenstein model [HG41], parameterized by the directionality parameter g .

2.2. Editing Translucency

While there is vast literature studying the perception of translucency—for a comprehensive review we refer the reader to the work of Gigilashvili et al. [GTHP21]—, editing of translucent materials in an intuitive way is a relatively unexplored topic. Most existing works have focused on optically thick materials, where the light is at the diffusion regime: In this case, editing is generally done via modifying the diffusion profile [XGL*07; STPP09; JZJ*15; Bur15], or via local manipulation of the single scattering albedo in volumes [HR13]. For a broader range of optical thickness, the general approach falls back to manipulating optical parameters via sliders, either directly or via simplified controls [WVH17]. Our work targets all ranges of optical thickness and poses editing as a low-dimensional intuitive navigation through the potential range of appearance, rather than low-level properties. Inverse rendering of participating media has proved a useful tool for manipulating translucent appearance. This includes optically thin [NGD*06] and thick materials [JMLH01; DLW*22], as well as more general materials using fully-differentiable volume path tracing [GZB*13; GLZ16]. Recently, neural-based approaches [CLZ*20; LNN23] have demonstrated good performance on inverting the BSSRDF. While we leverage inverse rendering to obtain an estimate of material properties, our focus is on providing an improved method to intuitively edit translucency.

2.3. Perceptual Spaces for Appearance Editing

A substantial body of work has been devoted to deriving low-dimensional perceptual manifolds of material appearance in which materials are organized along meaningful, interpretable dimensions. Often, a key motivation behind these is to facilitate editing and control of appearance, shielding the user from the high-dimensional, non-linear spaces spanned by the optical or physically-based parameters of BSDF models.

A seminal work in this area is that of Pellacini et al. [PFG00], where they derive a two-dimensional perceptual space for gloss. The space is built from psychophysical data, making use of multidimensional scaling, and they establish a relationship between their two main perceptual dimensions. A close goal is pursued by Wills et al. [WAKB09], who build a two-dimensional perceptual space for gloss for measured BRDFs. More recently, Toscani et al. [TGG*20] used a similar methodology with a wider range of appearance, and with the addition of an experiment to interpret and label the three main dimensions of their derived space.

Instead of focusing on a single attribute (e.g., gloss) and deriving a perceptual model for it, other works have broadened the scope, often at the same time lifting the restriction of having a perceptually-linear space, where a linear change in space corresponds to a linear change in appearance, and settling for *intuitive*, perceptually-meaningful spaces, in which items are spatially sorted following perceptual criteria, but without any extra relationship between space in the manifold and perceptual space. Matusik [Mat03] evaluated different dimensionality reduction techniques for measured BRDF data, and defined a set of 15 attributes (which they call perceptual traits) which are then used to navigate the space of appearance. Inspired by this, Serrano et al. [SGM*16] generated intuitive manifolds for a carefully-curated list of attributes, building on

the dimensionality reduction proposed by Nielsen et al. [NJR15]. Shi et al. [SWSR21] follow a similar approach, but do not establish specific attributes a priori, and instead extract dimensions that explain perceived differences in appearance. They also propose an editing interface, discussed in Section 2.4.

Akin to these works, we seek to build a low-dimensional space that is linked to the optical properties of the material appearance model in a manner that allows for appearance editing. We draw inspiration from their methods, but while their focus is on opaque BRDFs, we deal with translucent materials, a much less explored domain.

In the context of translucent materials, the closest work to ours, which identifies a two-dimensional perceptually-meaningful manifold for a set of physical scattering parameters, is that of Gkioulekas et al. [GXZ*13]. They specifically explore the impact of the phase function on translucent appearance, and relate the two axes of their space to moments of the phase function. We build on their findings, including the distance metric they use, but we tackle a wider range of appearance, and focus on providing an intuitive space for editing it. Recently, Liao et al. [LSX22] proposed an unsupervised learning-based model that identifies perceptually relevant dimensions for translucent material appearance from images, and show that manipulations in the latent space of the model can lead to modifying the appearance of the object. While they work in image space, we seek to retain the mapping to optical parameters of the material, required for physically-based rendering.

2.4. Material Design Interfaces

Different interfaces and interaction paradigms have been proposed in the context of material appearance design and editing. They mainly rely on one of two paradigms: those based on a set of sliders that allows modifying a series of parameters, or those based on navigation of a structured collection of images. Both were analyzed by Kerr and Pellacini [KP10] in their comparison of material design interface paradigms for novice users. A third category relies on directly painting with brush tools, and has been used for opaque BRDFs, particularly highlights [CPK06; PGSP08], but requiring custom material models.

Slider-based models can rely on *physical* sliders, that expose and control the optical properties of the underlying BSDF model (as used by many commercial, off-the-shelf modeling tools [Aut23]), or on *perceptual* sliders, allowing editing along the dimensions of perceptually-meaningful spaces of appearance like the ones introduced in Section 2.3. Some BSDF models, like Disney's Principled BSDF [Bur12] are already designed to provide an optimal combination of controllability, expressiveness, and robustness.

Hybrid approaches that combine image navigation with subsequent tuning of parameters are commonly found in everyday applications such as Adobe Filter gallery [Ado23] or Microsoft PowerPoint template gallery [Mic23]. They offer good balance between a rapid selection of the desired result, and the fine-tuning precision given by sliders. In the context of material appearance, the recent work by Shi et al. [SWSR21] proposed to use such a hybrid paradigm for authoring opaque metallic-like BRDFs. We draw inspiration from this approach, and explore the use of such hybrid

editing interfaces for translucency, which is a higher-dimensional problem with a more complex interaction between the low-level parameters used to define appearance.

3. A Distance Measure for Translucent Appearance

Measuring the perceptual distance between the appearance of two materials is an open problem, and multiple measures have been proposed [LBFS21]. These measures can operate in material space or in image space. The former have the obvious advantage that they do not require rendering the images, but they have shown to correlate poorly with perception, especially when dealing with suprathreshold differences [NDM06; SGM*16]. Image-based metrics allow to factor in the influence of the geometry and the illumination, critical in material appearance [LMS*19], and in particular in the case of translucency [GWA*15; XWG*14]. In this section, we briefly describe our choice of distance measure: we curate a set of metrics based on the literature (Section 3.1), and then conduct a perceptual study to select the best one for our particular type of stimuli (Section 3.2). For a more detailed description of the metrics and the experiment, we refer the reader to the supplemental material.

3.1. Image-based Metrics

While the literature in the area of image-based appearance metrics is vast, the specific case of translucent materials is not as well studied as that of opaque BRDFs. Studies range from analyzing existing metrics proposed for BRDFs for their use in translucent stimuli, to the proposal of models specifically designed to predict human responses to those translucent stimuli. We explore the use of representative metrics of both, as well as other commonly used image-based metrics. In particular, we evaluate: the *cubic root metric* [NDM06], which Gkioulekas et al. [GXZ*13] found to correlate well with the perception of translucency; the *anisoshading ratio* proposed by Kiyokawa [KNYK22] for translucency; and a set of image-based metrics including the well-known *MS-SSIM*, widely used for image comparison, the recent *FovVideoVDP* [MCR*21], and the learning-based metric *LPIPS* [ZIE*18]. Details on these metrics can be found in the supplemental material.

