

Generating Materials for Augmented Reality Applications using Natural Language

Sebastian Buntin
sebastian@buntin.de

ABSTRACT

In this paper, I present a novel method in parametrizing the BRDF (and other BxDF) as well as the texture map of a material by using natural language. The visual properties of a material can be described by the user using rather complex phrases in real time. These phrases then will be parsed and mapped to BRDF and texture parameters. This allows an easy way to create and specify material representations for various applications. In the context of this paper, I focus on the application in augmented reality.

CCS Concepts

• **Computing methodologies** → *Reflectance modeling; Mixed / augmented reality;*

Keywords

1. INTRODUCTION

The appearance of a material is the visual impression a human perceives by looking at a certain object. In computer graphics, the creation of a photo-realistic image is achieved by modeling this appearance. This modeling is also one of the main limiting factors for achieving photo realism in computer generated images or scenes. A lack of accuracy in the scene description, specifically in the material description or the material properties, will inevitably yield unconvincing images. In local illumination models, the appearance of a material is based on (other than the materials' intrinsic parameters) the viewing angle, the lighting set-up and the scale on which the object is viewed at. The material itself can be characterized by three important components: a spectral component that describes the color; a directional component, describing the shininess, haziness, and glossiness; and a spatial component that describes spatial variations of the material (e.g. small cracks and bumps in the material surface, variations in the color). For people without a strong background in the theory of computer graphics, this modeling process can be quite complicated.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

AH 2016, February 25 - 27, 2016, Geneva, Switzerland

© 2016 Copyright held by the owner/author(s). Publication rights licensed to ACM. ISBN 978-1-4503-3680-2/16/02...\$15.00

DOI: <http://dx.doi.org/10.1145/2875194.2875195>

In general conversation, it can be observed that people without a strong computer graphics background describe the appearance of a material using common language. Without a detailed knowledge of the mathematics behind reflectometry, people use a large variety of adjectives and adverbs to describe the visual appearance of a material. A material is also described by the **present** attributes (e.g. a surface is usually described as *smooth* rather than *not rough*) which, in natural language, are represented by adjectives (smooth, rough, anodized, blue, glossy,...) which can be further refined by adverbs (lightly, heavily,...). As an example, if one says “lightly brushed, blue anodized aluminum”, most people can imagine how the material will appear. The material is specified by an abrasive machining technique (*brushing*) and a surface treatment (*anodizing*).

The word “material” in the context of this paper refers to its interpretation in the context of computer graphics. A 3D-object can be rendered using a “material”. This can be a simple color, a texture, a shader or a combination of some or all of them. The “material” specifies how the incoming light is reflected, scattered and absorbed by the object’s surface and does not have to refer to an actual physical material, e.g. “lightly brushed, blue anodized aluminum” is a “material” in this context.

The novelty of this paper is the possibility to parametrize a BRDF and a texture synthesizer using natural language to describe the appearance of a material.

2. RELATED WORK

2.1 Materials

A comprehensive overview of the digital modeling of materials, especially the mathematical background, the material models, measurement of BRDF and BTF, further effects and some rendering techniques can be found in [6]. [10] discusses the measurement of different appearance attributes as well as the mathematical and physical background in great detail. The science of vision and perception is discussed in [18] as an interdisciplinary approach in the fields of physiology, psychology, physics and computer science.

For the rendering of 3D-objects using “materials”, there are a number of different BRDF models for different usage scenarios. For diffuse only materials, [14] provides a simple but powerful model widely used in computer graphics. This model was generalized by [17] for rough surfaces by having the diffuse light component dependent of the viewing angle. The most widely used models for the specular reflection (especially for real-time rendering) for isotropic reflectance are

described in [3], [5], [9]. Those models are based on different aspects of the bidirectional reflectance distribution functions such as their representation, measurement or derivation. Models for anisotropic reflectance are described in [22] and [2] where the latter also provides a non-lambertian diffuse term to calculate specular and diffuse lighting.

[7, 230-293] provides a detailed overview of using those BRDF in real time rendering.

