

MODELING AND PERCEPTION OF THE APPEARANCE OF  
VIRTUAL HUMANS

CARLOS ALIAGA

SUPERVISORS: DIEGO GUTIÉRREZ & ADRIÁN JARABO

Tesis Doctoral - Ingeniería  
Departamento de Informática e Ingeniería de Sistemas  
Escuela de Ingeniería y Arquitectura  
Universidad de Zaragoza

March 2017



To my parents, my sister and Andrea.



## ABSTRACT

---

Representing virtual humans is very important in computer graphics imagery due to its ubiquitous use in the entertainment industry and their increasing applications in other areas like surgery, psychology, forensics, prosthetic design, fashion or computer vision. However, given the sensitivity of our brain to human depictions, representing virtual humans is still a very challenging problem, with multiple ramifications including geometry modelling, animation, simulation, appearance definition and perception. In this thesis we focus on the later two: modeling and perception of the appearance of virtual humans.

The first part of this thesis aims to understand what makes a character appealing and expressive, depending on the level of stylization of key factors like its shapes and materials. To this end, a set of perceptual experiments are performed over two characters where we combine different levels of stylization -from very realistic to very stylized- for shape and material.

Then, the work focuses in skin and cloth. These two materials are crucial to depict photorealistic humans, and both exhibit a highly complex structures and volumetric properties. We present the first model capable to simulate the changes in the appearance of elderly skin, accounting for the variations in biophysical parameters based on real measurements from medical and tissue optics literature. This predictive nature can broaden the models' range of applications to cosmetics or dermatology.

Regarding the computer generated cloth, a perceptual study is conducted to disentangle the relative importance of dynamics and appearance in the perception of cloth. The goal is to efficiently distribute the resources in production, since both dynamics and appearance simulation are orthogonal processes that require great computational costs and manual intervention by artists and technicians. Last, we present an appearance model for cloth at the fiber level, developing high-quality scattering functions for cloth fibers that rely on real-world physical parameters, such as the shape of the fiber's cross-section or the amount of dye and its absorption rates in the visible spectra. The reflectance fields obtained are specified by actual manufacturing parameters used in the textile industry, what makes the model useful beyond computer graphics, in applications related to cloth design and prototyping.

## MEASURABLE CONTRIBUTIONS

---

This thesis has led to the following results:

- 3 JCR-indexed journal publications [246, 111, 10] (1 of them ACM Transactions on Graphics [246]).
- 1 JCR-indexed journal accepted with revisions [11].
- 1 peer-reviewed conference publication [8].
- FPI Grant from the Regional Government, Diputación General de Aragon (4-year PhD grant).
- 1 Research Stay Grant (CAI Europa).
- 2 research internships (nine months in total) at *Walt Disney Animation Studios* in Burbank, Los Angeles.
- 1 research internship (three months) at *Technicolor* in Rennes, resulting in another peer-reviewed conference publication [62].
- 1 research stay (two months) at *Porto Interactive Center* in *Faculdade de Ciências da Universidade do Porto*.
- 1 patent in progress derived from the internship at *Technicolor*.
- Bronze at *ACM Student Research Competition* in *Siggraph 2016*.
- 5+ invited talks.
- Participation in 3 research projects.
- Reviewer for 3 journals and 4 international conferences, and local committee member of 2 conferences (1 international and 1 national).
- 2 supervised undergraduate students.

## RESUMEN

---

La representación de humanos virtuales es muy importante en el área de la informática gráfica, debido a su uso extensivo en la industria del entretenimiento y sus crecientes aplicaciones en otras áreas como la cirugía, la psicología, las ciencias forenses, el diseño de prótesis y animatrónes, el diseño de moda o la visión por computador. Sin embargo, dada la gran sensibilidad de nuestro cerebro a la hora de percibir representaciones de humanos, la generación de humanos digitales es todavía un problema complejo, con múltiples ramificaciones que incluyen el modelado geométrico, la animación, la simulación, el modelado de la apariencia y la percepción. En esta tesis nos centramos en las dos últimas: el modelado y la percepción de la apariencia de humanos virtuales.

La primera parte de esta tesis se centra en entender los factores que intervienen en el atractivo y la expresividad de un personaje virtual. Con este fin, se han llevado a cabo una serie de experimentos psicofísicos sobre dos personajes que combinan distintos niveles de estilización (desde muy estilizado a muy realista) tanto en forma como en material, dos factores clave a la hora de definir la apariencia.

El resto de la tesis se centra en dos materiales: la piel y la tela. Estos dos materiales son cruciales para representar humanos fotorrealistas, y ambos presentan una estructura compleja y propiedades volumétricas. Se ha desarrollado el primer modelo biofísico de piel que tiene en cuenta los cambios en su apariencia debidos a la edad, mediante el modelado de los cambios estructurales más importantes, como el adelgazamiento progresivo de las capas de la piel, o el descenso en la concentración de cromóforos como la melanina o la hemoglobina. La naturaleza predictiva de este modelo puede hacerlo extensible a otros campos como la cosmética o la dermatología.

Respecto a la apariencia de las prendas digitales, se ha realizado un estudio perceptual para desentrañar la importancia relativa del movimiento y la apariencia de la tela. El objetivo es poder distribuir los recursos de manera eficiente en producción, ya que tanto la apariencia como el movimiento de las prendas se generan en procesos completamente ortogonales, que requieren de grandes costes tanto computacionales (simulaciones y *renderings* físicos) como humanos, en forma de trabajo manual de artistas y técnicos para ajustar los parámetros de dichos procesos.

Por último, la tesis desarrolla un modelo de apariencia de tela a nivel de fibra, a través de funciones de dispersión de las principales fibras comerciales (algodón, seda, poliéster y lana). Dichas funciones se especifican mediante parámetros reales en la industria textil, como la

cantidad y tipo de tinte empleado, lo cual puede extender el uso del modelo más allá de los gráficos en aplicaciones relacionadas con el diseño y prototipado de tejidos.

## ACKNOWLEDGEMENTS

---

This thesis is the result of the support and help of many people.

*Diego*, for teaching me how to improve my research skills and become a better professional. I joined the lab as a complete newbie, not only in computer graphics but as a researcher in general. I specially value the wide view *Diego* has about the field and his extreme attention to detail, what makes you improve everyday. And *Adrián*, for his continuous support and endless discussions, this thesis would not have been possible without his patience and exceptional talent.

*José Antonio Iglesias Guitián*, for his compromise during my Masters Thesis. And his good mood, working with him was trully fun.

My hosts during my internships at Walt Disney Animation. *Rasmus Tamstorf* and *Carol O'Sullivan*, for giving me the first opportunity and teaching me many things. *Matt Chiang* and *Brent Burley*, for the fruitful discussions. Particularly *Matt*, I learned a lot from his skills at solving problems in production. And not only for the work but also for the great times in California and Spain.

My hosts at Technicolor *Tania Pouli* and *Patrick Perez*, for their help and insightful conversations. They made me feel like home.

The members of the Graphics and Imaging Lab. The daily work, the deadlines, the trips, the parties... were awesome with them, with special mention to *José Ignacio Echevarría* and *Elena Garces* for being the closest to me during last year, and for the fun also with *Cristina Tirado*. *Julio Marco*, for the great times in Glendale and his unforgettable black Disney polo, *Adolfo Muñoz* for the early bird caffeinated discussions, and *Susana Castillo* for the extremely late meetings.

My coworkers and colleagues during internships, specially *Yuanfang Chen*, *Simon Kallweitt*, *Rajaditya Mukherjee*, *Ozan Cetinaslan*, *Sylvain Duchene*, *Carlos Castillo* and *Jorge López*. And the crew of spaniards at Disney Animation, with spetial mention to *Luis Labrador* and *José Gómez*.

The students I supervised, *Balma Félez* and *Carlos Guillén*. I hope they learned from me as I did from them.

My friends, for the fun and disconnection from work, and my family, for always caring, especially *Dionisio* and *Manuela*.

My sister, for always trusting me and never let my artistic side die.

*Andrea*, for the infinite good times and support, I owe her part of this thesis.

My parents, for everything. I won't be able to give them back all I got from them.

This thesis has been funded by the Gobierno de Aragón, Walt Disney Animation Studios, the European Research Council (Consolidator Grant, project Chameleon, ref. 682080) and the Spanish Ministry of Economy and Competitiveness (project LIGHT SLICE, ref. TIN2016-78753-P).

# CONTENTS

---

1	INTRODUCTION	1
1.1	Goals and Overview	5
1.2	Effect of Stylization on the Perception of Virtual Characters	6
1.3	Modeling Human Skin	7
1.4	Simulating The Appearance of Cloth	8
1.5	Contributions and Measurable Results	10
2	PERCEPTION OF COMPUTER GENERATED FACES	15
2.1	Introduction	16
2.2	Related work	18
2.3	Stimuli Creation	20
2.4	Experiment Design	22
2.5	Importance of Shape, Material, Shading and Texture	24
2.5.1	Experiment 1a: Shape and Material	24
2.5.2	Experiment 1b: Shading and Lighting	27
2.5.3	Experiment 1c: Texture	29
2.5.4	Conclusion	31
2.6	Experiment 2: Further Investigation of Shape and Material	34
2.7	Experiment 3: Effect of Expressions	37
2.7.1	Intensity and Recognition of Expressions	37
2.7.2	Effect of Expression on Realism and Appeal	39
2.8	Discussion	40
2.8.1	Limitations and Future Work	42
2.9	Appendix A. Effect of Expressions in Experiments 1a and 1b	44
2.10	Appendix B. Additional Diagrams	46
2.10.1	Experiment 1a: <i>Shape and Material</i>	46
2.10.2	Experiment 1b: <i>Shading and Lighting</i>	48
2.10.3	Experiment 2: <i>Further Investigation of Shape and Material</i>	51
2.10.4	Experiment 3: <i>Recognition and Intensity</i>	52
3	A BIOPHYSICALLY-BASED MODEL FOR SKIN AGING	53
3.1	Introduction	53
3.2	Previous Work	54
3.3	Anatomy and Optical Properties of Skin	55
3.3.1	Skin aging	57
3.4	Our Skin Aging Model	58
3.4.1	Layered structure	58
3.4.2	Surface reflection	59
3.4.3	Absorption	61
3.4.4	Scattering	63

3.5	Using Our Model	66
3.6	Results	67
3.6.1	A simplified model	68
3.7	Conclusions and Future Work	70
3.8	Appendix A. Modelling of sebum	73
4	PERCEPTION OF COMPUTER GENERATED CLOTH	75
4.1	Introduction	76
4.2	Related Work	77
4.3	Stimuli Creation	78
4.4	Experiments	84
4.4.1	Experiment One: Ground Truth comparison	84
4.4.2	Experiment Two: Identifying Mismatches	91
4.4.3	Experiment Three: Perceived Cloth Features	93
4.5	Conclusions	96
5	AN APPEARANCE MODEL FOR TEXTILE FIBERS	99
5.1	Introduction	100
5.2	Related Work	101
5.3	Light Scattering from Textile Fibers	104
5.3.1	Fiber properties	107
5.3.2	Obtaining the BCSDFs	108
5.4	Analysis & Discussion	109
5.5	Rendering	112
5.6	Results	114
5.6.1	Validation	116
5.7	Conclusions	118
6	CONCLUSIONS & FUTURE WORK	121
	BIBLIOGRAPHY	127

## LIST OF FIGURES

---

Figure 1	The Uncanny Valley	2
Figure 2	Device for capturing a real human and its real-time rendering	3
Figure 3	Realistic hair rendering and artist-friendly editable models	4
Figure 4	Light and skin interactions for realistic skin rendering	5
Figure 5	The multi-scale nature of cloth	8
Figure 6	High resolution details needed for realistic cloth rendering	9
Figure 7	Examples of the stimuli of the experiments	15
Figure 8	Face scanning setup	20
Figure 9	Examples of stylizations created for the study	22
Figure 10	Stimuli used in Experiment 1a	23
Figure 11	Results of Experiment 1a for realism and appeal	25
Figure 12	Results of Experiment 1a for eeriness and familiarity	26
Figure 13	Stimuli for Experiment 1c	28
Figure 14	Results of Experiment 1c: effect of texture blurring	30
Figure 15	Stimuli for Experiment 2. Shape stylizations and expressions for the male	32
Figure 16	Stimuli for Experiment 2. MAterial stylizations and expressions for the female	33
Figure 17	Results of Experiment 2: realism and appeal for shape and material stylizations	35
Figure 18	Results of Experiment 3: effect of shape on expressivity	38
Figure 19	Results of Experiment 2: expression over realism and appeal	39
Figure 20	Realism vs appeal diagram of participant ratings	41
Figure 21	Results of Expressions in Experiment 1a	44
Figure 22	Results for Experiment 1a	46
Figure 23	Results for Experiment 1a	47
Figure 24	Results for Experiment 1b	48
Figure 25	Results for Experiment 1b	49
Figure 26	Results for Experiment 1b	50
Figure 27	Results for Experiment 2	51
Figure 28	Results for Experiment 3	52

Figure 29	Renders of young and old, female and male faces	53
Figure 30	Skin anatomy	56
Figure 31	Pictures of skin structural changes with age	57
Figure 32	The influence of the hypodermis	59
Figure 33	Evolution of surface roughness and sebum	60
Figure 34	Effects of roughness variation on the specular reflection	60
Figure 35	Spectral absorption for the main chromophores and per skin layer	62
Figure 36	Influence of water on the spectral absorption of skin	63
Figure 37	Influence of the decrease of melanin with age	64
Figure 38	Influence of the decrease of haemoglobin with age	65
Figure 39	Distributions of collagen fibers diameter	66
Figure 40	Skin aging simulations for the six types of skin according to the Fitzpatrick scale	68
Figure 41	Differences on skin translucency and overall color between 30, 55 and 80 years old	69
Figure 42	Aging effect depending on body location	70
Figure 43	Differences when using a simplified, three-layered version	70
Figure 44	Some of the fabrics created for the experiments	75
Figure 45	Comparison between the real fabrics and the CG replicas	80
Figure 46	Lighting studio setup and comparison of real and CG dynamics	81
Figure 47	All CG stimuli from Experiments 1 and 2	83
Figure 48	Experiment 1 screen layout	84
Figure 49	Experiment 1 results	85
Figure 50	Experiment 1 results sorted by Appearance and Dynamics	87
Figure 51	Experiment 1 results, showing the number 'R-fabric' depending on different appearances and dynamics	88
Figure 52	Map of the perceptual distances of the fabrics	90
Figure 53	Two screen layout for Experiment Two	92
Figure 54	Experiment 2 results, exploring mismatch type	92
Figure 55	Frequency of chosen fabrics when isolating factors	94
Figure 56	One screen layout for Experiment 3	95
Figure 57	Results of Experiment 3, perceptual attributes	95
Figure 58	Volumetric rendering of three fabric samples using our model	99