3.2. Perceptual Study

We seek to find, for the type of stimuli that we work with, which metric correlates best with human judgements. As mentioned, previous work has analyzed the performance of different image-based metrics on the perception of translucency [GXZ*13]. However, they did so with a set of stimuli that focused on studying the role of the phase function. Since their focus is on densely sampling the space of phase functions, they fix the other scattering parameters to isolate the effect of the former (σ_a and σ_t take the values of the green channel of marble [JMLH01], with perturbations of them around these default values; roughness is also fixed), leading to a more reduced range of translucent appearance (in terms of σ_a , σ_t and roughness coverage) than the one we cover in the present work. Therefore, we build on their work, and assess through a perceptual study whether their proposal of the cubic root metric applies to a larger range of appearance, even when compared to more recent and sophisticated image-based metrics.

3.2.1. Stimuli

Our stimuli are images of a translucent object against a uniform background; we use a fixed geometry, and vary the material's optical parameters and the lighting conditions. Images are rendered with Mitsuba 0.6 [Jak10], using the default volumetric path tracer with no limit to the number of bounces. We fix the geometry to the *Lucy* statue, since it has both thick and thin parts and shows a balance of medium and high-frequency details, which are helpful cues [GTPH21]. We add a background horizontal plane, to project caustic patterns that are also used to assess translucency [GDHP20].

Since the perception of a translucent material is heavily affected by where the light is positioned with respect to the object, we render our stimuli with three different variations in lighting: *front*, *side* and *back* [XWG*14]. In all three cases, we use a strong area light to generate a directional light source from the specific direction, and add a filling environment map (*Ennis* [ICT10]).

We model translucent materials as a medium enclosed by a surface, yielding five optical parameters that determine material properties (see Section 2.1). We fix the index of refraction (IOR) to $\eta = 1.5$, since most translucent materials have an IOR close to this number [GTPH21]. We explore the whole range of phase function asymmetry $g \in [-1, 1]$ and single scattering albedo $\alpha \in [0, 1]$, while we bound both surface roughness and extinction to $\rho \in [0.01, 0.6]$ and $\sigma_t \in [0.01, 10000]$, respectively. Given that the parameter space $\pi = [\rho, \sigma_t, \alpha, g]$ is four-dimensional, a grid-based sampling would be impractical; instead, we sample the space using a low-discrepancy quasi-random distribution (Halton). The use of a low-discrepancy sampler allows to have a uniform coverage of the sampling space and scales well in multidimensional spaces. Even within our scheme, the optical parameters sampled have a highly non-linear relationship with translucent appearance. Based on findings from previous work on translucency perception [GUT*22], we choose a power-like distribution for the single scattering albedo α and extinction σ_t and empirically set the parameters of the power function to provide a balanced coverage of appearances. Specifically, we follow $\alpha = \xi^5$ and $\sigma_t = \xi^{-5}$, with ξ a random number, skewing the sampling towards lower extinction values and including optically thinner appearances that would otherwise be under-represented.

3.2.2. Procedure

We avoid ranking methods and resort to a two-alternative forced choice (2AFC) experimental paradigm. This is common practice in complex appearance evaluation tasks, where the measured effect may not follow a linear scale. As in standard 2AFC experiments, the participants need to select which of the two *candidate* images is most similar to a *reference* one.

We sample sixteen sets of material properties, under three different lighting conditions. This yields a total of 5,040 triplets, each evaluated by at least five different participants, leading to a total of 25,200 trials. Given the large number of trials, we use Amazon Mechanical Turk, implementing control trials to discard unreliable participants (slightly under 25% of participants were discarded).

Table 1: Accuracy (both majority and raw) of the various image-based metrics when predicting human similarity judgements, for the three lighting conditions: side, back, and front. Performance is generally consistent across lighting conditions. The cubic root metric outperforms the rest, and is thus our distance metric of choice.

Metric	Majority			Raw		
	Side	Back	Front	Side	Back	Front
Cubic root metric	0.77	0.81	0.84	0.70	0.73	0.75
Anisoshading ratio	0.61	0.65	0.69	0.50	0.50	0.63
MS-SSIM	0.29	0.27	0.25	0.22	0.19	0.18
FovVideoVDP	0.29	0.27	0.25	0.23	0.20	0.17
LPIPS	0.75	0.77	0.77	0.69	0.71	0.72

3.2.3. Results

Following previous works [GAGH14; WAKB09; LMS*19], we compute the *accuracy* of each image-based metric when predicting human responses using both its *raw* and *majority* modalities. The raw accuracy considers each response as a correct answer for each triplet, while the majority accuracy considers that only the majority’s response for each triplet is correct. Differences between both modalities exist when there is low agreement between participant responses for a given triplet. Table 1 shows the accuracy results for each of the five metrics when predicting human responses, in both its *raw* and *majority* modalities.

Image quality metrics do not provide a good prediction in this scenario, even in the case of those designed to detect supra-threshold differences. On the other hand, we see how the cubic root metric initially developed as a metric for BRDF comparison [NDM06] offers the best results, consistent with what Gkioulekas et al. [GXZ*13] found in their exploration of the phase function. These are complementary findings to the work of Gkioulekas et al., since we span a wider range of optical parameters and appearance, going from crystal-like to darker materials (see supplemental material for representative subset of the materials used). Therefore, the cubic root metric, initially proposed for BRDFs, will be our distance metric of choice.

4. A Perceptually Meaningful Space for Translucent Appearance

Editing material appearance is generally done by moving along non-orthogonal non-linear dimensions via sliders. In our case, that means a four-dimensional space, which is a challenging task for novice (or even experienced) users.

Instead, our goal is to provide a more natural navigation along the manifold of translucent appearance. That imposes a set of design constraints for our manifold: 1) The manifold should be perceptually meaningful, so that similar appearances are close together; 2) the manifold should be continuous, to allow navigation along the space of possible translucent appearances; and 3) the manifold should be two-dimensional, so that it can be easily navigated by the user.

Mathematically, we want to find a two-dimensional manifold $\mathcal{M} \subseteq \mathcal{R}^2$ that is a convex hull of appearances. Each point $\mathbf{y}_i \in \mathcal{M}$

maps to a set of optical parameters $\boldsymbol{\pi}_i$, thus establishing a relationship $m(\mathbf{y}_i) = \boldsymbol{\pi}_i$. For that manifold to be perceptually meaningful, we impose that for a pair of points $\mathbf{y}_i, \mathbf{y}_j \in \mathcal{M}$, it holds that

$$|\mathbf{y}_i - \mathbf{y}_j| \propto d(\mathbf{i}_i, \mathbf{i}_j), \quad (2)$$

with $d(\cdot, \cdot)$ being the cubic root metric, and $\mathbf{i}_i = \text{Render}(\boldsymbol{\pi}_i)$ the generated appearance with optical parameters $m(\mathbf{y}_i) = \boldsymbol{\pi}_i$. Note that as we describe later, \mathbf{i}_i is not necessarily a single image; in fact, in the construction of our manifold and its navigation, we use a triplet of three images with the same optical parameters and under different illumination conditions. In the following, we describe our manifold construction from a discrete set of samples, which we navigate by using an interpolation scheme.