In the area of layered material (e.g. varnished wood), [8] were one of the first to introduce a method to render multi-layered materials with sub-surface scattering. [11] provides a general method for computing BSDFs of layered materials. Neither method is suited for real-time rendering. [23] on the other hand provide a method using stacked micro-facet (e.g. Cook-Torrance-BRDFs) to simulate layered surfaces. The description of their method is quite easily adaptable for real-time rendering using ordinary rasterization hardware.

2.2 Material Generation

For the design of a material (i.e. finding the parameters for a mathematical BRDF model) based on a given picture or object, [12] evaluated three different methods for designing a material with a BRDF model (in their case [5]). They compare “physical sliders”, where the parameters of the BRDF model are edited directly; “perceptual sliders”, a method presented in [19] which use uses a perceptual parametrization of the Ward-BRDF where each parameter represents a “perceptually meaningful dimension” of the surface appearance; and “image navigation”, a method where variations of the material can be viewed by browsing thumbnail images.

[24] developed a correspondence domain for the parameters of the Phong, Ward and Cook-Torrance BRDF based on gloss and haze. Furthermore, they describe a reflectance model for use with goniochromatic surfaces (surfaces with viewing angle dependent color, e.g. pearlescent car paint) and the measurement for those parameters.

[1] focuses on the perception and appearance of “stuff” (sic), meaning the perception of materials by humans and machines in a general context. He explains the importance of materials to vision and, more important, he tries to give a vocabulary to describe the visual appearance of a material.

2.3 Natural Language in Computer Graphics

One of the first appearances of natural language in a computer graphics context was [4]. He combines speech recognition and position sensing technologies to command simple shapes on a large screen display using a set of simple sentences and commands. [15] introduces the concept of a “Knowledge Representation layer”, using entities, attributes and relations to describe the “knowledge” (objects, their visual and non-visual attributes, their relationships, modelling operations etc.) of the system using a semantic representation. [16] uses natural language to specify the parameters of a music synthesizer. Although this is not computer graphics, the basic concepts of parametrizing an entity with a set of attributes is comparable to our method. He uses a loose phrase structure grammar to be able to describe the sound (filters,...) of a synthesizer using words like “fat”, “soft” etc.

3. MATERIAL DESIGN USING NATURAL LANGUAGE

3.1 Finding the right words

The basic vocabulary was found by analyzing corresponding literature ([1], [6], [18], [10]) and our own experience. It is known that a vocabulary “based on our experience” might be a highly questionable method but it gives us, combined with the vocabulary extracted from the literature, a good starting point that can easily be adapted in the future. One of the main issues is that the vocabulary can be highly complex and different terms with the same meaning could be used. [16] showed that the number of adjectives to describe (in his case) sound rises quickly and becomes quite diverse when taken from user experience or user interviews. This shows that the unrestricted usage of “every day language” is not feasible and that an approach using controlled natural language ([13]) which reduces the the vocabulary to a predefined subset should be used.

Using only adjectives to describe the visual properties of a material is not detailed enough. “Varnished wood”, for example, leaves far too much room for interpretation as, in this example, only the type of coating is specified and not the coating color. The imminent question that arises is: “What color does the varnish have?”. “White varnished wood” on the other hand specifies the color of the coating, giving a more detailed description of the visual appearance of the “material”, so adverbial used adjectives will be taken into account too.

After choosing a basic set of adjectives and adverbs, several services can be used to extend the vocabulary such as the Wortschatz (wortschatz.uni-leipzig.de). The Wortschatz provides a web interface for finding synonyms (related meaning) and antonyms (opposite meaning) in the German language.

3.2 Parsing the user input

Depending on the application, the user “enters” the verbal description of the material using Human-Computer-Interface technologies (e.g. speech recognition, keyboard). In the current state of the system, these phrases have to follow a (relatively) simple grammar required by the underlying natural language processing system.

Using the provided vocabulary (excerpts are provided in the appendix), a natural language processing system will parse the user input. As said, there are currently only limited grammatical constructs available.