Figure 59	SEM images of textile fibers	104
Figure 60	Coordinates and frame used by the model	104
Figure 61	Light absorption in fibers and threads	108
Figure 62	Comparisons of our BCSDFS and related work	110
Figure 63	Volumetric renders of a knitted stockinette fabric	111
Figure 64	Renders of yarns with different BCSDFS distributions	112
Figure 65	Comparison against related work	114
Figure 66	Volumetric renderings comparing our BCSDFs to related work	115
Figure 67	Renders of the virtual replicas of the acquired yarns	116
Figure 68	Further comparison of our BCSDFS to related work	116
Figure 69	Measurement setup	118
Figure 70	Scattering of yarns; measured vs simulated	119

## LIST OF TABLES

---

Table 1	Typical thickness of the skin layers and their range of variation per decade	59	
Table 2	Chromophore concentrations of a typical adult		62
Table 3	Input parameters for the results of the chapter.	71	
Table 4	Intermediate and output model parameter for all the images of the chapter	72	
Table 5	Contribution of each chromophore to the final spectral absorption coefficient of a given layer	73	
Table 6	Wavelength-dependent absorption equations for each layer of our skin model	74	
Table 7	Manufacturer parameters of the real fabrics		79
Table 8	Outgoing field of the BCSDf of a set of cloth fibers	105	
Table 9	Measurements of the physical and optical properties of four of the most common fabrics		107
Table 10	Computation and storage costs	109	
Table 11	Sewing yarns specifications	117	

## INTRODUCTION

---

Virtual characters are ubiquitous in movies and video games industries, and they are increasingly present in other domains of application such as biomedical simulation or virtual and augmented reality. Digital humans have proven to be useful for phobia treatments and therapies for a range of psychological disorders [235, 175], medicine and surgery [206], cultural heritage [4], and can potentially give birth to new learning and training paradigms [107]. Other applications like design and prototyping of prosthetics and animatronics [27], forensics [140, 29], virtual make-up systems [196], fashion [225], or facial detection in computer vision [197] can benefit from advances in the process of creating computer generated humans. This thesis presents contributions in several aspects relevant to human appearance simulation and perception. We perform psychophysical studies aimed to disentangle the factors involved in the perception of a digital character, and propose several solutions for accurately modeling and rendering humans. In particular, the presented models are focused in skin and cloth, some of the most complex materials found in humans due to their multi-scale structure and their complex interaction with light.

Achieving believable humans that convey the right emotions not only lays on solving the intricated animation of a human face. Modeling and rendering the subtle details and complex materials that are specially present in skin, eyes, hair and cloth is also crucial. Despite the complexity of this process, advances along decades have allowed to create digital replicas indistinguishable from their real counterparts, as in the recent cases of Carrie Fisher and Peter Cushing in *Rogue One* (2016), or Arnold Schwarzenegger in *Terminator Genisys* (2015), to name a few. But it has been a long way since the very first uses of computer generated imagery (CGI) in feature films like *Futureworld* (1976), for depicting an animated hand and a face, or *Tron* (1982), with more than fifteen minutes of digital effects in total. In the case of CG characters, the first virtual actors in movies like *Rendezvous in Montreal* (1987) and the baby of the *Pixar's* short *Tin Toy* (1988) were important milestones, both still very far from photorealism. In 1995, *Toy Story* became the first full-length CG film with compelling characters, thanks to its talented cartoonish stylizations. In the context of realistic depictions, *Star Wars Episode I* (1999) extensively used CGI for crowds and non human creatures, but the resulting look was still synthetic. In this regard, the challenge of creating photo-realistic human actors was courageously accepted by *Square Pictures Studios*

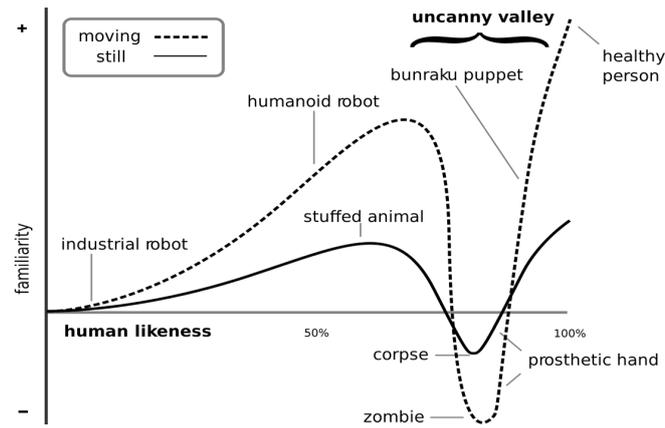


Figure 1: The Uncanny Valley Theory [158] states that a certain point or real- isn, an artificial human starts to fall in a deep valley of disgusting reactions for the audience until it comes out from the valley when perfectly matches reality. Image from Wikipedia Commons.

with the release of *Final Fantasy* (2001). This feature film pushed the boundaries of CG film production by using a huge render farm to create highly detailed characters with high quality geometry, textures and materials (up to 60000 hairs on the head of the main character). However, the achieved look of the characters provoked the rejection of the audience, showcasing how the closer a virtual human gets to a real one, the easier it can lead to disgusting reactions if any subtle detail is wrong or missing. This well known phenomena was first hypothesised by Mori as the *Uncanny Valley* [158], a theory initially developed for humanoid robots but also applicable in other forms of artificial humans. It states that reactions of rejection and discomfort can easily rise in the audience when an artificial human, either manually crafted, manufactured or virtually created, reach a considerable realism but is not able to keep the illusion for the viewers of being in front of a real human (Figure 1). In the particular case of digital characters, advances in the fields of capturing, animation, simulation and rendering have allowed to reach high levels of realism, but it is at such levels when any subtle imperfection or missing detail can trigger negative responses [205], and some variables involved in the perception of virtual faces remain unclear. *The Polar Express* (2004) represents another example of this kind of aesthetic failure, despite the movie introduced the use of motion capture in feature animation. Their combination of stylized characters with captured movements and realistic rendering was not visually pleasant for the audience. Other attempts for more realistic humans suffered from similar problems, like the case of *Beowulf* (2007). In the meanwhile, the use of digital doubles of real actors was introduced by *Matrix Reloaded* (2003) and *Matrix Revolutions* (2003), suitable for certain action scenes and with some limitations. In most cases, either the lack of realism or the rejection from the audience were mainly due to both the capturing techniques



Figure 2: Light Stage used for capturing the appearance of a real human, rendered on the right in real time. Images from *USC Institute for Creative Technologies* and *Activision RD* [7].

employed and the simplifications of the materials used for rendering. The latter is specially critical for complex materials like skin and cloth, what will be further discussed in Sections 1.3 and 1.4. In this regard, a successful example of the use of realistic skin is the character *Gollum* in *The Lord of the Rings: The Fellowship of the Ring* (2001), which became the first virtual actor to win an academy award. With respect to the capturing techniques, later films like *The Curious Case of Benjamin Button* (2008) or *Avatar* (2009) presented facial capture rigs that really made affordable the extensive use of realistic virtual characters in film production (Figure 2).

The aforementioned milestones in film production are based on research work of companies and academic institutions that focus their efforts in human appearance capture and modeling. More precisely, in material modeling, seminal papers like Marschner’s light scattering for hair [149] or Jensen’s subsurface scattering (SSS) [121] (Figure 3) allowed to depict plausible fur in *King Kong* (2005) and realistic translucent skin for the character *Gollum*, respectively (the latter was awarded with a Technical Achievement Award). At the same time, complex acquisition devices were developed aiming to capture every detail and subtleties needed when depicting digital humans, including facial hair or realistic eyes [22, 25, 23, 134], and several tools emerged for stylizing such captured data for later fabrication purposes [63]. In this regard, the canonical example of a capture system is the technology developed around the *Light Stage* [49, 48]: a spherical set-up of light sources and cameras capable of efficiently capturing the face of an actor from any viewing position when lit from every possible direction (Figure 2). This allowed the extensive use of doubles in feature films, like the mentioned case of *Benjamin Button*, the first extended performance of a digitally rendered actor. Thus, within current techniques it is feasible to obtain accurate reproductions of the geometry of the face even at the micro-scale level [163], and skin’s features like its color, spatially varying features or translucency [233, 61]. Then, this high resolution captured data can be combined with both offline and real-time rendering techniques

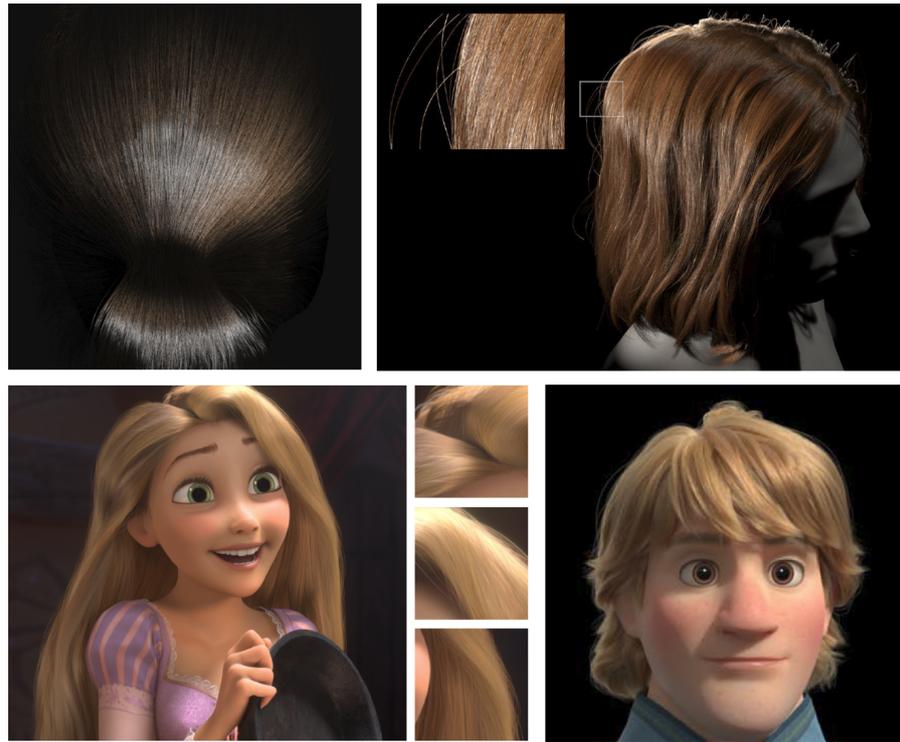


Figure 3: Examples of realistic hair models in the top row [148, 54], and artist friendly and practical models for production pipelines [188, 38]. Models capable to simulate the complex interaction of the light with hair fibers are needed to achieve realistic results. Similarly, part of this thesis deals with modeling the scattering of textile fibers.

[7, 125], allowing to depict photorealistic characters in video-games and virtual reality applications (Figure 4). In addition to skin modeling, complimentary but crucial materials like hair and cloth have also experimented an extraordinary evolution, within simulations even at the yarn level [42]. However, there are many open problems in material modeling despite all these achievements. When creating abstract models that simulate the appearance of real materials, the border between geometry and material becomes very fuzzy, due to the multi-scale structure of most materials present in nature. Current models often rely on simplifications that create a disconnection from reality. Such disconnection is usually solved through skilled manual intervention of artists and technicians, who spend many time and resources changing parameters in a trial and error procedure, so that the desired look is achieved. In this context, this thesis attempts to better understand the perceptual impact of some parameters involved in such processes, and then focuses on building bridges between the models suitable for computer graphics and the reality itself. This is done by proposing solutions that are based on real measurements, reaching a level of detail up to the smallest structural features and



Figure 4: Realistic real-time skin renderings, where the close-up comparison without (bottom right) and within the model (top right) shows how the lack of light and skin interactions ruins the realism of the image (image from Jimenez et al. [126]). Such interactions are further explored along this thesis.

optical properties of each material. It particularly focuses on skin and cloth, two materials that remain challenging to simulate in practice. Sections 1.3 and 1.4 will discuss in more depth the existing work related to modeling and rendering such complex materials and the contributions this thesis presents in this regard.

## 1.1 GOALS AND OVERVIEW

The main goal of this thesis is to develop models for accurately simulating realistic human appearance while providing useful guidelines for the many different options that arise in the process of creating digital characters. Appearance is determined by the physical properties of the objects and materials, and their interaction with light. But human appearance is particularly affected by our visual perception mechanisms, since the brain is especially sensitive when observing other humans. This is a well-known fact traditionally used by artists to exploit perceptual tricks to depict pleasant characters, like stylization, softening or exaggeration. However, this knowledge is not easily transferable to computer-generated humans, a fact that this thesis tries to address during its first part through experimental studies. We subsequently focus on developing appearance models for detailed depictions of skin and cloth, two key components of digital humans whose intricate complexity makes them difficult to model and render accurately. The purpose is to obtain representations of such materials with a direct link to real physical properties, like the variation of chromophore concentrations due to skin aging or the cross-sections of the fibers in the case of cloth, keeping the models tractable in any existing rendering system. In both cases, their predictive nature opens a range of applications beyond computer graphics, such as cosmetics

and dermatology, or fashion design and prototyping.

## 1.2 EFFECT OF STYLIZATION ON THE PERCEPTION OF VIRTUAL CHARACTERS

Chaper 2 of this thesis is devoted to the perception of human faces. The appearance of real world objects is not only determined by factors like their structure at different scales or the properties of the materials they are made of. It also depends on the way the human brain interprets the visual information. The same happens in the case of digitally created content: the eye does not work purely like a camera, but behaves like an extension of the brain. In addition, the brain is particularly sensitive when perceiving other human beings, due to biological adaptation and evolutionary reasons. This fact has been extremely well exploited by artists for centuries; quoting Pablo Picasso: "*Art is the lie that reveals the truth*". From the classics, artists are able to represent mundane real things in a way that they become irresistible and pleasant for the human brain, being able to retain the essence of such things at the same time. Studies demonstrate how humans can recognize abstractions of reality with great distortions and exaggerations better than the reality itself [1], specially when performed by experts capable to keep the most important features to resemble the original real counterpart. This is known as *peak-shift effect*. Caricatures are a great example of how exaggerating salient features enhances the facial recognition process that occurs in our brain.

Knowledge learned from traditional art is hardly applied to digital humans directly. Both 2D and 3D cartoons have traditionally avoided to depict realistic characters, opting for exaggeration and beautification to enhance the expressiveness and appeal of the characters [152]. As in the case of traditional art, stylized versions are preferred over realistic faces [85] when depicting virtual humans in order to avoid the *Uncanny Valley*, as exposed in the introduction. But there are other psychological factors that play a role, like the familiarity of the characters to the viewers [56], or observers' level of training in computer graphics [68] and their sensitivity, increased over the years, to detect the techniques employed for rendering virtual humans [214]. In this context, the effect the rendering style [154], post-production 2D filters [227], or the level of anthropomorphism [205] were already studied. This thesis focuses instead on two key variables when stylizing a 3D virtual character; namely, shape and material. Starting from scans of two real people performing several basic facial expressions, we create different stylization levels for both shape and material with the help of artists, and analyze how different combinations affect the perception of the characters. This work attempts to extract the relative

importance of shape and material stylization, providing interesting insights for future research in perception of human faces and also useful guidelines for generating appealing 3D characters.