4.1. Manifold Construction

In order to build our manifold, we leverage dimensionality reduction techniques. This allows us to leverage the cubic root metric, which we showed that correlates best with human perception in our scenario (Section 3), so that materials with similar appearance are closer together, and different materials are pushed further away. A benefit of this approach is that similarity relations between optical parameters are implicitly handled, as well as non-linearities between optical parameters and resulting appearance.

We sample $M = 100$ tuples of optical parameters $\hat{\boldsymbol{\pi}}$ using the sampling procedure described in Section 3.2, and render the *Lucy* statue under three lighting conditions (front, back and side). Then we compute a $M \times M$ distance matrix relating the perceptual distance between each pair of appearances, using the average distance of all three lighting conditions. By associating each point to three different lighting conditions we enforce the manifold to place closer together points that carry optical parameters that generate similar appearances under different lighting conditions, making the manifold more robust.

Finally, we feed the IsoMap algorithm [TdL00] with the distance matrix, and set the number of neighbors to five and the dimensionality of the output space to two. With these settings, we build a two-dimensional manifold where each image triplet $\hat{\mathbf{i}}_i$ generated with optical parameters $\hat{\boldsymbol{\pi}}_i$ has a coordinate $\hat{\mathbf{y}}_i \in \mathcal{M}$ in the manifold.

Discussion While other algorithms for dimensionality reduction exist (e.g., multidimensional scaling or kernel principal component analysis) we chose the IsoMap algorithm since it works well with arbitrary distance matrices and handles well non-linear spaces [TdL00]. Other alternatives using neural networks, following previous works on perceptual spaces for material appearance [LMS*19; DLC*22; SL23], were considered, but we found that IsoMap works well in our case, and produces stable and meaningful manifolds that can be easily navigated, while the latent spaces of neural-based methods might lead to unintuitive spaces and would require re-training with extensive labeled data for translucent appearance.

4.2. Manifold Navigation

The manifold defined above provides a convex hull of appearances for translucent materials. However, the discrete set of samples it

is built with fill the 2D manifold unevenly, and do not allow for smooth navigation. To build a continuous manifold inside the convex hull of the manifold from a discrete set of samples, we follow an approach similar to Wills et al. [WAKB09], and estimate the optical parameters π_i from the finite set $\overline{\mathcal{M}}$ of data points $(\bar{\mathbf{i}}_j, \bar{\pi}_j, \bar{\mathbf{y}}_j)$ used for building the manifold.

Computing π_i directly from the optical parameters associated with the finite set of data points $\bar{\pi}_j$ does not necessarily result into a perceptually meaningful appearance manifold. The reason is the complex interactions between optical parameters and final appearance, arising due to the non-linear effect of each different parameter, as well as appearance metamerism due to well-known similarity relations between optical parameters [ZRB14], where very different optical parameters might result in a similar appearance. Instead, we compute the interpolation between datapoints in appearance space (i.e., in image space), and compute \mathbf{i}_i as

$$\mathbf{i}_i = \frac{1}{K_N} \sum_{j=1}^N K(|\mathbf{y}_i - \bar{\mathbf{y}}_j|) \bar{\mathbf{i}}_j, \quad (3)$$

with $\bar{\mathbf{y}}_j$ and $\bar{\mathbf{i}}_j$ the position and appearance of the N nearest data points to \mathbf{y}_i (we set $N = 5$), $K(\cdot)$ a radially symmetric kernel (we use a triangle kernel), and $K_N = \sum_{j=1}^N K(|\mathbf{y}_i - \bar{\mathbf{y}}_j|)$ the normalization factor. This approach has two benefits: First, it makes navigation inside the manifold smooth and predictive. Second, it imposes that the relationship in Equation 2 holds for the generated appearance \mathbf{i}_i at \mathbf{y}_i and the rest of the manifold, assuming a dense enough set of datapoints generating a manifold. Figure 2 shows the manifold for the back-lit condition.

Unfortunately, Equation 3 results into a relatively simple image-space interpolation, with no direct relationship between \mathbf{y}_i and π_i . Instead, for building the map between \mathbf{y}_i and π_i we need to compute \mathbf{i}_i , and from that compute $\pi_i = \text{Render}^{-1}(\mathbf{i}_i)$, as we describe in the next subsection.

Discussion While this interpolation has some similarities to the one performed by Wills et al. [WAKB09], note that they interpolate the optical parameters in the parameter space, which is only feasible in the low-dimensional space of gloss in opaque BRDFs. In the case of translucent materials, this approach leads to suboptimal results, and causes non-smooth transitions inside the manifold. Additionally, while the interpolation in Equation 3 does not guarantee that there is a set of parameters π_i that would generate the appearance in \mathbf{i}_i (i.e., \mathbf{i}_i might not be physically meaningful), we found that there is always a close match between π_i and \mathbf{i}_i . Both cases are shown in Section 4.4.1 and Figure 3.

4.3. Optical Parameters Retrieval

We leverage physically-based inverse rendering for computing the inverse problem $\pi_i = \text{Render}^{-1}(\mathbf{i}_i)$, by solving

$$\min_{\pi_i} d(\mathbf{i}_i, \text{Render}(\pi_i)), \quad (4)$$

with $d(\cdot, \cdot)$ our distance metric (cubic root metric), averaged over the three different illuminations. We solve for π_i using the Adam optimizer [KB14] (learning rate $lr = 0.025$, 100 iterations), with

gradients computed using the physically-based differentiable rendering in Mitsuba 3 [JSR*22]. We initialize the optimization by interpolating the optical parameters following

$$\pi_{i,0} = \frac{1}{K_N} \sum_{j=1}^N K(|\mathbf{y}_i - \bar{\mathbf{y}}_j|) \bar{\pi}_j, \quad (5)$$

which we empirically found works well in our case. To make our problem tractable, we inverse rendered images with a resolution of 128×128 pixels, as it was a good balance between the quality of results and the time to convergence. In our tests, we found that increasing the resolution for inverse rendering did not have a major impact on the reconstruction error, but resulted in significantly longer reconstruction times. In particular, we found that performing the optimization for a native resolution (512×512) and the lower-resolution one (128×128) provided similar error values in both the recovered optical parameters (MAE of 0.16 vs 0.14, respectively) and error in image space (cubic root metric of 0.037 vs 0.048, respectively). We further analyze the quality of the inverse rendering stage in Section 4.4.

4.4. Analysis of the Space

The resulting two-dimensional manifold can be seen in Figure 2. Roughly, the manifold shows two main axes of appearance: From top-left to bottom-right, we can see that the appearances are distributed from glass-like (top-left) to a more opaque diffuse-like appearance (bottom-right) similar to, e.g., marble or jade, which agrees with the findings by Gkioulekas et al. [GXZ*13]. On the other hand, the horizontal axis distributes appearances from dark (left) to bright (right). Since the main purpose of our generated manifold is to enable navigation of the space of translucent appearance, we evaluate this through a user study by integrating our manifold into an appearance editing interface in Section 5. Prior to this, in the remainder of this section, we analyze the resulting manifold in terms of *validity*, *stability* and *generality*. Finally, we discuss the relationship of this manifold with the underlying optical parameters.