The following grammatical constructs can be parsed :

- Noun (e.g. aluminum, wood, brass)
- Adjective Noun (e.g. anodized aluminum, varnished wood, polished brass)
- Adverbial-adjective Adjective Noun (e.g. blue anodized aluminum, white varnished wood, lightly brushed brass)
- several 'Adverbial-adjective Adjective' terms, separated by comma; Noun (e.g. lightly brushed, blue anodized aluminum)

In the first step, a part-of-speech tagger annotates (tags) the phrase to find adverbs, adjectives and nouns. The noun typically denotes the base material. Adjectives and adverbs provide more detail on the visual appearance (see Appendix). Several groups of adverbs and adjectives, separated by commas (as used in natural language), can be provided (e.g. “slightly brushed, blue anodized aluminum”)

For example, “aluminum” [Noun] denotes a silverish, highly specular material. “Anodized aluminum” [Adjective Noun] adds (or enhances) the top oxide layer of the Aluminum (by about $1\mu m$) resulting in more diffuse reflection properties. “Blue anodized aluminum” [Adverbial-adjective Adjective Noun] further describes the anodization process. The result will be a blue metallic surface with slightly diffuse reflection properties. Combining this phrase with another adjective (“brushed, blue anodized aluminum”) or adverbial-adjective term (“lightly brushed, blue anodized aluminum”), the description of the visual appearance of the material can be further refined, in this case the brushing results in an even higher diffuse value with anisotropic properties. In its current state, connecting multiple phrases with conjunctions (*and*) or disjunctions (*or*) are not available at the moment but will be implemented in a future version.

In the second step, the inflection (especially if working with the German language) of the words needs to be analyzed and the base form has to be found. This is done using a simple baseform lexicon.

The biggest problem with the specification of attributes using natural language is “intent vs. interpretation”. This means, that the user might intent something different with his input than the framework “interprets” by parsing and analyzing it. For example, when asked about the color “blue” people might think of colors ranging from “light blue” to “navy blue”, all being valid shades of “blue”. Since the earliest speech processing systems in computer graphics ([4]), this remains a more or less unresolvable issue, as different people have a different personal definition of the meaning behind descriptive words (e.g. adjectives like “blue”, “rough”, “dark”). To resolve this issue, the user can provide further modifiers (*more*, *less*, *stronger*, *darker*, *lighter*; e.g. “stronger brushing”, “darker blue”) in several refining steps. This frees the user from specifying discrete values for parameters, as this would contradict our intention. Using the further refinements, the system still has enough potential to find the users intended attribute.

3.3 Mapping the input

After tagging and parsing the user input, an appropriate “set of instructions” needs to be found that describes how to render the material. As multiple target frameworks can be supported (different real-time engines as well as raytracing software), this involves a target-independent intermediate layer: The tagged input is mapped to a BRDF, whose parameters are chosen according to the material (the noun) and its attributes (the adverbs and adjectives). To achieve this, an internal database is used to find the appropriate BRDF for a specified material and its coating properties (if given). For materials with spatially varying attributes (color, normals), the appropriate (texture) maps can be (procedurally) generated as well. For example, “brushed Aluminum” can have a visual structure in the surface resulting in modifications of the surface normal. This can be achieved with a procedurally generated normal map which adds the impression of the brush strokes. Other types of abrasive machinery can be simulated accordingly. For the material “wood”, a procedural texture (diffuse map) can be created providing the spatially varying color information, in this case, the wood grain. For materials that simulate aging or weathering, a parameter map changing the parameters of the BRDF can be generated (e.g. “rust”, in combination

with a texture map).

For every material that is available in the Application, an initial BRDF and a set of base BRDF parameters will be retrieved from the applications internal database. Every word can influence some or all parts of the output BRDF. This means, that some attributes change the color or glossiness of the target BRDF, or are also able to change the BRDF completely.

The default case will be an isotropic BRDF (see refsec:rendering). Some materials (e.g. velvet) or machining techniques (e.g. brushing) change the surface reflection properties to show anisotropic behavior. This will override the default isotropic BRDF and choose an anisotropic BRDF (Ward’s BRDF).

These parameters for the isotropic case are usually the *roughness* (m), *index of refraction* (η) and the *extinction coefficient* (k) if the material is a conductor. m is usually in the range $[0, 1]$, η in the range $[1, 5]$ and k in the range $[1, 10]$. In the anisotropic case, the parameters α_x and α_y (the standard deviation of the surface slope in x and y direction) as well as ρ_d (the specular reflectance of the material) have to be specified. (For a more detailed explanation of the BRDF, see Chapter 4.1).