### 1.3 MODELING HUMAN SKIN

Chapter 3 focuses on modeling the appearance of the skin, a central component of the human appearance. Thanks to modern 3D scanning devices, geometric and texture details such as wrinkles, scars, pores or freckles can be accurately captured, even the micro-geometry of the skin surface. But naively rendering these highly detailed models results in a very hard looking and unrealistic skin. Examples previously cited like *Final Fantasy* or *The Polar Express* films showcase how opaque materials used to mimic the skin appearance easily trigger negative reactions in the audience. This is due to the absence of *sub-surface scattering* (SSS): the skin is made of several translucent layers so that when a beam of light reaches the surface, it scatters, gets partially absorbed, and emerges from the skin back to the exterior at a different point than where it entered. The effect of such process is subtle but turns out to be crucial for rendering realistic humans. Several techniques [50, 98] have been presented to model SSS, since the first approaches [121, 120] that introduced the use of diffusion approximation in computer graphics. Some of them oriented to real-time rendering [51, 124, 125], some other capable to simulate the effect of different emotional and physical states on the appearance of the skin [123]. Another line of research was focused on the multilayered structure of the skin and its composition. The color of the skin is mostly due to the chemical compounds, called chromophores, present in different proportions for each of the layers of the skin. Mainly melanin in the outermost layers (epidermis) and blood in the inner ones (dermis). There are approaches that deal with multilayered translucent materials like skin [59, 52], and some models account for chromophore concentrations in a biophysically-based way [60, 135, 18]. Starting from an accurate static biophysical skin model, part of this thesis focuses on simulating the changes in the optical properties of the skin due to aging. In computer graphics, aging has only been addressed from the geometric perspective, such as the synthesis and simulation of wrinkles [28, 92]. Our model relies on the existing measurements from the vast medical and tissue optics literature to account for key factors like the slimming of the skin layers and other structural changes, or the variations in the chromophore concentrations. Together, all these changes modify the absorption and light scattering patterns of the skin, leading to different appearances that ultimately depend on age, gender, or life habits.

## 1.4 SIMULATING THE APPEARANCE OF CLOTH

Chaper 4 and Chaper 5 are centered on the perception and modeling of computer generated garments. In computer graphics, cloth remains as one of the most challenging materials due to the microstructures found at the fiber level and the complex light-material interactions at such scales (Figure 5). To obtain realistic animations of cloth requires a lot of computational resources in the form of costly physical simulations and renderings, and also a great amount of work of skilled artists and technical directors to manually tune the dynamics simulation and rendering parameters.

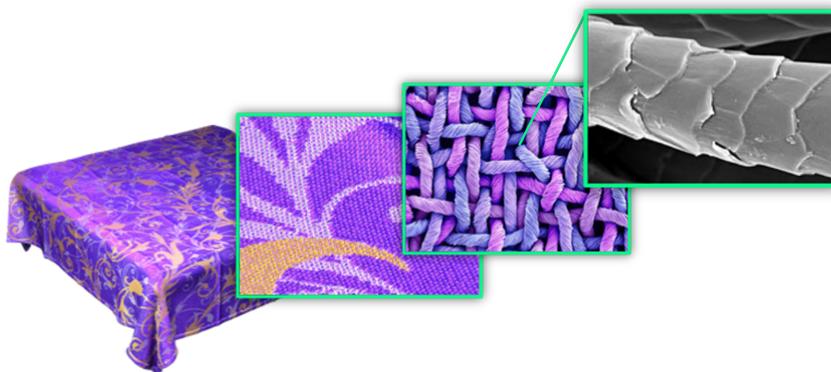


Figure 5: The overall appearance of cloth results from the aggregate effect of each fiber interaction with light at micro-scale level. Image modified from Zhao et al. [250].

**PERCEPTION OF CLOTH** In this scenario, a deeper understanding on how audience perceive digital cloth can help to balance the resources and decide whether to direct more or less efforts. A goal of this thesis is to gain knowledge about the relative importance of dynamics and appearance on the perception of virtual cloth. Both represent orthogonal processes in practice that require independent treatment. On the one hand, to achieve the desired movements depends on a trial and error procedure where skilled specialists are in charge of testing combinations of many parameters over costly dynamics simulations. Some of these parameters are not even linked to physical properties and are particular to each simulation engine, a fact that was addressed by recent approaches [208] in an attempt to build a generic control space for cloth simulation. On the other hand, rendering cloth is a very challenging task, due to the complexity of textile materials and the amount and heterogeneity of the available models and representations in computer graphics, what will be discussed later. In fact, cloth appearance has traditionally relied on the skills of artists to mimic the desired look, what makes the task very time consuming. In this thesis we present a study of the rela-

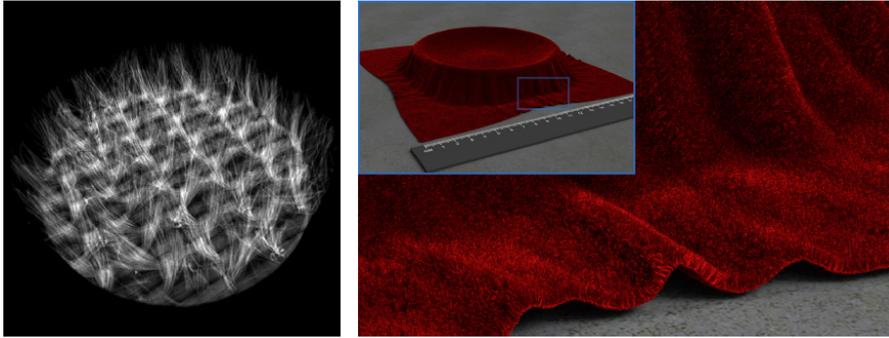


Figure 6: Recent cloth models reach a great level of detail for very small scales, some of them relying in x-ray computed tomography scanners to acquire high resolution volumetric samples. Such small details are crucial for achieving a realistic look of the cloth, and our work goes further in this direction. Image from Zhao et al. [248].

tive importance of cloth appearance and dynamics through several perceptual experiments to disentangle the effect of each factor under different conditions, what can help to efficiently distribute human and computational resources in production pipelines.

**MODELING CLOTH** The complexity of cloth modeling and rendering has been traditionally addressed by taking big assumptions and simplistic models while relying on the skills of shading artists. Taking again *Final Fantasy* as an example, some clothing in that movie was rendered using plastic and metal shaders, and even current productions mimic cloth appearance through a combination of textures and materials stacked in several shading layers. This is due to the difficulties involved to accurately model and render cloth, a hard problem due to its multi-scale structure and the complex interaction of the light with its smallest building block: the textile fibers. In fact, cloth rendering is a very active research field that has received a lot of attention during the past few years. Either by modeling the fiber assemblies [199] or capturing pieces of cloth through photographs [203] or Computed Tomography Scanners (CT) [247, 251], recent approaches represent a step further to accurately represent the cloth appearance at very small scales (see Figure 6). This is needed for capturing the subtleties and rich optical behavior of individual fibers, and determines the overall appearance of the cloth at coarser scales. Despite the extreme level of detail in terms of geometry, the models [255, 130] used to simulate the light scattering from textile fibers suffer from two key drawbacks. First, they over simplify the fibers' structure and resultant scattering patterns (e.g. assume circular or elliptical cross-sections). Second, current models lack of any connection with physical properties and manufacturing parameters used in reality, like the role of dyes and their compounds in the absorption of light and derived saturation of dyed cloth. The last part of this

this thesis is dedicated to develop high-quality scattering functions for cloth fibers capable to take into account the optical and structural features of real textile fibers. To this end, the proposed model relies on real-world measured data available in the textile research literature to build digital replicas of real fibers. Then, fiber scattering functions for different fibers (polyester, cotton, silk, wool) are obtained through simulations, so that such functions are specified by actual fabrication parameters used by manufacturers in real life. This makes the model not only relevant for the obvious application of CG films and videogames, but also for cloth design and prototyping.

## 1.5 CONTRIBUTIONS AND MEASURABLE RESULTS

### 1.5.0.1 Publications

Some part of the work presented in this thesis has been already published and some is under review. In particular:

- *Effect of Stylization on the Perception of CG Faces*. This work was accepted in SIGGRAPH ASIA 2015, and published in ACM Transactions on Graphics [246]. This journal has an impact factor of 4.218, and its position in the JCR index is 1st of 106 (Q1) in the category of Computer Science, Software Engineering (data from 2015). In this work, I led the stimuli creation and co-led the design of the experiments; the first author was in charge of the analysis of the results.
- *A Biophysically-Based Model of the Optical Properties of Skin Aging*. The work on skin modeling was accepted in EUROGRAPHICS 2015, and published in Compute Graphics Forum [111]. This journal has an impact factor of 1.542, and its position in the JCR index is 17 out of 106 (Q1) in the category of Computer Science, Software Engineering (data from 2015). The project was led by José Antonio Iglesias Guitián. My role was to lead the research over the medical and tissue optics literature needed to build the model, being also in charge of the rendering part in PBRT [179] to create the results of the paper using the diffusion profiles obtained from the model presented.
- *Cloth Perception*. The perceptual study of the digital moving garments led to a publication in the Symposium on Applied Perception 2015 [8] and is accepted with major revisions in ACM Transactions on Applied Perception (TAP) [11]. This journal has an impact factor of 0.561, and its position in the JCR index is 86 out of 106 (Q4) in the category of Computer Science, Software Engineering (data from 2015). I was the leading author, and the work was partially done during my first internship of

six months at Walt Disney Animation Studios, and the remaining work was done at Universidad de Zaragoza with the advise of Diego Gutiérrez.

- *Cloth Appearance Modeling*. The work devoted to modeling the appearance of cloth, where I had the role of leading author, resulted in several publications. First, a poster of the ongoing project presented at SIGGRAPH 2016 obtained the 3rd place at the Student Research Competition. Then, a full paper was accepted in EUROGRAPHICS Symposium on Rendering [10], whose proceedings are published in Computer Graphics Forum. This journal has an impact factor of 1.542, and its position in the JCR index is 17 out of 106 (Q1) in the category of Computer Science, Software Engineering (data from 2015).

In addition to these previous publications, during my PhD I have collaborated in other research projects directly or indirectly related to the topic of this thesis:

- *Display Adaptive Disparity Remapping*. I participated in a project related to glasses free 3D content retargeting. It was focused in remapping the depth of stereoscopic 3D scenes to fit different displays while maintaining the perceived original depth. This work was accepted to Computers & Graphics [150]. This journal has an impact factor of 1.12, and its position in the JCR index is 41st of 106 (Q2) in the category of Computer Science, Software Engineering (data from 2015).
- *Art Directable Micro-Appearance Modeling of Cloth*. My second internship at Walt Disney Animation Studios was devoted to build solutions for cloth appearance with enough geometric detail at fiber level but scalable at the same time. Two talks were given in the company, interested in taking the ideas developed during the internship for integration in the production pipeline.
- *Mixed Illumination Analysis in Single Image for Color Grading*. During my internship at Technicolor, I joined an ongoing project of mixed illumination estimation from single images, led by Sylvain Duchene. I focused on re-designing the clustering of pairs of radiance in the image in terms of the different hues of light present in the scene, and designing the strategy for efficiently sampling such pairs near the occlusion boundaries. The former technique is going to be patented by Technicolor, and the complete work was published at Expressive 2017 conference [62].

#### 1.5.0.2 Grants & Awards

We include here a list of awards and fellowships received throughout this thesis, that have allowed the realization of the work here presented:

- FPI Grant from the Regional Government, Diputación General de Aragón (4-year PhD grant).
- CAI Europa Research Stay grant for the 2 months at Porto Interactive Center.
- Funding from Disney Research to extend the collaborative work after the first internship in 2014.
- Bronze at the ACM Student Research Competition, for the poster published at SIGGRAPH 2016 [9].

#### 1.5.0.3 *Research Stays and Internships*

Two research stays, totaling 14 months, were carried out during this PhD in three different locations:

- July 2012 - August 2014 (two months): Early Stage Researcher within the GOLEM Marie Curie project at Porto Interactive Center (Portugal). Supervisor: Veronica Orvalho. Worked on realistic skin real-time rendering.
- June 2014 - December 2014 (six months): Research Intern at Walt Disney Animation Studios. Supervisors: Rasmus Tamstorf and Carol O’Sullivan. Worked on the perception of virtual cloth. A publication in a peer-reviewer international conference [8] and its extension to Transactions on Applied Perception (Chapter 4), accepted with revisions, result from this internship and further collaboration.
- June 2016 - September 2016 (three months): Research Intern at Walt Disney Animation Studios. Supervisors: Matt Chiang and Brent Burley. Worked on developing a scalable appearance model for cloth at fiber level.
- September 2016 - December 2016 (three months): Research Intern at Technicolor. Supervisors: Tania Pouli and Patrick Perez. Worked on estimating the illuminance of a scene under mixed illumination from a single image. The results were published in a conference paper [62].

#### 1.5.0.4 *Supervised Students*

During the development of this thesis I have supervised the Graduate Thesis of two students:

- Carlos Guillén (Computer Engineering, 2015). Realistic Hair Rendering.
- Balma Felez (Industrial Design, 2013). A Virtual Recreation of Asch Psychological Experiment.

#### 1.5.0.5 *Research Projects*

During my PhD studies I have participated in the following research projects:

- **GOLEM: Realistic Virtual Humans.** European Commission Marie Curie Industry–Academia Program, Seventh Framework. Grant no.: 251415. PI: Diego Gutierrez.
- **VERVE: Vanquishing fear and apathy through e-inclusion: personalised and populated realistic virtual environments for clinical, home and mobile platforms.** European Commission (FP7-ICT-2011-7). Grant no.: 288914. PI (in Spain): Diego Gutierrez.
- **LIGHTSLICE: Capture, analysis and applications of the multidimensional light transport (application to medical imaging).** Ministerio Español de Economía y Competitividad. PI: Diego Gutierrez.

#### 1.5.0.6 *Professional service*

During this thesis I had the chance to be reviewer for several international journals and conferences such as Computer Graphics Forum, Pacific Graphics, Computers & Graphics, or Transactions on Applied Perception, and I was on the local organizing committee for the Eurographics Symposium on Rendering (EGSR) 2013 and the Spanish Conference in Computer Graphics (CEIG) 2014, both held in Zaragoza and hosted by our group.



PERCEPTION OF COMPUTER GENERATED FACES

---

Virtual characters contribute strongly to the entire visuals of 3D animated films. However, designing believable characters remains a challenging task. Artists rely on stylization to increase appeal or expressivity, exaggerating or softening specific features. In this chapter we analyze two of the most influential factors that define how a character looks: shape and material. With the help of artists, we design a set of carefully crafted stimuli consisting of different stylization levels for both parameters, and analyze how different combinations affect the perceived realism, appeal, eeriness, and familiarity of the characters. Moreover, we additionally investigate how this affects the perceived intensity of different facial expressions (sadness, anger, happiness, and surprise). Our experiments reveal that shape is the dominant factor when rating realism and expression intensity, while material is the key component for appeal. Furthermore our results show that realism alone is a bad predictor for appeal, eeriness, or attractiveness.