4.4.1. Validity

To assess our approach to build the continuous manifold (Sections 4.2 and 4.3), we compare it to the case in which the manifold is generated by simply interpolating the optical parameters from the initial set of samples. Our measure of error is the cubic root metric, and results are shown in Figure 3. The error using inverse rendering is lower, and also shows that the image-based interpolation results in physically plausible appearance.

4.4.2. Stability

To analyze the dependence of our manifold on individual material samples, we perform a stability analysis [WAKB09]. We do this by building M manifolds following a leave-one-out procedure, with $M - 1$ manifold samples being used to build each manifold, leaving one material sample out each time a manifold is built (manifold construction is as described in Section 4.1). For each of these M manifolds, we compute its Procrustes distance [Ken89; Boo97] to our resulting, final manifold. This is done by computing the




Figure 2: Depiction of our continuous manifold for the Lucy geometry under the back lighting condition. The convex hull of the manifold is depicted by the dark-red line. Note that the manifold smoothly interpolates between different translucent appearances, naturally following two main high-level dimensions of translucency: glass-likeness (bottom-right to top-left) as described by Gkioulekas et al. [GXZ*13], and brightness (left to right). These axes are shown in the bottom-right inset. We refer to the supplemental material for the depiction of the manifold under other lighting conditions.




Figure 3: Error (cubic root metric) between the interpolated image, \mathbf{i}_i , and the rendered one, $\text{Render}(\boldsymbol{\pi}_i)$, for a set of uniformly distributed points along the manifold. We show the error when optical parameters $\boldsymbol{\pi}_i$ are computed through parameter interpolation (left), and by means of inverse rendering (right).

Euclidean transformation (rotation, reflection, scaling, and translation) between the two embeddings that minimizes the distance between them, i.e., the sum of squared differences between the points in both manifolds. The average normalized Procrustes distance across our M manifolds is $d_P = 0.017$; with distances in the range $[0, 1]$, this value is indicative of a stable manifold.

4.4.3. Generality

Since we build our manifold based on a single geometry (*Lucy*), we would like to assess whether such manifold remains valid for other geometries. For that, we compare our manifold with the manifolds generated for a set of additional geometries that show different ranges of spatial frequencies and distributions of thickness, which are relevant geometric features for both reflection and transmittance. These additional manifolds are built for the *Armadillo* and *Bunny* from the Stanford Repository, and a *sphere* and a *cube*. All these geometries were re-scaled to have a thickness similar to the pedestal in the *Lucy* statue, so that the level of translucency is comparable. The manifold for each of these geometries is generated using the same procedure as with *Lucy*, described in Section 4.1.

We analyze the generality of our approach by performing pairwise comparisons between the manifolds generated for all five different geometries, analyzing the correlation of the distance matrices

Table 2: Pairwise correlation between manifolds generated for different geometries and the Lucy statue, computed using the Mantel test. Values correspond to the Pearson product-moment correlation coefficient, which ranges from -1 to 1 , with 1 indicating a strong positive correlation and -1 a strong negative one. Results shown are statistically significant ($p < 0.05$).

	Armadillo	Bunny	Sphere	Cube
Lucy	0.96	0.94	0.88	0.82

using the non-parametric Mantel test [Man67]. Results are shown in Table 2, with 1.0 being perfect correlation: Even for very different geometries all manifolds exhibit great consistency. As expected, simple geometries (i.e., sphere and cube) show smaller correlation, in line with previous work [GSW*21] that noted that simple geometries are perceived differently from more complex geometries. However, as this complexity increases, the manifolds become extremely consistent, even for geometries with very different distributions of thickness and spatial frequency (e.g., *Bunny* vs *Lucy*).

4.4.4. Underlying Optical Parameters

Here we analyze the relationship between the underlying optical parameters and the perceptual manifold. These parameters are retrieved from the continuous manifold following the procedure described in Section 4.3. Given the complex interplay of optical parameters resulting in strong similarity relationships between phase function, albedo, and extinction [GXZ*13; ZRB14], we opt for analyzing the reduced coefficients (i.e., first-order similarity relationships [WPW89]). Figure 4 shows the distribution of the reduced extinction σ_t^* and reduced single scattering albedo α^* inside the manifold: As we can observe, there is a trend that relates the reduced extinction and single scattering albedo with the glass-like appearance (diagonal axis, from top-left to bottom-right), so that lower σ_t^* and α^* increases the glass-like look, and higher values increase the diffuseness, as expected. The brightness (horizontal axis




Figure 4: Visualization of the reduced optical parameters in our continuous manifold, obtained via inverse rendering. We show the reduced extinction coefficient σ_t^* (left, log scale) and reduced single scattering albedo α^* (right, linear scale), computed using first-order similarity. Please refer to the main text for an analysis of the observed trends.

of the manifold), on the other hand, is mostly related with the single scattering albedo. While these trends seem consistent, further exploration would be required to find a complete description of the high-level perceptual axes within the manifold, and their potential high-order relationships with the underlying optical parameters.

5. An Interface for Editing Translucent Appearance

This section seeks to assess whether our perceptually meaningful space can indeed be effectively used to edit translucent appearance. To do that, we build a prototype editing interface that leverages our manifold, and compare it to a standard editing method (based on tuning sliders for each of the optical parameters) in a user study comprising two editing tasks on a diverse set of materials and geometries.

5.1. Interface Overview

We build a prototype material editing interface that leverages our proposed translucency manifold. Our interface is integrated into the widespread, open-source modeling software Blender [Ble23]. The main goal of the manifold is to enable users to rapidly obtain a reasonable approximation to the desired final appearance. This approximation can be very hard to reach through individual tuning of four (or more) optical parameters that interact in complex ways. At the same time, it has been shown that sliders perform better than image navigation for *precise adjustments* in matching tasks [KP10].

Taking all of this into account, we propose a hybrid interface that combines image navigation through a zoomable version of our manifold, and subsequent fine-tuning of the optical parameters using a slider-based interface. Image navigation facilitates an initial good approximation to the desired appearance (which is hard to obtain with sliders alone), and fine-tuning through sliders enables refining this approximation to obtain an optimal match.

Our interface can be seen in Figure 5. It has three different areas: one shows the preview rendering of the object whose appearance is being edited, as in most editing interfaces (we use the Cycles built-in progressive renderer for display); another shows the manifold (or a part of it, depending on the selected scale), together with an inset that shows a large version of the selected manifold sample; and a third area is devoted to showing the target image that will be

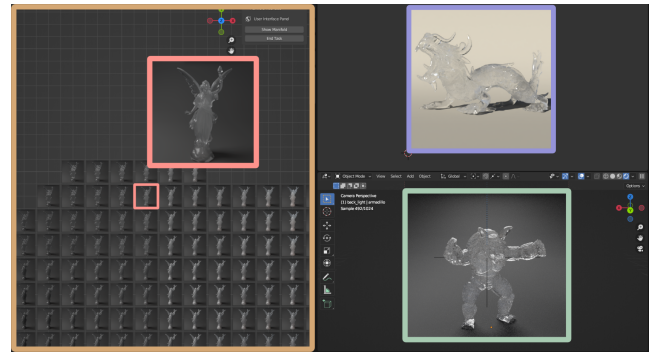


Figure 5: Screenshot of our proposed editing interface, as displayed in our user studies, while in image navigation mode. There are three distinct areas. Left: Manifold navigation area (orange box), which can be zoomed and panned, and includes a preview of the selected appearance at each instant (pink). Top right: Target image area, showing the target image (purple); this area is included for the purpose of our user studies, but can be hidden for normal interface usage. Bottom right: Preview area, showing the object being edited, rendered interactively as its appearance is changed (green).

used in the user studies. The third area is thus optional, and could be hidden resulting in a split screen-type interface. The arrow keys allow the user to traverse the manifold, whose scale can be changed to show a coarser or finer view of it. Finally, our interface allows to toggle between the three lighting conditions used to build the manifold, i.e., displaying the manifold with the *front*, *side* or *back* lighting condition (please also refer to the supplementary video for a demonstration of the interface).