The base BRDF can be modified using adjectives, especially if the material is coated with a relatively “thick” layer (e.g. varnished wood). The base parameters are modified according to the adjectives as well as adverbs. In its current implementation, these modifying factors are stored per adjective and adverb in the internal database. In the previous example, “anodizing” increases the value of the parameter m by 0.1. “Blue” will change the color attribute to a blue color (e.g. $RGB(0, 0, 1)$). The adjective “brushed” will trigger the procedural part and generate a normal map simulating abrasive machinery (in this case, ‘brushing’). The adverb “lightly” works as an input parameter for the procedural unit decreasing the strength of the gradient in the height map which is converted to a normal map.

In the anisotropic case, for the Ward-BRDF, three parameters are necessary, α_x and α_y in the range $[0, 1]$ and ρ , the value for the specular reflectance (the magnitude of the BRDF lobe). The α -Parameters can be matched with the roughness parameter m from the isotropic case, but the “strength” of the anisotropy, e.g. in which direction on the surface the spherical highlight splats out, depends on the difference between α_x and α_y , as they represent the roughness depending on the surface geometry (the width of the BRDF lobe).

In the case of a multi-layered material (e.g. a semitransparent layer of varnish on wood), a stacked microfacet BRDF according to [23] is selected. Both layers then need values for m , η , k . After the BRDF is chosen according to the entry in the database for the specified material, it is initialized with default parameters for the material. These parameters are also located in the database. For example, just entering “Aluminium” into the system would result in an isotropic BRDF with light grey color ($RGB(0.75, 0.75, 0.75)$), and the other parameters being $m = 0.1$; $\eta = 1.097$; $k = 6.794$ (Microfacet Model, see 4.1.1).

In the following steps, the adjectives and adverbs are mapped. An adjective can have a default formula assigned that modifies BRDF parameters or a per-material formula. For example, “brushed” increases the roughness by 20% ($m = m \cdot 1.2$), if the material is “Wood”, the roughness is increased by 25%

($m = m \cdot 1.25$). The same applies for adverbs. For example, "lightly brushed" will change the factor of 20% to only 10%.

The decision, as to whether 2D-maps are needed is done based on the noun (the material itself, e.g. wood and marble). If the material usually has visible spatial variances (wood grain, rusted metal), a (procedural) color (texture) map will be needed for these effects. For some materials, these procedural steps are predefined but can be varied by the user using natural language. The texture for wood, for example, could be altered by "less grain" or "higher contrast". The decision for a normal map is either based on the material (e.g. rough stone) or based on abrasive machining techniques applied to the surface (brushing).

After finding the parameters for the BRDF based on the users input, these values are combined with a shader implementing one of the BRDFs for the corresponding target platform.

4. IMPLEMENTATION

4.1 Rendering

4.1.1 Microfacet BRDF

Most physically based models (like [5]) are based on microfacet theory. In this theory, the surface can be modeled with very small micro-facets (too small to be seen), which are assumed to be perfect reflectors and their distribution can be described by statistical models. Thus, the BRDF itself is represented by an analytic function as seen in 1.

$$f(\vec{x}) = \frac{FGD}{4(\vec{N} \cdot \vec{L})(\vec{N} \cdot \vec{H})} \quad (1)$$

D represents the microfacet distribution function (2) where m specifies the roughness of the surface. This distribution describes the statistical distribution of the micro-facets normals. \vec{N} is the normal vector of the surface, \vec{V} is the direction to the viewer, \vec{H} is the half-vector between \vec{L} (direction of incoming light) and \vec{V} ($\vec{H} = \frac{\vec{L} + \vec{V}}{|\vec{L} + \vec{V}|}$).

$$D(m) = \frac{m^2}{\pi((\vec{N} \cdot \vec{V})^2(m^2 - 1) + 1)^2} \quad (2)$$

G represents the self-shadowing of the microfacets. It is also a statistical term and basically models the probability of shadowing of the microfacets between each other. The functions for D and G used here are the GGX functions proposed in [21].