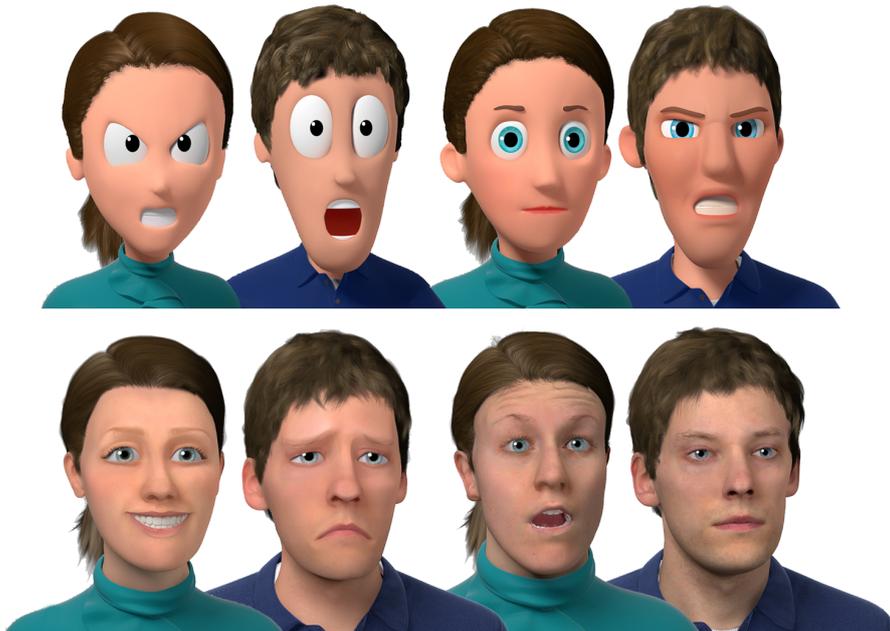


Figure 7: Some of the stimuli used in our experiments. From most abstract on the top left to the most realistic versions on the top right, different expressions of four of the five levels of stylization used in the study are shown in the image.

## 2.1 INTRODUCTION

Over the last years, advances in the field of computer graphics have allowed the entertainment industry to create very realistic virtual humans [6, 126]. However, depicting convincing and believable characters continues to be a difficult task. The reasons why people like the appearance of a particular character may be influenced by many psychological factors, such as the familiarity of the characters to the observers [56], the possible adaptation of the society to cartoon faces [37], the viewers' level of expertise in computer graphics [68], or their increasing ability to notice the tricks and techniques employed [214]. In this context, it is generally accepted that stylized versions are often preferred over realism [85], but the reasons for this are still unclear. Moreover, stylization allows the artist to explore possibilities beyond what is found in the real world (such as oversized eyes for cartoons, for instance) to enhance the appeal or expressiveness of the characters, particularly in the case of cartoons [152] or illustrations [93].

The challenge is to understand and translate the knowledge of artists into feasible guidelines for generating appealing 3D virtual characters. Some works have investigated the role of few of the many variables involved in the creation of 3D faces, such as the influence of rendering style [154], anthropomorphism [205], or applying different 2D filters on the images [227].

In our work, we focus instead on two of the main aspects that primarily define the stylization of a 3D character: shape and material (including texture and optical properties). Due to the high dimensional nature of the problem, experiments were performed in two rounds. We first analyze which of the many sub-dimensions of both shape and material affect the appearance of the character the most. To this end, we define three different levels of stylization along shape and material for a single male character: a realistic head obtained by state-of-the-art 3D-scanning and two stylized versions designed by artists. Moreover, each level includes five different facial expressions: anger, happiness, sadness, surprise, and neutral. We then create all combinations along these dimensions (shape, material, expression) and analyze the perceived realism, appeal, eeriness, and familiarity of each character by means of perceptual studies. We also analyze the effect of decoupling the material dimension into its main components, testing two different shaders, three illumination methods, and three progressively blurred albedo textures. Results and acquired knowledge from these tests are then used to guide a second round of experiments, where we deeply explore the space with more samples along the factors found as most important in the previous studies. For this, we substantially increase the stimuli to two characters (male and female), five stylization levels (of both shape and material), and five expressions. We then analyze the most significant scales of the

previous experiments (realism, appeal), and also evaluate how the combination of each of these dimensions affects the expressivity of the characters.

Our experiment design is inspired and justified by the current trends in feature animation, which have recently used different combinations of stylized shapes and materials to depict 3D characters. Examples include highly stylized shapes and textures in Pixar’s *Toy Story* movies, or the somewhat less stylized shapes but photo-realistic materials in *The Adventures of Tintin*. Furthermore, in a review of recent advances in facial appearance capture, Klehm and colleagues [133] mention the need for deeper insights into human perception of facial appearance. They note the complexity and the importance of focusing on important features, which we address by carefully isolating the effects of the parameters being studied in each test. We use static pictures as stimuli, as it has been found that much of the information that people use to evaluate virtual characters is available in a still image [154].

Our main findings are:

- Shape is the key attribute for perceived realism. Stylized materials decrease the perceived level of realism for realistic shapes; however, realistic materials do not increase realism of stylized shapes.
- On the other hand, appeal, eeriness, and attractiveness are mainly affected by the stylization level of material, and not shape; realistic materials reduce appeal in general. Within the materials, the albedo texture is the dominant factor.
- The degree of realism is a bad predictor for appeal or eeriness.
- The perceived intensity of expressions decreases with realism of shape, but is nearly independent of material stylization.
- Our results are consistent across all tested expressions. Only the anger expression has been constantly perceived as less appealing and more eerie.

To our knowledge, this is the first work attempting to evaluate how the combination of different levels of stylization in shape and material affect the perception of a virtual character across different expressions. Conclusions from this study are restricted to the stimuli specifically created for these experiments, but can help to provide useful insights on how to create believable characters. To foster further research, we will make our stimuli and data publicly available at: <http://graphics.uni-bielefeld.de/publications/sigasia2015/>

## 2.2 RELATED WORK

Creating computer-generated (CG) characters is a challenging task, since even small imperfections can trigger negative responses [205]. Many studies have investigated this effect in the context of the theory of the uncanny valley [158], but the variables involved in the perception of virtual faces are still mostly unknown. For instance, Dill et al. [56] evaluated CG characters on still images and videos, concluding that people often prefer familiar faces, while other studies analyzed how knowledge of computer graphics affects 3D character perception [217, 68]. For brevity, we refer the reader to the excellent discussion on this topic by Tinwell [215], and focus on works closer to our particular goal.

**STYLIZATION** Some of the first attempts to measure the likability of stylized and realistic characters was performed by morphing pictures [101, 144, 198, 56]. Schneider et al. [198] studied the effect of stylization on characters in Japanese video games, and found that it increased perceived attractiveness. All these studies used different characters, including confounding factors such as changing lighting and background. In contrast, we investigate the effects of stylization on the same character under identical conditions. Other works have focused on changing certain features and modifying proportions in the shape of digital faces. It has been shown that uncanniness emerges when abnormal features of the face become apparent for highly realistic characters [205, 35]. Green et al. [96] concluded that there is less tolerance to deviations from original proportions in the cases where faces are more attractive and human-like. Another existing work uses an Electroencephalograph to evaluate the perception of different categories of human faces with varying degrees of realism [161], and a recent study [80] investigates the perceptual effects of stylization over 3D body scans of female avatars. Also, a personality trait exaggeration system was developed for emphasizing the impression of human face in images, based on multi-level features learning and exaggeration [213]. Different from these works, we investigate the effects of global stylizations as commonly adopted by the animation industry.

Wallraven et al. [227] studied the perceived realism, recognition, sincerity, and aesthetics of real and computer-generated facial expressions using 2D filters to provide brush, cartoon, and illustration styles. They concluded that realistic depictions improve subjective certainty about the conveyed expression. Later, they evaluated the perceptual realism of computer-generated faces under progressively blurred normal vectors and textures, finding no effect with their setup [228]. A recent work also explores the personality perception of abstract virtual faces [73]. In contrast to the first examples, we do not employ Gaussian blurring for producing abstract stimuli, but instead use stylized

models produced by artists, in order to better match the character styles used in industry.

MacDorman et al. [145] showed participants several images of virtual faces, combining different textures (from realistic to simple lines) with *geometric* levels of detail (i.e. decreasing polygon counts). Results suggested that decreasing photo-realism can make the face look less eerie and more attractive. In our work, *shape* refers to the global, high-level features of the face, not to technical aspects such as polygon count. Closer to our goal, the recent study by McDonnell and colleagues [154] found that rendering style affects the appeal and trustworthiness of the characters. Additionally, a character rendered in an appealing style can be perceived to have more positive personality traits [254]. Recent studies focusing on neurocognitive mechanisms attribute negative appeal ratings to the difficulty of categorizing images in a particular category, resulting in competing visual-category representations during recognition [72]. Negative effects for such images occurs to the extent that selecting one interpretation over the other requires inhibition of the visual-category information associated with the non-selected interpretation. Following the conclusions from these studies, stylization affects pleasantness ratings and furthermore, some combinations of visual elements might result in negative effect. Therefore, we study the effects of combining different levels of stylization for shape and material, which are the two key parameters governing visual appearance.

**SKIN APPEARANCE** Taking into account previous work related to the perception of human skin appearance helps understanding effects of material stylizations. Many studies about attractiveness of human faces merged different photographs to achieve average appearance. There was speculation that this technique impacts ratings of attractiveness not just because it averages the shape, but also because it removes blemishes and other skin irregularities [12]. Several studies confirmed that texture changes do result in a significantly more attractive face [24, 138]. Publications in the cosmetics domain also help explain the observed effects on appeal: Fink and colleagues [78] created textures from photographs of women of different age and evaluated these textures on a single female virtual character. Renderings with pure skin have been rated as younger and more attractive than renderings with strong variations in skin pigmentation. This observation was confirmed in a follow-up study [77], which showed that blurring the skin texture can increase attractiveness. Similar suggestions can be found in many photograph retouching books (e.g. [167]).

**EXPRESSION** We are also taking into account the influence of the particular expression. Brain studies show that some areas in the brain respond differently to certain expressions of emotion, specifically the

amygdala, which tends to activate while viewing fearful and angry faces, as opposed to happy, surprised, and sad faces [36]. Since the amygdala region is activated in response to danger, it is believed that negative emotional expressions, such as anger and fear, trigger a defense response in the perceiver. Another example comes from studying the “uncanny valley” effect on CG characters, where modified expressions of emotion with negative valence (e.g. anger, sadness) increased the perceived uncanniness of the character [216]. A recent study explores the effect of stylization on the emotion recognition in autism spectrum disorder [209]. Additionally, given different hypotheses that iconic representation of faces increase the expressibility and the recognizability of expressions [152], we further analyze our stylization domain by evaluating whether different levels of stylization in shape and material, including mismatches between them, affect these scales.

### 2.3 STIMULI CREATION

Our initial experiments required the design of three levels of stylization of the same character. Additionally, for each stylization level we modeled four of the universal facial emotions: anger, happiness, sadness, and surprise [67], plus a neutral expression. We discarded disgust and fear because their status as basic expression was questioned recently [114] and they are harder to identify by observers.



Figure 8: Our face scanning setup (right) and comparison between photographs and virtual reconstructions of our actor (left).

Our realistic characters are based on real people of about average attractiveness without ethnic bias to the group of participants. To generate the realistic models we replicated the multiview-stereo face scanner of Beeler et al. [21], which reconstructs high-resolution textured point clouds from the photographs of six cameras arranged as pairs around a person (Figure 8). Since all photographs are taken simultaneously, the scanning process is instantaneous and therefore well

suited for capturing different facial expressions. Each pose representing one emotion was captured several times, and the most convincing one was selected by a group of about twenty people of different cultural backgrounds, while referring back to Ekman’s guidelines. For the stylized shapes, we did not intend for the artist to precisely match the emotional intensity across the shapes, but rather to create expressions that resembled the expressions of our scanned actors to the best of their ability (e.g., teeth showing slightly in a happy smile) given the available facial features.

Since the scanner only captures the frontal part of the face and fails to faithfully reconstruct eyes and hair, we fit a template head model to the measured point cloud using a non-rigid registration approach similar to Weise et al. [231]. Regions of missing data are therefore filled in by the template model, which additionally provides a 2D parameterization of the model. This parameterization is used for texture mapping, with texture images being generated automatically from the photographs. The hair style, the eyes, and the teeth were manually sculpted and adjusted to fit the scanned model. Figure 8 shows one example of our reconstructed models.

While a realistic character can be obtained from 3D scans of a real person, no automatic solution exists to generate increasingly stylized versions. Therefore professional 3D artists produced the required stylized shapes and materials from our realistic characters, taking inspiration from commercial animation films (see examples in Figure 9). For our first set of experiments (Section 2.5) we used three stylization levels for shape and material (see Figure 48). The extended stimuli for the later experiments (Section 2.6) used two more stylization levels (see Figure 15 and Figure 16).

We are interested in analyzing the effect and interaction of shape, material, and textures. Therefore, we transferred all material properties of the baseline characters to the other character shapes (Figure 48). The inter-surface mapping for the texture transfer was computed based on a dense correspondence map established using the non-rigid registration technique of Zell and Botsch [245].

Rendering of all stimuli was performed using Mental Ray, with each character being placed in front of a light gray background. The lighting setup consists of a key light and a rim light, and photon mapping is used for global illumination. For the skin all three characters use the same multi-layer skin shader with subsurface scattering, with diffuse albedo specified by a high-resolution texture map. The shader parameters vary between the models in order to closely resemble the targeted render styles (Figure 9).

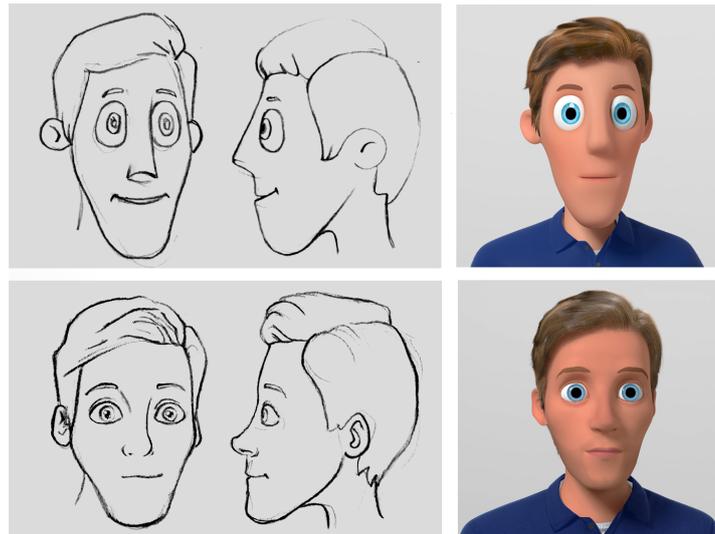


Figure 9: Two of the stylizations created for the study, showing the sketches provided to the artists on the left and their resulting stylized 3D models on the right. The designs are inspired by the films *Cloudy with a Chance of Meatballs* (top) and *Toy Story* (bottom).