5.2. User Study

We conduct a user study to validate the use of our manifold in an editing application scenario. Participants in the study perform editing tasks on a variety of translucent materials. Each editing task involves a *target* image and a *match* image, whose material appearance will be edited by the participants to match that of the target appearance. They conduct each task with two different interfaces: *ours* (described in Section 5.1), and a standard *slider-based* interface, in which the four optical parameters (α , σ_t , g and ρ) are exposed by means of sliders that the participants can tune to edit the material's appearance. The slider-based interface follows the same layout as ours (Figure 5), with the four sliders occupying the area reserved for the manifold in our interface.

Tasks The study features two types of tasks. The first type is a *matching task*, in which sixteen participants are asked to match the material of an object under fixed lighting to a target image of the *same* object and lighting. This provides a constrained task that facilitates quantitatively measuring users' performance [KP10]. In the second type, the *natural task*, ten participants are asked to match the material of an object under fixed lighting to a target image of a *different* object and under a *different* lighting. This is

designed to mimic a look-dev use case scenario, in which the user aims to match a certain appearance (e.g., from a photograph or any Internet image) on a new geometry and lighting.

Procedure For each interface and task, participants need to conduct a series of trials, one per target image that needs to be matched. We test four different materials per task; further, in the matching task, we test two different lighting conditions. This results in a total of eight target images in the *matching task* and four in the *natural task*, which need to be matched with each interface. We divide each experiment into two blocks, one per interface; the order of the blocks is randomized for each participant. Each block consists of four trials (i.e., matching four target images), one for each of the four target materials. Lighting conditions are pseudo-randomized so that each participant sees two trials with each lighting condition.

Participants start the experiment conducting a short demographics questionnaire, and then proceed to a training session. In the training session, they were asked to use the two interfaces to match the appearance of a glass-like material applied on the Stanford Bunny. Then, they perform one block, followed by a 5-minutes break, and later the second block. After each trial, participants answer a question regarding their level of satisfaction with the result, on a 1 to 7 Likert scale. In addition to that, we log time to completion, optical parameters of the final result, and toggling between lighting conditions events. There is no time limit to complete the tasks, and participants underwent a short questionnaire and debriefing session at the end of the study.

Participants Sixteen subjects participated in the matching task (four female, twelve male, none identifying themselves as non-binary, not listed, or preferring not to say). The average age was 24.8 years, with a standard deviation of 3.58 years. For the natural task, we recruited ten participants (four female, and six male, none identifying themselves as non-binary, not listed, or preferring not to say) among the sixteen initial subjects; in this way, participants were already familiar with both interfaces.

5.2.1. Matching Task

Stimuli We use the *Armadillo* statue for both the target image and match scene for this task. With this choice, we not only test the interface but also its ability to generalize to other geometries, since the manifold features the *Lucy* statue. We choose two lighting conditions: the *side* and *back* conditions described in Section 3.2.1, since they convey translucent appearance better than the *front* one. Stimuli are rendered with four target materials, selected to span a range of translucent appearance. Figure 6 shows the eight target images for the matching task (4 materials \times 2 lighting conditions).

Measured variables Since it is a matching task, we can compute the perceptual *error* (cubic root metric) between the target and the result obtained by each participant. We also analyze the *time* to completion of each trial, as well as the level of *satisfaction* provided by participants at the end of the trial (value in a 1 to 7 Likert scale).




Figure 6: Target images used in the matching task of the user study. We use the Armadillo statue with four different materials (columns), and each of these materials is rendered under two different light conditions: back (top row) and side (bottom row) lighting. Materials are selected such that they cover different regions of the manifold.

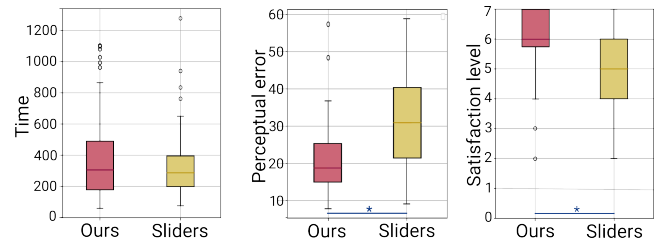


Figure 7: Comparison between interfaces for the matching task of the user study. We show, for each interface, the average time per trial (left), the error in image space computed using the cubic root metric (center), and the reported level of satisfaction with the result (right). Users are able to obtain better results, as shown by the increased satisfaction level and decreased error, using similar time than with the slider-based interface.

Results We use repeated measures ANOVA to test whether the interface used (two levels, *slider-based* vs. *ours*) had a significant effect on the perceptual error or the time to completion. We use Bonferroni correction, and set the threshold for significance at $\alpha = 0.05$. In the case of the satisfaction level, which is a rating on a 1 to 7 Likert scale, we use the non-parametric Wilcoxon signed-rank test to evaluate statistical significance.

Figure 7 shows all three measured variables for both interfaces tested. We can see that the results obtained are better in the case of our proposed interface than for the slider-based one, as shown by both the error and the level of satisfaction. Satisfaction (Figure 7, right) is not only significantly higher with our interface ($Z = 204.00$, p -value < 0.0001) but also has a lower standard deviation, while the error with respect to the target (Figure 7, center) is significantly lower for our interface ($F(1, 15) = 12.28$, p -value = 0.003). While the average editing time is longer with our interface, differences between both interfaces in terms of time to completion are not statistically significant; despite this, participants are more satisfied with our interface, possibly owing to the increased quality of the results.




Figure 8: Qualitative results obtained in the matching task with the interfaces tested. For two target images, we show the results obtained with each interface (slider-based and ours) by two different participants.




Figure 9: Fine-tuning with our interface. Given a target image (left), we show the final image selected in the manifold during image navigation (center left) and the match scene rendered with the corresponding optical parameters (center right). Its appearance can then be fine-tuned via sliders to reach a final result (right) closer to the target. Please refer to the text for more details.

We show qualitative results obtained by the participants in our study in Figure 8. For three different target images, we show the final appearance obtained by two participants, with both the slider-based and our interface. Within the complexity of the task, our interface enables more closely matching the target, while in the case of the slider-based interface, participants can really struggle to reach even a fairly similar appearance. Interestingly, when using our interface, in 42% of the trials participants did not fine-tune their results, and kept the result obtained through image navigation of the manifold. However, in a number of cases the fine-tuning step can play a significant role, as shown in the example in Figure 9.