$$G(m) = \frac{2}{1 + \sqrt{m^2 + (1 - m^2)}} \quad (3)$$

F represents the Fresnel part. It describes how much light is reflected and refracted (absorbed). If the material is a dielectric, equation (4) is used: For conductors, it is common to use Schlicks ([20]) approximation as seen in equation (5). η_1 and η_2 are the indices of refraction. If the object is rendered in "air", η_1 can be set to 1. k is the extinction coefficient of the conducting material.

$$F = F_0 + (1 - F_0)(1 - \vec{H} \cdot \vec{V})^5 \quad (4)$$

$$F_0 = \left(\frac{\eta_1 - \eta_2}{\eta_1 + \eta_2} \right)^2$$

$$F = \frac{(\eta - 1)^2 + 4\eta(1 - \cos\theta)^5 + k^2}{(\eta + 1)^2 + k^2} \quad (5)$$

This BRDF is used in nearly all modern-day game engines (Unreal Engine, Unity, Frostbite) to simulate physically based rendering as this BRDF allows to specify parameters for a big range of materials thus making it ideal for general material rendering.

The parameters which need to be 'extracted' out of the user's input by the language processor are the roughness m , the indices of refraction η_1 (which is 1 for air), η_2 and k if the material is a conductor (see Chapter 4.2).

The BRDF can then easily be implemented with shaders for different target platforms.

4.1.2 Ward's BRDF

In case of anisotropic materials, Ward's BRDF is chosen:

$$f = \rho \frac{\vec{N} \cdot \vec{L}}{4\pi\alpha_x\alpha_y} e^{-2 \frac{\left(\frac{\vec{H} \cdot \vec{X}}{\alpha_x}\right)^2 + \left(\frac{\vec{H} \cdot \vec{Y}}{\alpha_y}\right)^2}{1 + (\vec{H} \cdot \vec{N})}} \quad (6)$$

with

$$\rho = \frac{\rho_s}{\sqrt{(N \cdot L)(N \cdot V)}} \quad (7)$$

The Ward BRDF ([22]) is an empirical model to fit measured BRDF data and is able to handle anisotropic reflectance. It is also widely used in computer graphics.

4.1.3 Procedural Generation of Maps

The maps can be predefined using a provided bitmap or procedurally generated. For the procedural generation, a simple graph-like API was developed which supplies generators (Perlin noise, Voroni diagrams (also known as Worley Noise), gradients, tiles, checkerboards, bricks, functions like square, saw and sine) that can be connected using filters (blending, blur, color correction, colorize, distort, transform, bump) and targets (color map, normal map, parameter map) in a graph. All these generators, filters and targets represent nodes in the graph which can be parametrized. The output is always a 2D texture.

This gives a simple, yet quite powerful possibility to procedurally generate texture maps for e.g. wood and marble as well as normal maps to simulate visible 'bumps' in the surface generated by e.g. abrasive machinery techniques. A parameter map can be used to simulate effects like weathering, which changes the parameters of the material's BRDF. Using the noise-function based generators allows to predefine a graph for a material (e.g. wood) and generate billions of possible maps just by modifying the seed of the noise functions. Using the filters, the same input can generate a texture and a normal map for a material.

Materials like wood and marble come with predefined graphs that can be altered by language ("less grain", "higher contrast",...) thus providing an initial visual experience.

4.2 Language Processing

The language processing part of the application is based on a self-developed part-of-speech tagger which uses a lexicon of the German language to annotate the provided words. Whilst annotating the words, the word is transformed into its base form using the Wortschatz-Project (see Chapter 3.1) to allow easier processing.

After annotating and transforming a word to its base form, the adjectives are mapped to the BRDF parameters along

with their values. This is done using an attribute-value matrix. If an adverb was specified in front of the adjective, its 'influence' works as a weight to the numerical representation of the adjective. If we take the previous example, "blue anodized aluminum", "aluminum" represents a conductor, thus the equation (5) is used with a low value for $m(0.1)$, η with the refractive index for aluminum with a value of 1.097 and k will be 6.794. The adjective "anodized" will first be transformed to the base form, "anodize" and then mapped to a value for m of 0.3, overriding the original value for m . The adverb "blue" then maps the surface color to $RGB(0, 0, 1)$.