#### 2.4 EXPERIMENT DESIGN

The appearance of virtual humans is a function defined over a huge multi-dimensional space. While it is generally recognized that *shape* and *material* are the main contributors to the overall appearance of virtual characters, these two might be affected by several sub-dimensions. For example, material is the combination of shader, shader parameters, and textures, each of which having a potentially different influence on appearance. This makes the experiment design an extremely difficult task, given the large number of variables to explore.

Similar to previous work on rendering style [154], we want to analyze how different levels of stylization (e.g., shape and material) change the perception of a virtual character. Following previous work (e.g., [154, 106]) we employ (subsets of) the following scales for our experiments. The descriptions below are the ones given to the participants of the perception studies:

- *Extremely unappealing—Extremely appealing*: High appeal means that the character is one that is pleasant and you would like to watch more of. Unappealing means that you dislike to watch the character.
- *Extremely eerie—Extremely re-assuring*: Indicate if you find the character eerie, which means that they are gloomy and leave you with a sense of fear. Re-assuring means that the character restores a sense of security, confidence, calm in you.

- *Extremely abstract—Extremely realistic*: Indicate if you find the character’s appearance to be highly stylized like in cartoons, or close to photo-realistic as in real pictures.
- *Extremely unfamiliar—Extremely familiar*: Indicate if you find the character’s appearance familiar to you, in that you have seen something similar to it before, or if you find the character unfamiliar with an appearance that you haven’t seen anything like before.
- *Extremely unattractive—Extremely attractive*: Indicate whether you find the character unattractive and ugly or beautiful and attractive.

We model these properties as Likert scales, which are popular in psychology as they allow subjective conditions such as the attitudes of participants to be measured. We chose a seven-point scale in order to give participants more response options and to allow for comparison to previous studies. The Likert scales were numbered 1–7, with a description provided on both ends of the scale.

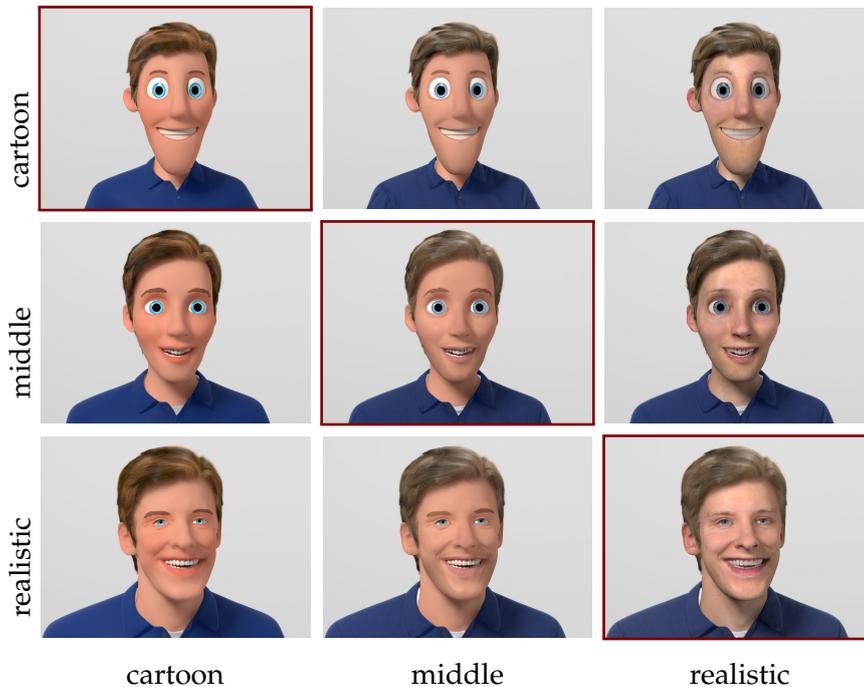


Figure 10: Stimuli used in Experiment 1a: three levels of shape and material stylization, shown here for the *happy* expression. The baseline stimuli are shown on the diagonal. Their textures have been transferred to the other shapes for producing the off-diagonal stimuli of mismatching stylization levels for shape and material. Please refer to the supplemental material for the full set of stimuli used in the experiment.

Since both the design and the analysis of our experiments share many similarities, we describe the general setup now and later only

mention deviations. The user’s task and the rating scales were explained on a written document to the participants before the experiment. Afterwards all stimuli were presented in a random order and shown for 3 seconds each. The display was calibrated, 20" wide and at about 50 cm distance from the participants. The renderings have a resolution of  $1024 \times 768$ , corresponding to approximately  $26.5 \text{ cm} \times 20.0 \text{ cm}$  on screen. After each stimulus presentation, participants were asked to rate it according to the above scales. In all experiments, the participants had normal or corrected-to-normal vision and were unaware of the final goal of the experiment. They were asked to report their 3D experience (how often they played video games, watched movies with visual effects, and how they would consider their knowledge of 3D graphics). We did not find any correlation between the reported 3D experience and the results of our tests, and thus omit this information for the rest of the Chapter.

For statistical analysis of each rating scale we conducted an n-way repeated measures Analysis Of Variance (rm-ANOVA). We run Mauchly’s test for validating sphericity of the data, and whenever it is significant we report results with Greenhouse-Geisser correction applied and marked with an asterisk “\*”. Whenever main interaction effects were found, we conducted a Tukey Honestly Significant Difference (HSD) test for the comparison of means to further explore the results [45].

## 2.5 IMPORTANCE OF SHAPE, MATERIAL, SHADING AND TEXTURE

The goal to investigate the influence of shape and material independently of the overall appearance of a CG character is motivated by differing design choices of recent animation films, ranging from cartoon shapes with cartoon materials (e.g., *Despicable Me*), to stylized shapes with realistic material (e.g., *The Adventures of Tintin*), to very realistic shapes and material (e.g., *Beowulf*). From a detailed analysis of character designs in commercial animation we identified three different recurrent stylization levels, which we denote by *cartoon*, *middle*, and *realistic*, where *Cloudy with a Chance of Meatballs* and *Toy Story* act as references for the two stylized versions, respectively.

### 2.5.1 Experiment 1a: Shape and Material

We first investigate the influence of shape and material, where we denote by material the combination of shader, shader parameters, and textures. The combination of each material with each shape style leads to a total of nine different versions of the character, times five different expressions, resulting in a set of 45 stimuli. Figure 48 shows the  $3 \times 3$  stimuli for the happy expression. We analyze the interaction between shape and material for the scales most frequently used in

previous work: realism, appeal, reassurance, and familiarity. Twenty-two volunteers participated in this first experiment: 14 female, 8 male, with age from 19 to 30 years (avg. 24.5).

In this section, we analyze the effects of shape and material only. Figure 11 and Figure 12 compare the ratings of the neutral expression with averaged ratings over all expressions. Despite a smaller offset and some noise, ratings for different expressions have been very consistent, which justifies averaging over all expressions. For statistical analysis, a rm-ANOVA with three factors (shape, material, and expression) was used.

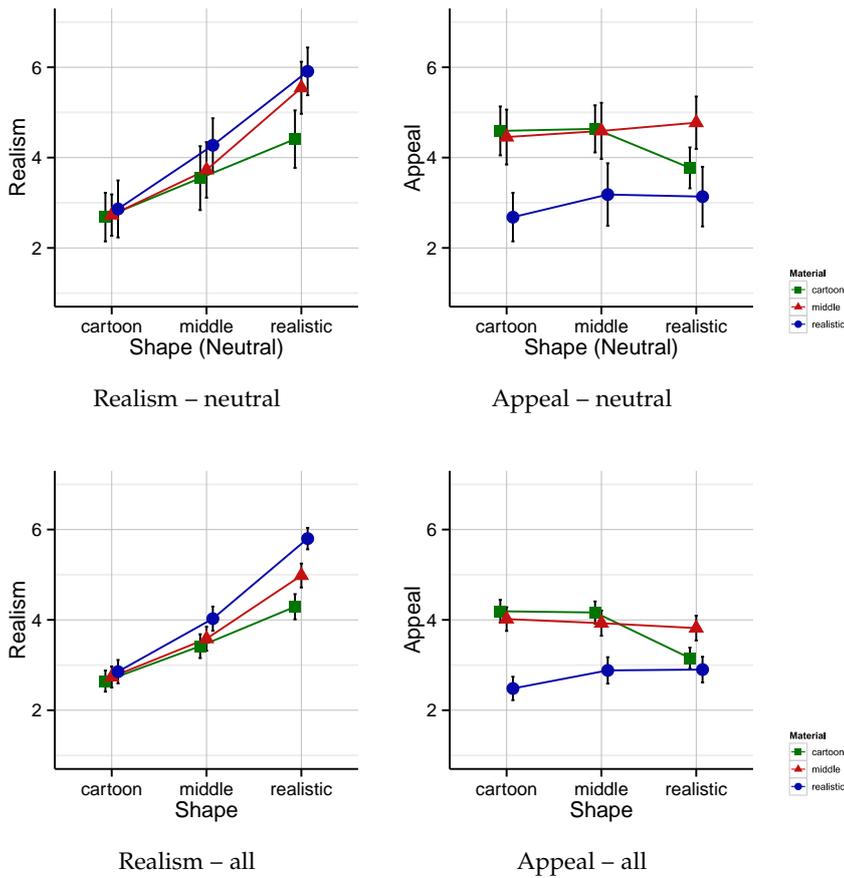


Figure 11: Results of Experiment 1a: Ratings for perceived realism and appeal for different shape and material stylizations. Top row show results for the neutral expression only, bottom row are averages over all expressions. Error-bars denote 95% confidence levels. Individual per-expression results are discussed in Section 2.9.

**REALISM** A main effect was found for shape ( $F(2,42) = 113.18$ ,  $p < 0.0001$ ) and material ( $F^*(1.47, 30.82) = 23.15$ ,  $p < 0.0001$ ,  $\epsilon = 0.734$ ), as well as for the interaction between shape and material ( $F(4, 84) = 11.14$ ,  $p < 0.0001$ ). Post-hoc tests show that the cartoon shape was perceived as least realistic, no matter which material was

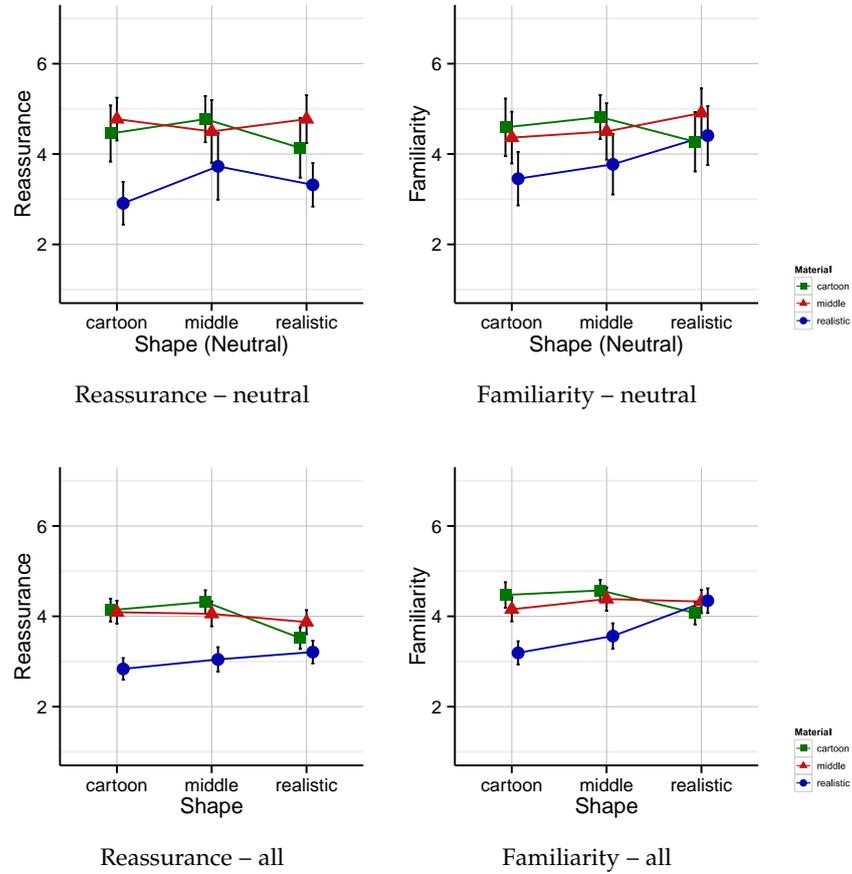


Figure 12: Results of Experiment 1a: Ratings for perceived eeriness and familiarity for different shape and material stylizations. Top row show results for the neutral expression only, bottom row are averages over all expressions. Error-bars denote 95% confidence levels. Individual per-expression results are shown in the supplementary material and discussed in Section 2.9.

used. Similarly, cartoon and middle materials did not make a difference for the middle shape (Figure 11 left), while the realistic material caused a more realistic perception for this shape ( $p < 0.002$  for both comparisons). In contrast, all material levels differ significantly for the realistic shape ( $p < 0.0002$ ). Interestingly, the most stylized shape does not reach the bottom of the realism scale, revealing that there is more potential for abstraction.

**APPEAL** We found a main effect of material on the ratings of appeal ( $F^*(1.41, 29.67) = 42.69$ ,  $p < 0.0001$ ,  $\epsilon = 0.706$ ), but no main effect of shape was found. An interaction between shape and material ( $F(4, 84) = 13.97$ ,  $p < 0.0001$ ) shows that a realistic material on a cartoon shape yields the least appealing combination, since a post-hoc analysis showed significantly lower ratings for this combination compared to all others ( $p < 0.02$  in all cases). The realistic material

is less favored on the middle shape as well, and the cartoon material on the realistic shape is similarly unappealing ( $p < 0.02$  in all cases except the combinations mentioned above). These results (see Figure 11 right) suggest that material contributes most to the perceived appeal of a CG character, and that strong mismatches in the level of stylization of shape and material can result in very unappealing characters. Furthermore, the middle shape was rated as equally appealing regardless of material, which could be due to the fact that it was never strongly mismatched with material. Our appeal ratings ranged from 2.5 to 4.2, which is similar to the appeal ratings reported by McDonnell et al. [154] for their static images.