Overall, our study indicates that, while sliders can be successfully tuned for fine edits (as shown in Figure 9), navigating the full space of translucent appearance with sliders is a daunting task for users. Meanwhile, image navigation provides a usable alternative, either to reach a good match, or to provide a good initial approximation in more challenging cases.

5.2.2. Natural Task

Stimuli The target images feature the *XYZ RGB Dragon* lit by an environment map, and are shown in Figure 10. The match scene is the same as the one in the previous task (see Section 5.2.1), therefore participants can select between three conditions for lighting



Figure 10: Target images used in the natural task of the user study. We employ the *XYZ RGB Dragon* geometry with four different materials, rendered with an environment map.




Figure 11: Comparison between interfaces for the natural task of the user study. We show, for each interface, the average time per trial (left), and the reported level of satisfaction with the result (right).

the scene (side, back, and front). Again, target images are rendered with four materials covering a wide range of appearance.

Measured variables In this case, in the absence of an objective measure of error (since the target and match images have different geometry and illumination), our measured variables are the *time* to completion of each trial, and the level of *satisfaction* provided by participants at the end of the trial.

Results We perform the same analyses as in the previous task for the two measured variables, time and level of satisfaction. Figure 11 shows these measures for both interfaces tested. While participants spent on average less time editing with our interface than with the slider-based one, this difference was not statistically significant. However, our interface yields a higher level of satisfaction with the result ($Z = 107.5$, p -value = 0.01). Both measures (especially level of satisfaction) exhibit less variance in the case of our interface. We also observe a drop in the time to complete each trial with respect to that of the matching task, with only a slight decrease in the level of satisfaction. Feedback during the debriefing sessions suggests that participants attempted to *approximate* the appearance, rather than exactly matching it. Sample qualitative results can be seen in Figure 12.

6. Discussion

We propose a novel way of exploring and editing translucent materials, via a continuous perceptually-motivated manifold of translucent appearance. To validate the benefits of this approach we have built an interface that leverages our manifold in an image navigation paradigm, followed by a fine-tuning step. An evaluation of its applicability is done by means of a user study, in which we compare our proposed interface to a standard, slider-based one, in two

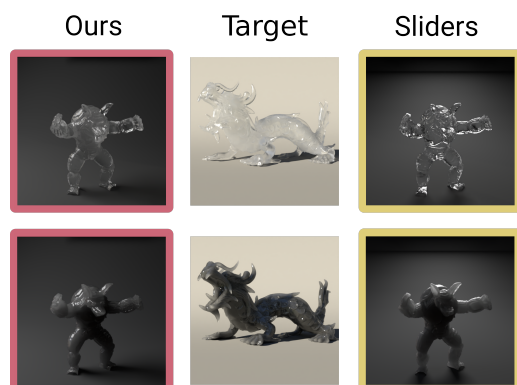


Figure 12: Qualitative results obtained in the natural task with the interfaces tested. In this task, participants had to mimic the appearance of a target image on a match scene that featured a different geometry and a different illumination. Despite the difficulty of the task, our interface led to significantly higher levels of satisfaction. The times to complete the task with our interface were 90s (top) and 218s (bottom), while for the slider-based interface were 325s and 265s, respectively.

editing tasks and for varied materials and geometries. Our interface enables non-expert users to reach a better result (based on both objective and subjective measures) with differences between both interfaces being statistically significant.

In particular, participants report higher levels of satisfaction, rating our interface as more intuitive and less frustrating than the traditional slider-based approach. Furthermore, for matching tasks, our navigation (plus slider-based fine-tuning) allowed a closer match to the reference, as well as a faster workflow for general look-dev tasks. In terms of workflow, the usual trend was to find a rough estimate of the appearance through navigation, and then locally explore this estimate, first using navigation, and then (sometimes) by fine-tuning using sliders. This is in line with the “block and refine” strategy found by Kerr and Pellacini [KP10]. At the same time, our results somewhat contradict the findings of Kerr and Pellacini, where a navigation-like paradigm was rated as the least effective, in that case for opaque reflectance modeling. This might be explained by the discrete nature of their image-based navigation through the reflectance parameters, as opposed to the continuous navigation on a perceptually-motivated manifold in our case, as well as by the fact that we focus on a different type of material.

The image-based navigation of our manifold requires computing the optical parameters via inverse rendering, which can lead to differences between the image in the manifold and the actual result. In most cases this was not noticeable, likely because translucency constancy is not good enough to perceive these small differences across different lighting or geometry. The supplemental material compiles detailed participants’ results illustrating this point.

The perception of translucency depends on geometry and lighting, and thus strictly speaking our analysis is only valid for the geometry and lighting considered. We address this by (i) including three lighting directions (commonly used in translucency perception); (ii) selecting a geometry with thin and thick parts, fine details

and sharp edges; and (iii) studying generalization to other geometries, including both complex and simpler ones (Section 4.4.3). We choose directional lighting because it better reveals the translucent nature of the material, and our natural task shows that users can leverage this manifold to edit complex geometries lit under an environment map. An analysis of the effect of other lighting conditions, including, e.g., more diffuse ones or other factors, could be done following the methodology laid out in this work.

Limitations The main limitation of image navigation leveraging our manifold, also pointed out by users, is its limited ability for very fine adjustments. This is particularly relevant for the matching task. This motivates the addition of a slider-based fine-tuning step complementing the image-based navigation, but it could potentially be addressed by a fully continuous navigation interface.

In terms of appearances, we have explored homogeneous, achromatic materials, and we have also limited our study to classical exponential media, to keep the problem tractable. The effect of colored albedo deserves further exploration, and the additional degrees of freedom of non-exponential transmittance [JAG18; BRM*18] have been proven useful in look-dev tasks in production [WVH17], at the cost of further increasing the complexity of editing. Having shown here the benefits that a perceptually-motivated manifold can have for editing translucent appearance, we hope these additional aspects will be explored in future work.

Acknowledgements

We thank the reviewers for their valuable feedback, and the participants that took part in the experiments. We also thank the members of the Graphics and Imaging Lab for their feedback, proof-reading the manuscript, and helping with the figures. This project has received funding from the European Union’s Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No 956585 (PRIME). The work was also partially funded by the MCIN/AEI/10.13039/501100011033/FEDER, UE through project PID2022-141539NB-I00.

References

- [Ado23] ADOBE. *Photoshop*. <https://www.adobe.com/products/photoshop.html>. 2023 3.
- [Aut23] AUTODESK. *Maya*. <https://www.autodesk.com/>. 2023 3.
- [Ble23] BLENDER FOUNDATION. *Blender*. <https://www.blender.org/>. 2023 8.
- [Boo97] BOOKSTEIN, FRED L. *Morphometric tools for landmark data*. Cambridge University Press, 1997 6.
- [BRM*18] BITTERLI, BENEDIKT, RAVICHANDRAN, SRINATH, MÜLLER, THOMAS, et al. “A radiative transfer framework for non-exponential media”. *ACM Trans. Graph.* (2018) 11.
- [Bur12] BURLEY, BRENT. *Practical Physically-Based Shading in Film and Game Production: Physically Based Shading at Disney*. SIGGRAPH 2012 Courses. 2012 3.
- [Bur15] BURLEY, BRENT. *Physically Based Shading in Theory and Practice: Extending the Disney BRDF to a BSDF with integrated subsurface scattering*. SIGGRAPH 2015 Courses. 2015 2, 3.
- [Cha60] CHANDRASEKHAR, S. *Radiative Transfer*. Dover Publications, 1960 2.