If the user then specifies "darker blue" in a further refining step, the RGB-color value is changed to $RGB(0, 0, 0.5)$. The input "less diffuse" would decrease the value for m to 0.2. If the user then, in a third input step would specify "more diffuse", the value for m would be increased to 0.25. This is achieved by internally specifying an interval between the 'old' value (in this case 0.3) and the current value (0.2) and then selecting the middle of said interval.

4.3 Augmented Reality

The augmented reality part in the application is implemented using ARToolkit5 (www.artoolkit.org). ARToolkit provides an API for the camera input and the feature tracking.

5. RESULTS




Figure 1: "Wood" in an augmented reality environment

The results show the materials "blue anodized Aluminum" and "wood" rendered on the Stanford Dragon in an augmented reality application using ARToolkit5. The rendering is in real time with 50 frames per second (the framerate is limited by the camera). The tracking is done using the Hiro-Pattern. The language processing takes less than a second on standard PC hardware (e.g. Quadcore Intel i5, 3GHz, nVidia Geforce GTX960), if no internet connection to the Wortschatz is used, which would be the main limiting factor. The wood texture is a simple texture based on Simplex Noise. The second image shows a screenshot out of the earlier development stage of the application. No augmented reality is used. The material shown here is "brushed brass".




Figure 2: "Blue anodized aluminum" in an augmented reality environment




Figure 3: brushed brass

All in all the application is fully suitable for real time rendering and evaluation. Even if the base form transformation is fully used (not cached), the results will be available within a few seconds. First tests on mobile platforms (Samsung Galaxy S5 + S6) using the Android operation system and Unity as a 3D engine have shown very good results too.

6. CONCLUSIONS

The method shown in this paper allows an easy and interactive way to describe materials in 3D applications, especially in augmented reality situations. Using natural language as an input for the material part of the rendering, inexperienced users are able to quickly generate photo-realistic materials without requiring knowledge in computer graphics.

7. FURTHER RESEARCH

7.1 General

To allow a wider range of materials, e.g. materials such as wax or skin which rely on sub-surface scattering effects or

translucency, the rendering has to be extended to not just using the BRDF (which only handles reflection) but to also use the BTDF (Bidirectional Transmittance Distribution Function) and the BSDF (Bidirectional scattering-surface reflectance distribution function) to account for effects like transmittance (e.g. in glass) or sub-surface scattering. This leads to a much higher implementation effort in the target engines because those effects cannot be achieved by a “simple” shader but rely heavily on multipass rendering.

For real time apps on mobile platforms, a voice recognition module can be used on top of the input system to allow the user to describe the material verbally. This leads to a much easier handling of the application.

Finally, the speech processing system can be extended to allow conjunctions (*and*) and disjunctions (*or*).

7.2 User studies

As this paper describes an initial proof-of-concept if a verbal description of a material is possible for computer graphics, further studies have to be conducted. Especially the initial parameters and available adjectives and adverbs rely mostly on our own experience or on predefined values from literature, a user study to find suitable parameters for a wider range of materials is being prepared at the moment. In this study, the user will be shown several phrases describing a material and related images of renderings of this material. The user then has to select the image that best fits the phrase. The results will be statistically analyzed and shall provide a better set of initial parameter values.