**REASSURANCE** Similar to the appeal ratings, we found a main effect of material on the ratings of reassurance ( $F^*(1.51, 31.70) = 49.07$ ,  $p < 0.0001$ ,  $\epsilon = 0.755$ ), but no main effect was found for shape. An interaction between shape and material is present ( $F(4, 84) = 12.02$ ,  $p < 0.0001$ ) and post-hoc analysis showed significantly lower ratings of reassurance especially in shape-material combinations that reduce appeal as well—realistic materials on all shape levels and cartoon materials on the realistic shape ( $p < 0.02$ ). The realistic material on the cartoon and middle shape was perceived most eerie. A Cronbach's alpha value of  $\alpha = 0.88$  confirms high similarity between the appeal and the reassurance scale (see Figure 11 right and Figure 12 left).

**FAMILIARITY** Again, a main effect has been found for material ( $F(2, 42) = 12.58$ ,  $p < 0.0001$ ), but not for shape. Furthermore, there is also a significant interaction between shape and material ( $F(4, 84) = 17.99$ ,  $p < 0.0001$ ). The results of the post-hoc test for familiarity are less similar than between the appeal and eeriness ratings. Even though the combination of realistic material and realistic shape is unappealing and eerie, it was not rated significantly less familiar than other combinations. Realistic materials on cartoon and middle shapes result in the least familiar combinations ( $p < 0.02$  in all cases). See Figure 12 right.

### 2.5.2 Experiment 1b: Shading and Lighting

The above experiment reveals a strong influence of material, in particular on the appeal and reassurance ratings. The realistic material was rated as the least appealing for all character shapes, while the middle material was the most appealing for the realistic shape. Materials are controlled by a large number of shader parameters, and testing each of them is infeasible. In addition, only certain parameter combinations are meaningful and would be used in a real-world scenario. We note that all shader parameters are mainly responsible for light-material interaction, while albedo textures control primarily the color.

Instead of varying certain shader parameters within certain ranges, we modify the light transport more drastically by altering shading and lighting technique.

In an experiment similar to the previous one, we tested the initial baseline characters (three matching shape/material stylization, five expressions) with two different shaders and three illumination methods. For shading we tested a simple Phong shader in addition to the sophisticated skin shader. The lighting categories were (i) global illumination and soft shadows, (ii) ambient light and soft shadows, (iii) ambient light and hard shadows. All questions and scales were the same as for the previous experiment.

Twenty new volunteers participated in this second experiment (15 female, 5 male, ages from 19 to 30 years). A rm-ANOVA with three factors (shading, lighting, expression) was used for statistical analysis. While there was a main effect of lighting on realism ( $F(2, 38) = 6.66$ ,  $p = 0.003$ ), with global illumination being rated more realistic than soft shadows ( $p = 0.020$ ) and hard shadows ( $p = 0.004$ ), the difference was very small (means are  $3.95 \pm 0.1$ ). Besides the effect of lighting on realism, we did not find any other significant effects, neither for the other scales nor for the different shaders. These results suggest that textures have more influence than shader parameters on appearance, and therefore we explore them more in depth in the following.

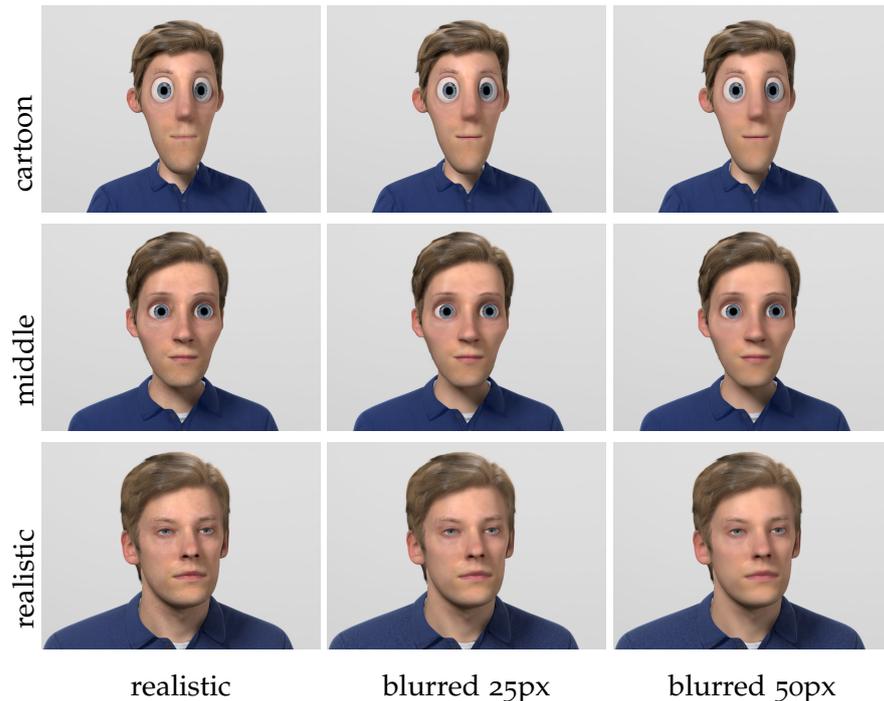


Figure 13: Stimuli for Experiment 1c: Realistic material with realistic texture and two variants with blurred textures (Gaussian kernels of 25 and 50 pixels), for the three shape stylizations.

### 2.5.3 Experiment 1c: Texture

One possible explanation of why the middle material was rated the most appealing for the realistic shape could be the reduced pigmentation variation as reported by Fink and Matts [77], as discussed in Section 2.2. In order to analyze whether their findings on attractiveness can also explain our effects on appeal and reassurance, we designed a variation of the first experiment of Section 2.5.1. Our third experiment should then:

- test whether it is possible to influence appeal or realism by changing only the albedo texture,
- show a possible correlation between attractiveness and appeal/reassurance, and
- reveal whether appeal can be increased without sacrificing realism too much, simply by filtering a photo-realistic texture.

To this end, we created two additional textures with reduced skin details by applying uniform Gaussian blur of kernel sizes 25 and 50 pixels (for 4k textures), respectively. The 50px kernel covers barely 1 cm of the face, which translates into around four pixels in image-space. Lips and skin were filtered independently in order not to blur the boundary inbetween; eyebrows were not filtered. These three textures (realistic, blurred 25px, blurred 50px) were used in combination with the realistic material. To enable a comparison with Experiment 1, we also included the cartoon and middle materials (with their original textures only). This results in a set of 5 materials, which were also transferred to the middle and cartoon shapes, as shown in Figure 13.

For this experiment we tested these 5 materials on the 3 shape stylizations, but used the neutral expression only, leading to 15 stimuli in total. Note that the three realistic materials differ in their (blurred) texture only. The presentation of the stimuli was repeated three times with different random orderings. After each stimulus, participants were asked to rate it according to the previously described scales for appeal, reassurance, and realism, plus a new scale *attractiveness*. Twenty-one new volunteers (13 female, 8 male), average age 24.6 years, participated in the experiment. For statistical analysis, a rm-ANOVA with three factors (shape, material, and expression) was used. All results from Section 2.5.1 were confirmed, and thus we only describe the main effects related to the added material levels.

**REALISM** Although a main effect was found for shape ( $F^*(1.29, 25.78) = 124.98, p < 0.0001, \epsilon = 0.645$ ), material ( $F(4, 80) = 17.52, p < 0.0001$ ) and an interaction between shape and material ( $F(12, 240) = 6.42, p < 0.0001$ ), the post-hoc shows that this is not related to the added textures. The ratings for the two blurred textures are between the realistic and the middle texture, but are not significantly different for any

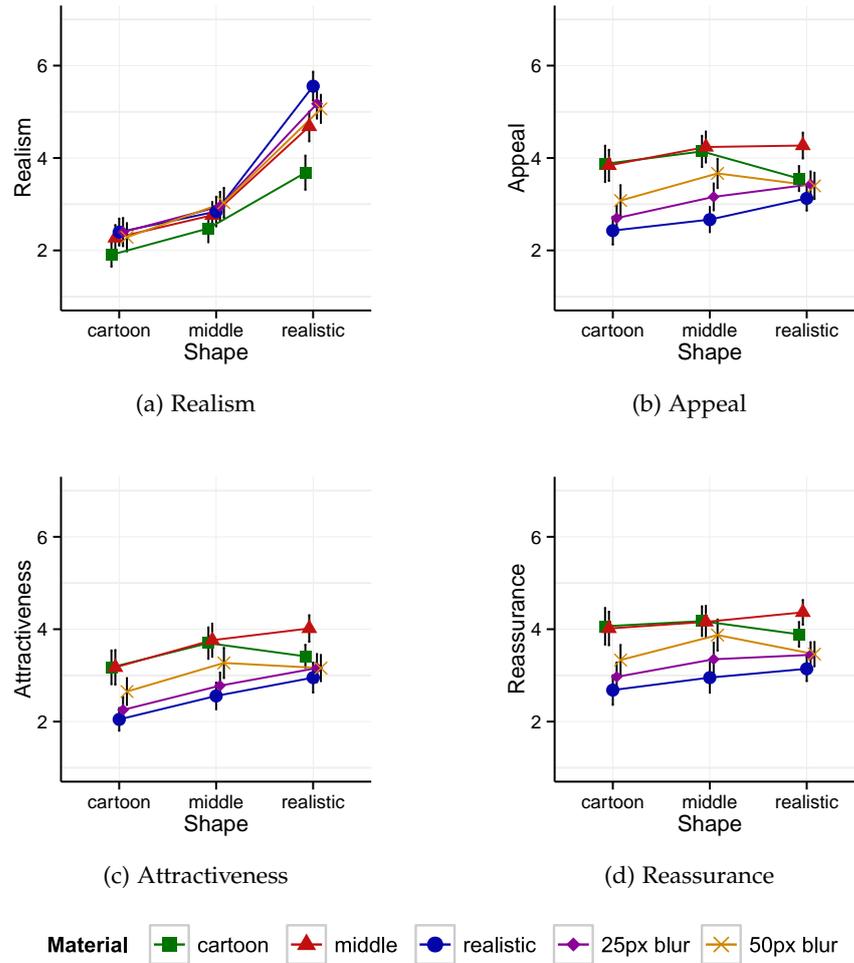


Figure 14: Results of Experiment 1c: While there is nearly no difference between the realistic and blurred textures for the realism scale, the blurred textures increase appeal and attractiveness and reduce eeriness.

shape. This confirms our initial assumption that blurring a realistic texture only slightly reduces the perceived realism of a character.

**APPEAL AND ATTRACTIVENESS** Due to the high similarity between appeal and attractiveness (Cronbach's  $\alpha = 0.87$ ) we report these results together. A main effect was found for shape for attractiveness ( $F^*(1.33, 25.54) = 5.36, p = 0.021, \epsilon = 0.665$ ) but not for appeal. Material was significant in both cases (Appeal:  $F^*(1.68, 33.60) = 27.17, p < 0.0001, \epsilon = 0.421$ ; Attractiveness:  $F^*(1.56, 31.26) = 16.72, p < 0.0001, \epsilon = 0.391$ ). The interaction between shape and material is significant (Appeal:  $F^*(7.05, 94.03) = 4.99, p < 0.0001, \epsilon = 0.588$ ; Attractiveness:  $F(12, 240) = 2.88, p = 0.005$ ). As we hypothesized, the blurred textures were rated higher than the realistic texture. This effect is stronger for the cartoon and middle shapes and a significant

difference between the realistic and 50px blurred version was found ( $p < 0.003$  in all cases). For other comparisons between the blurred and realistic textures no significant difference was found. However, the graphs in Figure 14 show that the two blurred textures were rated equally appealing for the realistic shape. In contrast, a stronger blur is preferred for cartoon and middle shapes. We therefore conclude that blurring realistic skin textures is a reasonable approach for increasing appeal or attractiveness, without losing too much realism. Although the results of our tests are not significant in some cases, these findings are in line with research of Fink and Matts [77]: We generalize their findings to character shapes of different stylization levels.

**REASSURANCE** Although the graphs of reassurance and appeal are similar (Figure 14;  $\alpha = 0.89$ ), a main effect was found for material only ( $F^*(1.44, 28.72) = 24.55$ ,  $p < 0.0001$ ,  $\epsilon = 0.359$ ), but not for shape. In addition, there is an interaction between shape and material ( $F^*(7.128, 142.46) = 2.66$ ,  $p = 0.029$ ,  $\epsilon = 0.594$ ). The two blurred textures have been rated less eerie than the realistic version. Significant differences have been found between the realistic texture and the 50px blurred version for cartoon and middle shapes ( $p < 0.0001$ ). Thus, blurring a texture does not only increase appeal, but also reduce eeriness.

#### 2.5.4 Conclusion

The three tests described above allow us to draw the following main conclusions on the tested dimensions:

- Shape is the main descriptor for realism, while material is more important for perceived appeal, reassurance, and attractiveness. Strong mismatches in stylization between material and shape affect negatively the appeal and attractiveness of the characters and make them more eerie.
- Texture has stronger influence on appeal and attractiveness than shading or illumination models. Blurring a realistic texture does not significantly reduce realism but increase appeal and attractiveness.
- Ratings for appeal, reassurance, and attractiveness measure similar concepts ( $\alpha > 0.87$  in all experiments), but do not correlate with the realism scale ( $\alpha < 0.5$  in all experiments).

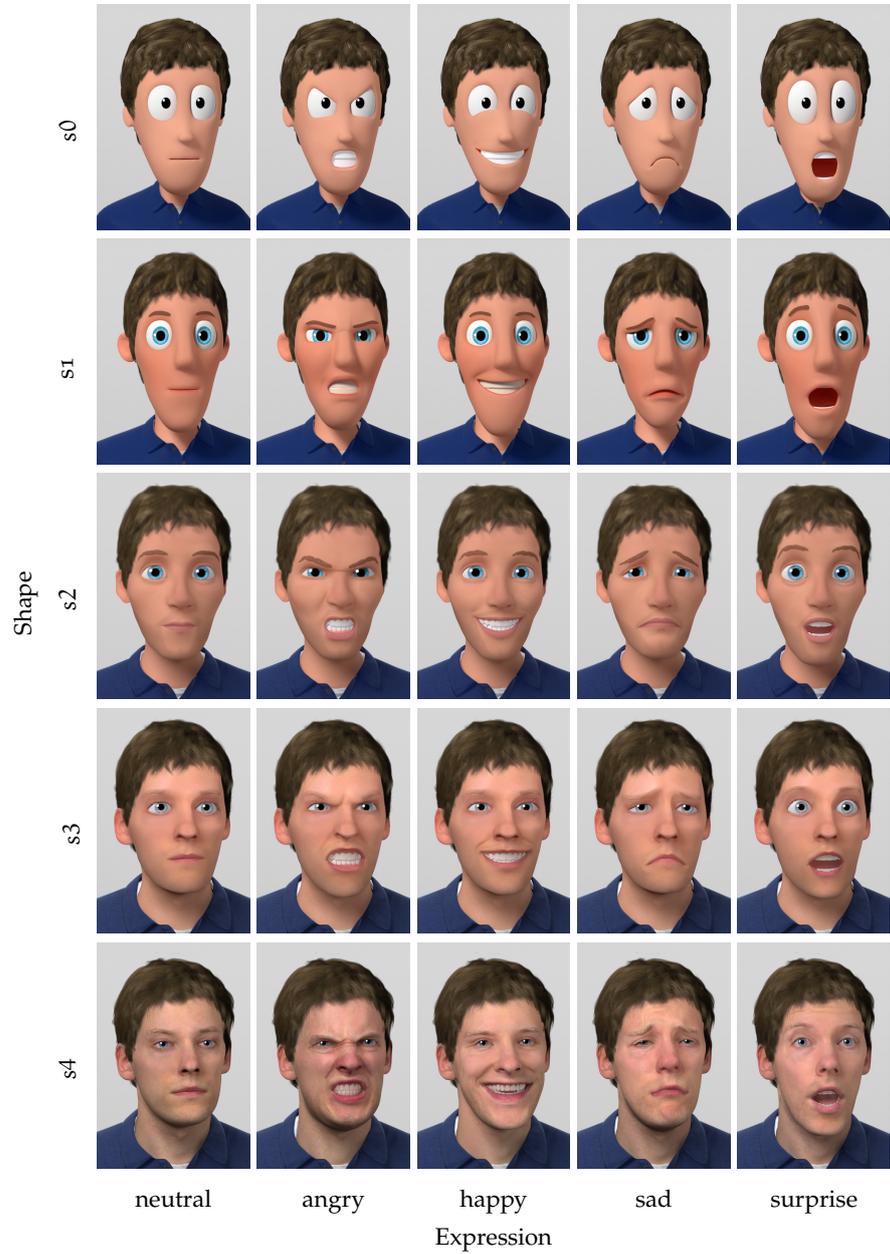


Figure 15: Stimuli for Experiment 2: renderings of the male character for different stylizations (rows) and basic emotions (columns).