- [CLZ*20] CHE, CHENGQIAN, LUAN, FUJUN, ZHAO, SHUANG, et al. "Towards Learning-based Inverse Subsurface Scattering". *Proc. ICCP* (2020), 1–12 3.
- [CPK06] COLBERT, MARK, PATTANAİK, SUMANTA, and KRIVANEK, JAROSLAV. "BRDF-Shop: Creating Physically Correct Bidirectional Reflectance Distribution Functions". *IEEE Comput. Graph. Appl.* 26.1 (2006), 30–36 3.
- [DLC*22] DELANOY, JOHANNA, LAGUNAS, MANUEL, CONDOR, J, et al. "A Generative Framework for Image-based Editing of Material Appearance using Perceptual Attributes". *Computer Graphics Forum*. Vol. 41. 1. 2022, 453–464 5.
- [DLW*22] DENG, XI, LUAN, FUJUN, WALTER, BRUCE, et al. "Reconstructing Translucent Objects using Differentiable Rendering". *Proc. SIGGRAPH* (2022) 3.
- [FB05] FLEMING, ROLAND W and BÜLTHOFF, HEINRICH H. "Low-level image cues in the perception of translucent materials". *ACM Trans. Appl. Percept.* 2.3 (2005), 346–382 2.
- [FJB04] FLEMING, ROLAND W, JENSEN, HENRIK WANN, and BÜLTHOFF, HEINRICH H. "Perceiving translucent materials". *Proc. SAP*. 2004, 127–134 2.
- [GAGH14] GARCES, ELENA, AGARWALA, ASEEM, GUTIERREZ, DIEGO, and HERTZMANN, AARON. "A Similarity Measure for Illustration Style". *ACM Trans. Graph.* 33.4 (July 2014) 5.
- [GDHP20] GIGILASHVILI, DAVIT, DUBOUCHET, LUCAS, HARDEBERG, JON YNGVE, and PEDERSEN, MARIUS. "Caustics and translucency perception". *Electronic Imaging* (2020), 33–1 4.
- [GLZ16] GKIIOULEKAS, IOANNIS, LEVIN, ANAT, and ZICKLER, TODD. "An evaluation of computational imaging techniques for heterogeneous inverse scattering". *Proc. ECCV*. Springer. 2016, 685–701 3.
- [GSW*21] GIGILASHVILI, DAVIT, SHI, WEIQI, WANG, ZEYU, et al. "The Role of Subsurface Scattering in Glossiness Perception". *ACM Trans. Applied Perception* 18.3 (2021), 1–26 7.
- [GTHP21] GIGILASHVILI, DAVIT, THOMAS, JEAN-BAPTISTE, HARDEBERG, JON YNGVE, and PEDERSEN, MARIUS. "Translucency perception: A review". *Journal of Vision* 21.8 (Aug. 2021), 4–4 3.
- [GTPH21] GIGILASHVILI, DAVIT, THOMAS, JEAN-BAPTISTE, PEDERSEN, MARIUS, and HARDEBERG, JON YNGVE. "On the appearance of objects and materials: Qualitative analysis of experimental observations". *Journal of the International Colour Association* 27 (2021), 26–55 2, 4.
- [GUT*22] GIGILASHVILI, DAVIT, URBAN, PHILIPP, THOMAS, JEAN-BAPTISTE, et al. "The Impact of Optical and Geometrical Thickness on Perceived Translucency Differences". *Journal of Perceptual Imaging* 5 (2022) 4.
- [GWA*15] GKIIOULEKAS, IOANNIS, WALTER, BRUCE, ADELSON, EDWARD H, et al. "On the appearance of translucent edges". *Proc. CVPR*. 2015, 5528–5536 2, 4.
- [GXZ*13] GKIIOULEKAS, IOANNIS, XIAO, BEI, ZHAO, SHUANG, et al. "Understanding the Role of Phase Function in Translucent Appearance". *ACM Trans. Graph.* 32.5 (2013) 3–7.
- [GZB*13] GKIIOULEKAS, IOANNIS, ZHAO, SHUANG, BALA, KAVITA, et al. "Inverse volume rendering with material dictionaries". *ACM Trans. Graph.* 32.6 (2013), 1–13 3.
- [HG41] HENYEV, LOUIS G and GREENSTEIN, JESSE LEONARD. "Diffuse radiation in the galaxy". *The Astrophysical Journal* 93 (1941), 70–83 2.
- [HR13] HAŠAN, MILOŠ and RAMAMOORTHY, RAVI. "Interactive albedo editing in path-traced volumetric materials". *ACM Trans. Graph.* 32.2 (2013), 1–11 3.
- [ICT10] ICT VISION & GRAPHICS LAB. *High-Resolution Light Probe Image Gallery*. <https://vgl.ict.usc.edu/Data/HighResProbes/>. 2010 4.
- [JAG18] JARABO, ADRIAN, ALIAGA, CARLOS, and GUTIERREZ, DIEGO. "A radiative transfer framework for spatially-correlated materials". *ACM Trans. Graph.* 37.4 (2018), 1–13 11.
- [Jak10] JAKOB, WENZEL. *Mitsuba renderer*. 2010 4.
- [JMLH01] JENSEN, HENRIK WANN, MARSCHNER, STEPHEN R, LEVOY, MARC, and HANRAHAN, PAT. "A practical model for subsurface light transport". *Proc. CGIT*. 2001, 511–518 3, 4.
- [JSR*22] JAKOB, WENZEL, SPEIERER, SÉBASTIEN, ROUSSEL, NICOLAS, et al. *Mitsuba 3 renderer*. Version 3.0.1. <https://mitsuba-renderer.org>. 2022 6.
- [JZJ*15] JIMENEZ, JORGE, ZSOLNAI, KÁROLY, JARABO, ADRIAN, et al. "Separable subsurface scattering". *Computer Graphics Forum* 34.6 (2015), 188–197 3.
- [KB14] KINGMA, DIEDERIK P and BA, JIMMY. "Adam: A Method for Stochastic Optimization". *arXiv e-prints* (2014), arXiv–1412 6.
- [Ken89] KENDALL, DAVID G. "A survey of the statistical theory of shape". *Statistical Science* 4.2 (1989), 87–99 6.
- [KNYK22] KIYOKAWA, HIROAKI, NAGAI, TAKEHIRO, YAMAUCHI, YASUKI, and KIM, JUNO. "The perception of translucency from surface gloss". *Vision Research* (2022), 108140. ISSN: 0042-6989 2, 4.
- [KP10] KERR, WILLIAM B. and PELLACINI, FABIO. "Toward evaluating material design interface paradigms for novice users". July 2010 3, 8, 11.
- [LBFS21] LAVOUÉ, GUILLAUME, BONNEEL, NICOLAS, FARRUGIA, JEAN-PHILIPPE, and SOLER, CYRIL. "Perceptual quality of BRDF approximations: dataset and metrics". *Computer Graphics Forum* 40 (2021) 4.
- [LJM22] LANZA, DARIO, JARABO, ADRIAN, and MASIA, BELEN. "On the Influence of Dynamic Illumination in the Perception of Translucency". *Proc. SAP*. 2022, 1–9 2.
- [LMS*19] LAGUNAS, MANUEL, MALPICA, SANDRA, SERRANO, ANA, et al. "A Similarity Measure for Material Appearance". *ACM Trans. Graph.* 38.4 (2019) 4, 5.
- [LNN23] LI, CHENHAO, NGO, TRUNG THANH, and NAGAHARA, HAJIME. "Inverse Rendering of Translucent Objects using Physical and Neural Renderers". *Proc. CVPR*. 2023, 12510–12520 3.
- [LSX22] LIAO, CHENXI, SAWAYAMA, MASATAKA, and XIAO, BEI. "Unsupervised learning reveals interpretable latent representations for translucency perception". *PLoS Computational Biology* 19 (2022) 3.
- [Man67] MANTEL, NATHAN. "The detection of disease clustering and a generalized regression approach." *Cancer research* 27 2 (1967), 209–20 7.
- [Mat03] MATUSIK, WOJCIECH. "A data-driven reflectance model". PhD thesis. Massachusetts Institute of Technology, 2003 3.
- [MCR*21] MANTIUK, RAFAŁ K., CHAPIRO, ALEXANDRE, RUFO, GIZEM, et al. "FovVideoVDP: a visible difference predictor for wide field-of-view video". *ACM Trans. Graph.* 40 (2021), 49:1–49:19 4.
- [Mic23] MICROSOFT. *Microsoft PowerPoint*. <https://www.microsoft.com/es-es/microsoft-365/powerpoint>. 2023 3.
- [NDM06] NGAN, ADDY, DURAND, FRÉDO, and MATUSIK, WOJCIECH. "Image-driven Navigation of Analytical BRDF Models". *Symposium on Rendering*. Ed. by AKENINE-MOELLER, TOMAS and HEIDRICH, WOLFGANG. 2006 4, 5.
- [NGD*06] NARASIMHAN, SRINIVASA G, GUPTA, MOHIT, DONNER, CRAIG, et al. "Acquiring scattering properties of participating media by dilution". *Proc. SIGGRAPH*. 2006, 1003–1012 3.
- [NGHJ18] NOVÁK, JAN, GEORGIEV, ILIYAN, HANIKA, JOHANNES, and JAROSZ, WOJCIECH. "Monte Carlo Methods for Volumetric Light Transport Simulation". *Computer Graphics Forum* 37.2 (May 2018) 2.
- [NJR15] NIELSEN, JANNIK BOLL, JENSEN, HENRIK WANN, and RAMAMOORTHY, RAVI. "On Optimal, Minimal BRDF Sampling for Reflectance Acquisition". *ACM Trans. Graph.* 34.6 (2015) 3.