8. REFERENCES

- [1] E. H. Adelson. On seeing stuff: the perception of materials by humans and machines. In B. E. Rogowitz and T. N. Pappas, editors, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 4299, pages 1–12, June 2001.
- [2] M. Ashikhmin and P. Shirley. An anisotropic phong brdf model. *J. Graph. Tools*, 5:25–32, February 2000.
- [3] J. F. Blinn. Models of light reflection for computer synthesized pictures. In *Proceedings of the 4th annual conference on Computer graphics and interactive techniques*, SIGGRAPH ’77, pages 192–198, New York, NY, USA, 1977. ACM.
- [4] R. a. Bolt. Put-that-there. *Proceedings of the 7th annual conference on Computer graphics and interactive techniques - SIGGRAPH ’80*, pages 262–270, 1980.
- [5] R. L. Cook and K. E. Torrance. A reflectance model for computer graphics. In *Proceedings of the 8th annual conference on Computer graphics and interactive techniques*, SIGGRAPH ’81, pages 307–316, New York, NY, USA, 1981. ACM.
- [6] J. Dorsey, H. Rushmeier, and F. Sillion. *Digital Modeling of Material Appearance*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2007.
- [7] W. Engel. *Programming vertex geometry and pixel shaders, 2nd Edition*. Programming Series. Charles River Media, 2008.
- [8] P. Hanrahan and W. Krueger. Reflection from layered surfaces due to subsurface scattering. In *Proceedings of the 20th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH ’93, pages 165–174, New York, NY, USA, 1993. ACM.
- [9] X. D. He, K. E. Torrance, F. X. Sillion, and D. P. Greenberg. A comprehensive physical model for light reflection. *SIGGRAPH Comput. Graph.*, 25:175–186, July 1991.
- [10] R. S. Hunter and R. W. Harold. *The measurement of appearance / Richard S. Hunter, Richard W. Harold*. Wiley, New York :, 2nd ed. edition, 1987.
- [11] W. Jakob, E. D’Eon, O. Jakob, and S. Marschner. A comprehensive framework for rendering layered materials. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2014)*, 33(4), 2014.
- [12] W. B. Kerr and F. Pellacini. Toward evaluating material design interface paradigms for novice users. *ACM Trans. Graph.*, 29:35:1–35:10, July 2010.
- [13] T. Kuhn. A survey and classification of controlled natural languages. *Computational Linguistics*, 40(1):121–170, March 2014.
- [14] J. H. Lambert. *Photometria sive de mensura de gratibus luminis, colorum umbrae*. Eberhard Klett, 1760.
- [15] M. E. Latoschik, P. Biermann, and I. Wachsmuth. Knowledge in the Loop: Semantics Representation for Multimodal Simulative Environments. *Proceedings of the 5th International Symposium on Smart Graphics 2005*, pages 25–39, 2005.
- [16] J.-t. Milde. Using natural language to specify sound parameters. *Proceedings of the ICMC-SMC*, pages 1627–1630, September 2014.
- [17] M. Oren and S. K. Nayar. Generalization of lambert’s reflectance model. In *In SIGGRAPH 94*, pages 239–246. ACM Press, 1994.
- [18] S. E. Palmer. *Vision Science, Photons to Phenomenology*. MIT Press, 1999.
- [19] F. Pellacini, J. A. Ferwerda, and D. P. Greenberg. Toward a psychophysically-based light reflection model for image synthesis. In *SIGGRAPH*, pages 55–64, 2000.
- [20] C. Schlick. An Inexpensive BRDF Model for Physically-based Rendering, 1994.
- [21] B. Walter, S. Marschner, H. Li, and K. Torrance. Microfacet models for refraction through rough surfaces. *Eurographics*, pages 195–206, 2007.
- [22] G. J. Ward. Measuring and modeling anisotropic reflection. In *Proceedings of the 19th annual conference on Computer graphics and interactive techniques*, SIGGRAPH ’92, pages 265–272, New York, NY, USA, 1992. ACM.
- [23] A. Weidlich and A. Wilkie. Arbitrarily layered micro-facet surfaces. In *GRAPHITE 2007*, pages 171–178. ACM, Dec. 2007.
- [24] H. B. Westlund and G. W. Meyer. Applying appearance standards to light reflection models. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, SIGGRAPH ’01, pages 501–51., New York, NY, USA, 2001. ACM.

9. APPENDIX

Nouns (excerpt):

Wood	highly diffuse, procedural texture
Marble	highly specular, procedural texture
Aluminium	specular, grey color
Brass	highly specular, golden color

Adjectives (excerpt):

anodize	slightly more diffuse than base material
varnish	less diffuse
brush	diffuse, probably normal map for brush strokes, anisotropic

Adverbs (excerpt):

lightly	the following adjective is only applied lightly
strongly	the following adjective is applied with a higher factor than the default
intensely	
color qualifier	RGB color according
(e.g. blue, yellow)	to HTML color name