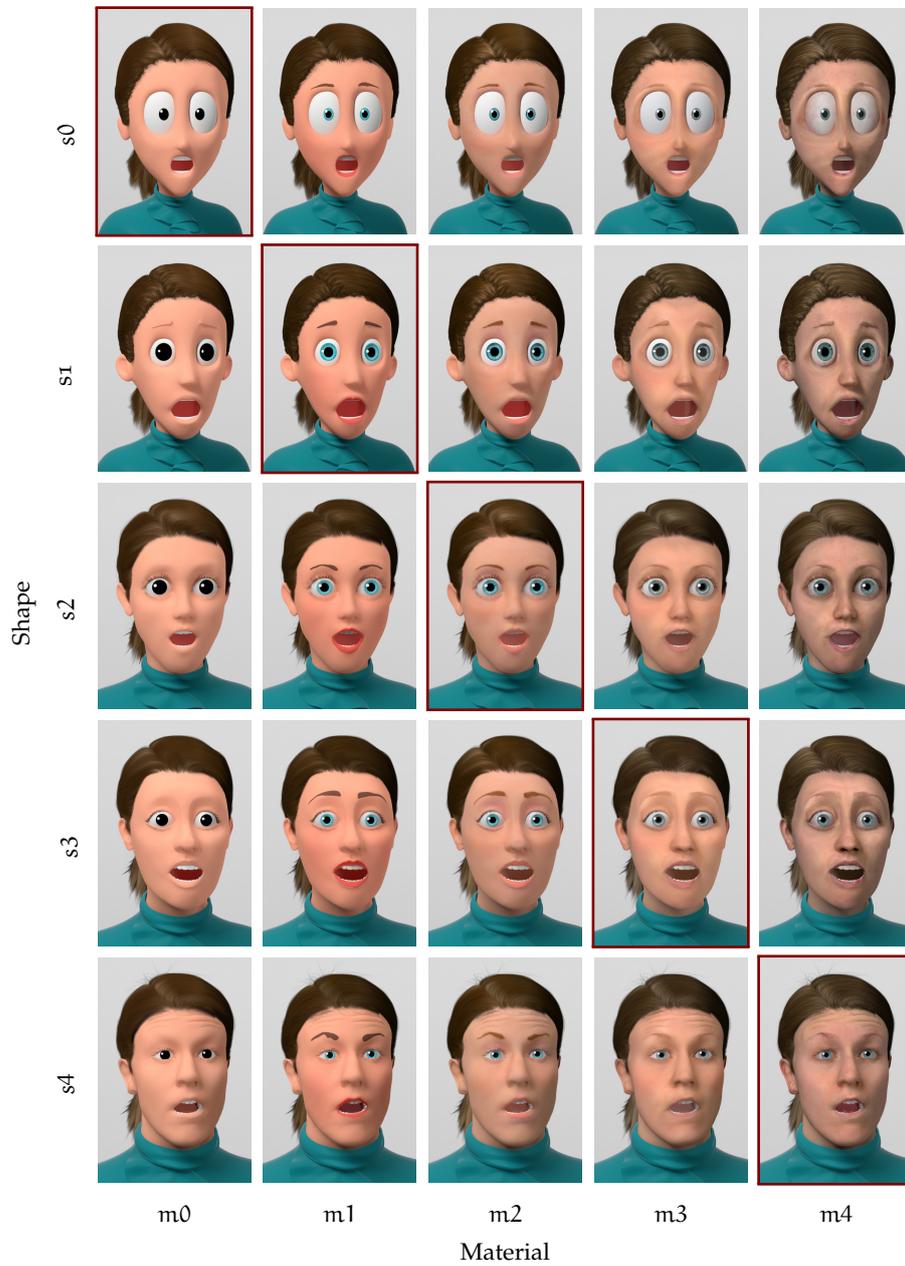


Figure 16: Stimuli for Experiment 2: combinations of shape and material stylization for the female character (surprise expression), with baseline stimuli on the diagonal.

## 2.6 EXPERIMENT 2: FURTHER INVESTIGATION OF SHAPE AND MATERIAL

The experiments in Section 2.5 indicate that different stylization levels of material and shape have a big impact on perceived appeal or realism. However, our set of stimuli contained only a single character, and the realism scale was not densely sampled. A more stylized character might reveal that big mismatches between material and shape cause unappealing results, or a stylization level between middle and realistic might cause uncanny reactions. To allow for a more generalized conclusion about different stylization levels, further investigation is required.

In the following experiment we analyze the effect of varying stylizations on shape and material, including matching and mismatching levels of stylization, on a significantly extended set of stimuli. In particular, we seek answers to the following questions:

- Can our findings be observed on other characters as well?
- Does a strong mismatch between material and shape create unappealing results only for realistic shapes or for all shapes?

**STIMULI** We extended our initial stimuli with another character of different gender, because this adds by design a clearly distinctive person. For each character, two additional stylizations were created, yielding five stylization levels from *level 0* (most stylized) to *level 4* (highly realistic). We distinguish between stylizations in material and shape by using the prefix *m* and *s* respectively. The new stylizations (level 0 and level 3) have been particularly designed by the artists to fill the gaps for perceived realism in the stylization scale. For these levels our character designs are inspired by *Pocoyo* and *Tangled*. We also changed the hairstyle of the virtual male character in order to allow a better comparison with a photograph of the actor. This provides us with baseline ratings on appeal and realism for the real person. The new set of stimuli is composed of 2 characters times 5 shape stylizations times 5 material levels times 5 expressions, leading to a total of 250 images. A representative subset of the stimuli is shown in Figure 15 and Figure 16, for the five expressions and matching shape/material levels of the male character (Figure 15), and the 25 combinations of material and shape for the female character (Figure 16).

**PROCEDURE** The largely extended stimuli require a reduction of the scales in order to keep the experiment tractable. Given that the appeal, reassurance, and attractiveness scales measure similar concepts, and that the familiarity scale did not provide much information, we decided to keep only the realism and appeal scales for this experiment. Furthermore, we increased the display time of the stimuli to

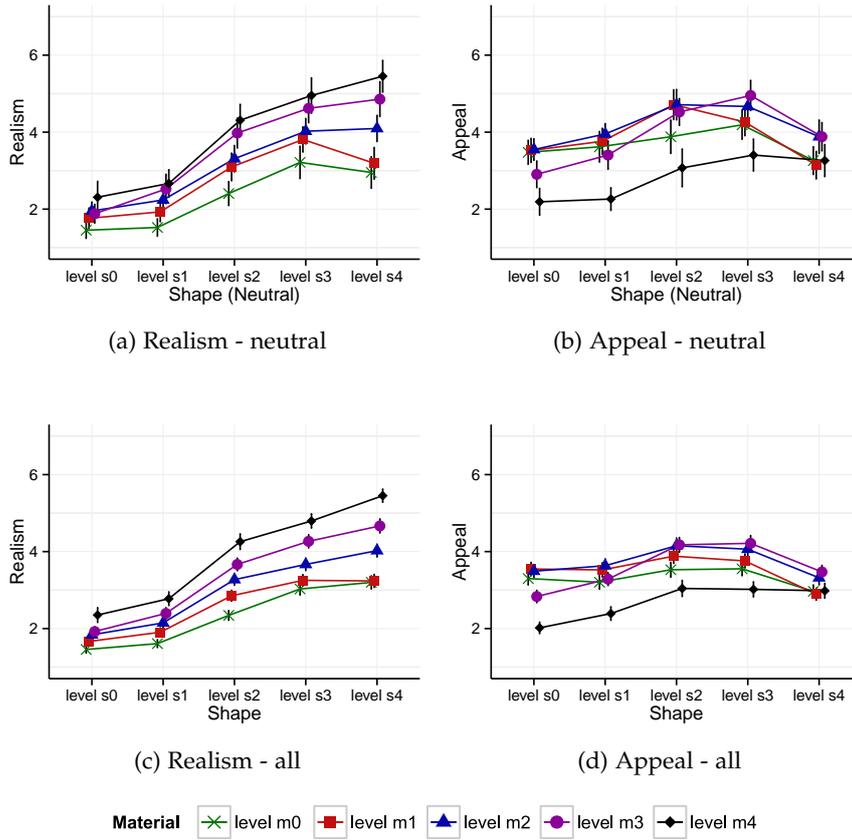


Figure 17: Results of Experiment 2: Ratings for perceived realism and appeal for different shape and material stylizations. Upper row: neutral expression averaged over male and female characters. Bottom row: averaged over all expressions and characters.

4 s, and showed the neutral male and female baseline characters before the experiment, such that participants could better estimate the range of characters from the beginning on. At the end of the actual experiment, participants rated a photograph of the real characters in neutral expression. The rest of the experiment remains similar to the previous one. With all these changes, participants finished the experiment within 50 minutes or less. Twenty-one new different volunteers (17 female, 4 male) took part, average age 23.4 years.

Our results are summarized in Figure 53 and are mostly consistent across male and female. Repeated measures ANOVA with four factors (character, shape, material, and expression) was used for statistical analysis. Differences between the two characters were significant, but since they were rather small and/or inconsistent, we exclude them from further analysis. In the following we present an in-depth discussion of the realism and appeal ratings, and report the impact of expression in Section 2.9.

**REALISM** A main effect of shape ( $F^*(1.98, 39.6) = 178.67, p < 0.0001, \epsilon = 0.495$ ) and material ( $F^*(1.33, 26.4) = 73.92, p < 0.0001, \epsilon = 0.333$ ) was found as well as an interaction between shape and material ( $F^*(6.71, 134.1) = 11.59, p < 0.0001, \epsilon = 0.419$ ). Post-hoc analysis shows that all shapes ( $p < 0.004$ ) and most of the materials ( $p < 0.003$  except for level m0 and m1) differ significantly from each other. The 25 groups resulting from the combinations of shape and material differ also significantly in more than 80% of the cases. Most non-significant comparisons can be found for the shape level s0 (see Figure 53). For example, increasing the material from level m1 to m2 or from level m2 to m3 does not cause a significant difference. This contrasts with the case of the realistic shape levels s3 and s4 ( $p < 0.002$ ). This is in line with the results from Section 2.5.1, and confirms that as the shape becomes more realistic, the material stylization becomes more dominant for perceived realism.

**APPEAL** The main effects of shape ( $F^*(2.58, 51.6) = 20.97, p < 0.0001, \epsilon = 0.645$ ) and material ( $F^*(1.88, 37.6) = 20.39, p < 0.0001, \epsilon = 0.470$ ) are comparable. There is a slightly weaker interaction between shape and material ( $F^*(6.06, 121.3) = 14.29, p < 0.0001, \epsilon = 0.379$ ). Post-hoc analysis reveals that shape levels s2 and s3 were perceived more appealing than the other shape levels ( $p < 0.0002$  in all cases between the two groups).

For the materials, only the most realistic version (level m4) was significantly less appealing than all other materials ( $p < 0.0002$ ). This supports our assumptions from Section 2.5.3 that smooth(ed) skin pigmentations are perceived more appealing. For the abstract shape s0, material levels m0, m1, and m2 form a cluster without any significant difference; this cluster is found significantly more appealing than material levels m3 and m4 ( $p < 0.03$ ). On the other hand, shape level s3 is rated significantly higher with matching material levels (m2 and m3), with both more stylized (m0 and m1) and more realistic (m4) materials being rated significantly lower. These results support that in all cases a strong mismatch between shape and material is perceived as unappealing.

**PHOTOGRAPH** At the end of the experiment, participants rated a photograph of the real actor in neutral pose. As expected, the average realism rating is very high (6.98, SD = 0.15). The average appeal rating was 4.5 (SD = 1.40), which is higher than the average ratings for the realistic s4/m4 characters (3.26, SD = 1.33). This dip in appeal rating for the s4/m4 character is in agreement with the uncanny valley theory [158]. However, appeal for stylizations s2/m2 and s3/m3 (4.71, SD = 1.25 and 4.95, SD = 1.25) were rated highest. In addition, Figure 20 depicts that realism *alone* is a bad predictor for appeal; instead, our results show that the compatibility of shape and material styl-

izations, i.e., their matching degrees of realism, has a stronger (and predictable) influence on appeal.

## 2.7 EXPERIMENT 3: EFFECT OF EXPRESSIONS

In previous experiments, we have analyzed the overall effect that shape and material have on the perception of faces. Here, we first analyze whether different levels of stylization in shape and material, including mismatches between them, affect the recognition and intensity of expressions (*anger, happy, neutral, sad* and *surprise*). We then discuss how ratings are affected by particular expressions (Figure 15 and Figure 16). This is interesting since previous findings suggest that valence of emotion affects character perception [36, 216], making negative expressions to be rated less appealing than positive expressions. In particular, we seek answers to the following questions:

- Does the level of stylization affect the intensity of expressions? Are they easier or more difficult to recognize?
- Do negative expressions affect the perceived appeal of characters? Is this influenced by stylization of shape or material?

### 2.7.1 Intensity and Recognition of Expressions

As discussed previously, stylization is a well-known tool for artists to enhance the expressivity of 3D characters, removing unnecessary details and enhancing specific features. In this experiment we explore how the different stylizations of shape and material affect *recognition* and the perceived *intensity* of the expressions, and which of the two dimensions is dominant for expression recognition. The extended 250-stimuli set from Experiment 2 is used again.