- [PFG00] PELLACINI, FABIO, FERWERDA, JAMES A, and GREENBERG, DONALD P. "Toward a psychophysically-based light reflection model for image synthesis". *Proc. SIGGRAPH*. 2000, 55–64 2, 3.
- [PGSP08] PACANOWSKI, ROMAIN, GRANIER, XAVIER, SCHLICK, CHRISTOPHE, and POULIN, PIERRE. "Sketch and Paint-based Interface for Highlight Modeling". *Eurographics Workshop on Sketch-Based Interfaces and Modeling*. Ed. by ALVARADO, CHRISTINE and CANI, MARIE-PAULE. 2008 3.
- [SGM*16] SERRANO, ANA, GUTIERREZ, DIEGO, MYSZKOWSKI, KAROL, et al. "An intuitive control space for material appearance". *ACM Trans. Graph.* 35.6 (2016) 2–4.
- [SL23] SUBIAS, J DANIEL and LAGUNAS, MANUEL. "In-the-wild Material Appearance Editing using Perceptual Attributes". *Computer Graphics Forum*. Vol. 42. 2. Wiley Online Library. 2023, 333–345 5.
- [STPP09] SONG, YING, TONG, XIN, PELLACINI, FABIO, and PEERS, PIETER. "SubEdit: A Representation for Editing Measured Heterogeneous Subsurface Scattering". *ACM Trans. Graph.* 28.3 (2009) 2, 3.
- [SWSR21] SHI, WEIQI, WANG, ZEYU, SOLER, CYRIL, and RUSHMEIER, HOLLY. "A Low-Dimensional Perceptual Space for Intuitive BRDF Editing". *Proc. EGSR*. 2021 2, 3.
- [TdL00] TENENBAUM, JOSHUA B., DE SILVA, VIN, and LANGFORD, JOHN C. "A Global Geometric Framework for Nonlinear Dimensionality Reduction". *Science* 290.5500 (2000), 2319–2323 5.
- [TGG*20] TOSCANI, MATTEO, GUARNERA, DAR'YA, GUARNERA, GIUSEPPE CLAUDIO, et al. "Three Perceptual Dimensions for Specular and Diffuse Reflection". *ACM Trans. Appl. Percept.* 17.2 (2020) 2, 3.
- [WAKB09] WILLS, JOSH, AGARWAL, SAMEER, KRIEGMAN, DAVID, and BELONGIE, SERGE. "Toward a perceptual space for gloss". *ACM Trans. Graph.* 28.4 (2009), 1–15 2, 3, 5, 6.
- [WCW*08] WANG, RUI, CHESLACK-POSTAVA, EWEN, WANG, RUI, et al. "Real-Time Editing and Relighting of Homogeneous Translucent Materials". *Vis. Comput.* 24.7 (July 2008), 565–575 2.
- [WMLT07] WALTER, BRUCE, MARSCHNER, STEPHEN R, LI, HONGSONG, and TORRANCE, KENNETH E. "Microfacet Models for Refraction through Rough Surfaces." Vol. 2007. 2007 2.
- [WPW89] WYMAN, DOUGLAS R, PATTERSON, MICHAEL S, and WILSON, BRIAN C. "Similarity relations for the interaction parameters in radiation transport". *Applied optics* 28.24 (1989), 5243–5249 7.
- [WVH17] WRENNINGE, MAGNUS, VILLEMEN, RYUSUKE, and HERY, CHRISTOPHE. "Path traced subsurface scattering using anisotropic phase functions and non-exponential free flights". *Tech. Rep.* Pixar Inc., 2017 3, 11.
- [XGL*07] XU, KUN, GAO, YUE, LI, YONG, et al. "Real-time homogeneous translucent material editing". *Computer Graphics Forum* (2007) 3.
- [XWG*14] XIAO, BEI, WALTER, BRUCE, GKIOULEKAS, IOANNIS, et al. "Looking against the light: How perception of translucency depends on lighting direction". *Journal of vision* 14.3 (2014), 17–17 4.
- [ZIE*18] ZHANG, RICHARD, ISOLA, PHILLIP, EFROS, ALEXEI A., et al. "The Unreasonable Effectiveness of Deep Features as a Perceptual Metric". *Proc. CVPR* (2018), 586–595 4.
- [ZRB14] ZHAO, SHUANG, RAMAMOORTHI, RAVI, and BALA, KAVITA. "High-order similarity relations in radiative transfer". *ACM Trans. Graph.* 33.4 (2014), 1–12 2, 6, 7.