Each stimulus was presented for 4 seconds in random order; participants were first asked to classify the expression according to the following options: *anger, happy, neutral, sad, surprised*. After each answer (except for *neutral*), a follow-up question asked to rate the expression intensity with respect to a seven-point Likert scale bounded by *extremely low* and *extremely high intensity*. When participants rated an expression as *neutral*, its intensity was set to the lowest value. Twenty-four new volunteers (16 female, 8 male, 23.6 years old on average) took part in this experiment. Results are shown in Figure 18 and again a rm-ANOVA with four scales (character, shape, material, and expression) was used for statistical analysis.

**RECOGNITION** We found a main effect of expression ( $F^*(1.22, 28.04) = 74.00, p < 0.0001, \epsilon = 0.305$ ), as well as several interaction effects between expression and shape ( $F^*(4.56, 104.9) = 41.3, p < 0.0001, \epsilon = 0.285$ ), texture ( $F(16, 368) = 4.97, p < 0.0001$ ) and model ( $F^*(2.3, 51.26) =$

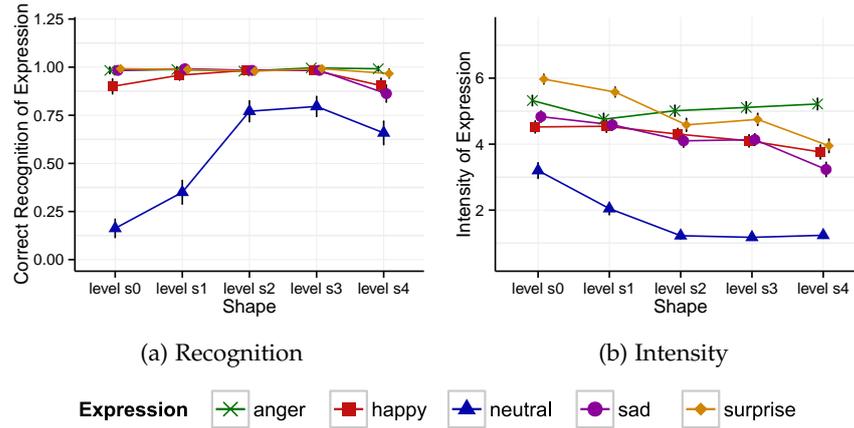


Figure 18: Results of Experiment 3: Effect of shape on the recognition and intensity of the expression. All expressions, except neutral, have been recognized well or outstandingly well independent of the shape. However the intensity reduced continuously with higher shape stylization levels.

4.23,  $p = 0.016$ ,  $\epsilon = 0.557$ ). The neutral expression is mainly responsible for all these effects; its recognition rate was lower ( $p < 0.002$ ) than the other expressions, varying strongly across different shape levels. This neutral expression was in general poorly recognized for the more stylized shapes (s0 and s1): For instance, some participants reported that the big round eyes made them look surprised. This might be explained by the fact that cartoons are usually designed to enhance expressivity, not to be posed displaying a neutral emotion.

We also found a main effect for shape ( $F(4, 92) = 44.23$ ,  $p < 0.0001$ ), which is mainly determined by the neutral expression, as discussed above, and a main effect for material ( $F(4, 92) = 10.09$ ,  $p < 0.0001$ ). The material level m4 reduced the recognition rate significantly ( $p < 0.015$ ) but only by 2%.

**INTENSITY** Main effects of shape ( $F^*(2.11, 48.61) = 91.40$ ,  $p < 0.0001$ ,  $\epsilon = 0.528$ ) and material ( $F^*(2.47, 56.90) = 30.46$ ,  $p < 0.0001$ ,  $\epsilon = 0.618$ ) were found. Apart from the angry expression, the perceived intensity of expressions is continuously reduced with increasing realism of shape ( $p < 0.0002$ ). Only in the case of shape levels s2 and s3 does the intensity remain constant. In the case of material, the absolute difference was very small (0.5 between the lowest and highest mean) and only the material level m4 had a higher intensity ( $p < 0.0002$ ). This matches previous research [227, 228], which found that details such as wrinkles increase the expressivity of realistic characters, although in our case the effect is weaker.

In addition, a main effect of expression ( $F^*(2.57, 59.10) = 204.6$ ,  $p < 0.0001$ ,  $\epsilon = 0.642$ ) and interactions between shape and expres-

sion ( $F^*(5.78, 132.94) = 19.00, p < 0.0001, \epsilon = 0.361$ ), material and expression ( $F(16, 368) = 5.04, p < 0.0001$ ), and expression and model ( $F(4, 92) = 19.55, p < 0.0001$ ) were found. In particular, the happy, sad, and surprise expressions are perceived with lower intensity as the realism of shape increases. This difference is significant in the majority of cases for shape levels s3 and s4 ( $p < 0.01$ ), but is less frequent for lower shape levels. The perception of the angry expression, on the other hand, remains constant along shape abstractions.

Overall we found that expressions of cartoon shapes are perceived as more intense, which confirms that properly stylizing features helps increase expressivity. The neutral expression is hard to read for very stylized character shapes, which suggests that low-intensity subtle expressions are harder to convey in abstract characters designed to enhance expressivity. Additionally, we found no or small impact of material on the intensity or expression recognition, which indicates that shape is the dominant dimension when designing expressive characters.

### 2.7.2 Effect of Expression on Realism and Appeal

In our previous experiments on material and shape with the five basic expressions (Section 2.5), we found that appeal and eeriness measure similar concepts, while effects for familiarity were generally small. We focus here on the effect of expressions on realism and appeal with the extended stimuli set. The rest of the analysis can be found in the supplemental material. Figure 21 shows the results, which we analyze below.

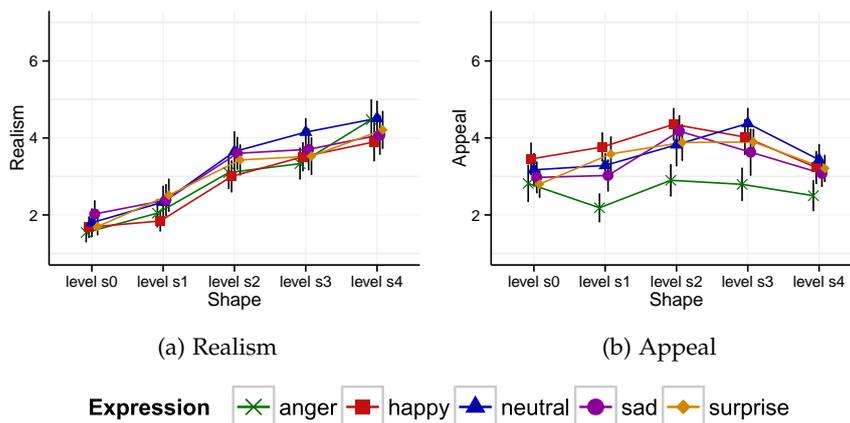


Figure 19: Results of Experiment 2: While emotions do not differ in realism, the anger expression was perceived as more eerie and unappealing for all stylization levels.

**REALISM** A main effect of expression ( $F(4, 80) = 10.38, p < 0.0001$ ) was found, which could be mainly attributed to the neutral and sad expressions, which have been perceived as more realistic ( $p < 0.006$ ). Because the means are located within a small range ( $\pm 0.16$ ), we classify this effect as noise and omit similar examples for the rest of this section. Nevertheless, equal realism ratings confirm that expressions were well designed by the artists.

**APPEAL** A main effect was found for appeal ( $F^*(1.56, 31.36) = 19.34, p < 0.0001, \epsilon = 0.392$ ), which is primarily caused by the anger expression ( $p < 0.001$ ). These results reveal that anger is rated much lower with respect to appeal. Previous studies reported that negative emotions trigger unpleasant responses from the observers [36]; our results confirm these studies. Moreover, this effect is maintained even in the presence of highly stylized and appealing characters, suggesting that negative expressions are perceived as unappealing independent of stylization level.

Additionally, we observe that ratings are unsteady across different stylization levels for the rest of the expressions. In many cases, interaction effects between expression and shape or material are found with  $p < 0.0001$ . Zhu et al. [253] showed with photographs that different instances of the same expression do indeed vary in perceived appeal. We believe that this might also be the primary reason for the variations in our ratings. We rule out recognition as an error source, since all the expressions were recognized outstandingly well (see previous subsection).

## 2.8 DISCUSSION

Shape and material are two of the main aspects that define the appearance of virtual characters, which in turn are crucial when defining the visual look of animated feature films. We have analyzed the perceptual effects of different stylizations along these dimensions on computer-generated faces. In particular, we have studied five different stylizations of two virtual characters—male and female—ranging from very realistic to highly stylized, varying both the shape and the material.

Our results show that the main contributor for the perceived *realism* is shape, and the effect of material stylizations grows when shape realism is increased. This implies that mismatches in material and shape are less prominent on abstract characters. The resulting asymmetry is shown in Figure 53 (a,c), where the curves spread out as the level of realism increases.

On the other hand, we have found that material is the main factor for perceived *appeal*, specifically the albedo texture. In general, *appeal*, *attractiveness*, and *eeriness* are highly dependent on the material styl-

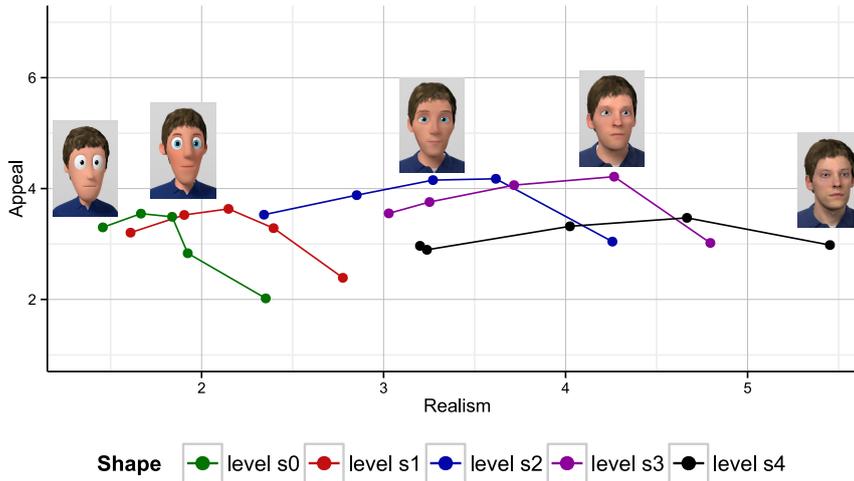


Figure 20: Participant ratings for our male stimuli plotted on a realism-vs-appeal diagram, similar to [217, 145]. Each graph corresponds to one shape stylization, while graph nodes correspond to material levels. The icons are placed above the nodes of matching shape/material levels. The diagram reveals that perceived realism is a bad predictor for appeal. Instead, it is the compatible degree of realism of both shape and material that matters.

ization. Matching levels of stylization of geometry and material cause the highest ratings of appeal, while strong mismatches (e.g., very realistic material on a stylized shape) result in unappealing and eerie characters.

Interestingly, as shown in Experiment 1c and later backed-up in Experiment 2, subtle stylization of a realistic material (edge-preserving blur on the albedo texture) increases appeal without sacrificing realism. These stylizations de-emphasize unwanted skin impurities, pores, and wrinkles, and our results are in accordance to empirical knowledge regarding the effect of makeup. Moreover, our results relate with previous findings on face perception showing that smooth homogeneous skin is generally rated more attractive, since it is a good estimate of a young and healthy subject [76]. However, this trend is only observed for mild stylizations, and stronger ones quickly reduce realism.

Our results are consistent across all tested *expressions*, except for anger, which was consistently rated less appealing and more eerie. This can be explained by negative or aggressive expressions triggering a defense response and a negative reaction of the viewer [36]. Our results are also consistent between different *characters*. Although small differences between the characters exist, all reported trends are consistent and well visible.

Realism alone was shown to be a bad predictor for appeal (Figure 20), which is not well aligned with the theory of the uncanny valley, although a similar finding was reported for rendering style [154].

One possible explanation is that some of our characters were difficult to categorize by the participants, due to their mismatched appearance parameters [195, 154].

Finally, our experiments show how stylization affects the *intensity of expressions*, and that *shape* is the main factor in this case, whereas material has no significant influence for stylized shapes. This confirms previous knowledge on modeling or drawing expressive stylized characters, where expressivity is mainly determined by the global shape of the character. However, for realistic shapes, we have observed that material stylization slightly, but significantly, reduces the perceived intensity of expressions. Another possible explanation, which also merits further investigation, is that realistic characters make *suspension of disbelief*<sup>1</sup> harder to maintain, and therefore observers find it more difficult to emotionally connect with the virtual character. These results are consistent with previous work [227, 228] and may explain the conscious disturbing effect of stylizing hyper-realistic characters in some movies (e.g., *A Scanner Darkly* or *Renaissance*).

### 2.8.1 Limitations and Future Work

As in all user studies, our results are only strictly valid for our particular set of stimuli. We have focused on a specific set of stylizations for two realistic characters, varying shape and material following typical designs used in feature animation. This of course limits the universality of the conclusions, which may not generalize if the character styles differ greatly from ours. However, since our design space was densely sampled and the observed trends are consistent between the different characters, we believe that our observations can be used as valid guidelines for creating digital characters within a reasonable range of styles.

In our statistical analysis, we employed a common significance threshold of  $p < 0.05$ . With the amount of results we report, it might be that some significant results are false positives. Because we only focus on clear, reoccurring trends, and since many significances have  $p < 0.001$ , it is unlikely that one of our main conclusions is a false positive.

Investigating the effect of realistic and stylized animation was outside the scope of this work, due to the number of stimuli combinations that would have to be tested, and the technical complexity of creating a scale of stylized animations. Previous work has shown no difference in ratings for realism, and only small differences in appeal ratings for static or motion-captured characters [154]. However, we would expect a more complex interaction between motion and appeal when combining characters and animations of different levels of

<sup>1</sup> In fiction, the suspension of disbelief is a semi-conscious decision by the viewer to accept as real what clearly is not. This allows her to connect with the story.

stylization. Therefore, a full investigation of motion stylization is an important future direction.

Finally, note that we analyzed clear *peak* expressions, avoiding the less attractive transitions between expressions common in the real world [253]. Evaluating the impact of these transitions for different stylizations could be also an interesting avenue of future work.

In summary, we believe that our work provides many interesting insights for the creation of virtual characters, and offers a set of guidelines that we hope will help practitioners and inspire future research on the perception of virtual characters.

## 2.9 APPENDIX A. EFFECT OF EXPRESSIONS IN EXPERIMENTS 1A AND 1B

In this section, we discuss the results of Experiment 1a (shape and material) and Experiment 1b (shading and lighting), analyzing the particular *expressions* of the stimuli. For the analysis we conducted a three-way repeated measure ANOVA. A Tukey HSD test was used for pairwise comparisons within each experiment. Figure 21 shows the results, which we proceed now to analyze.

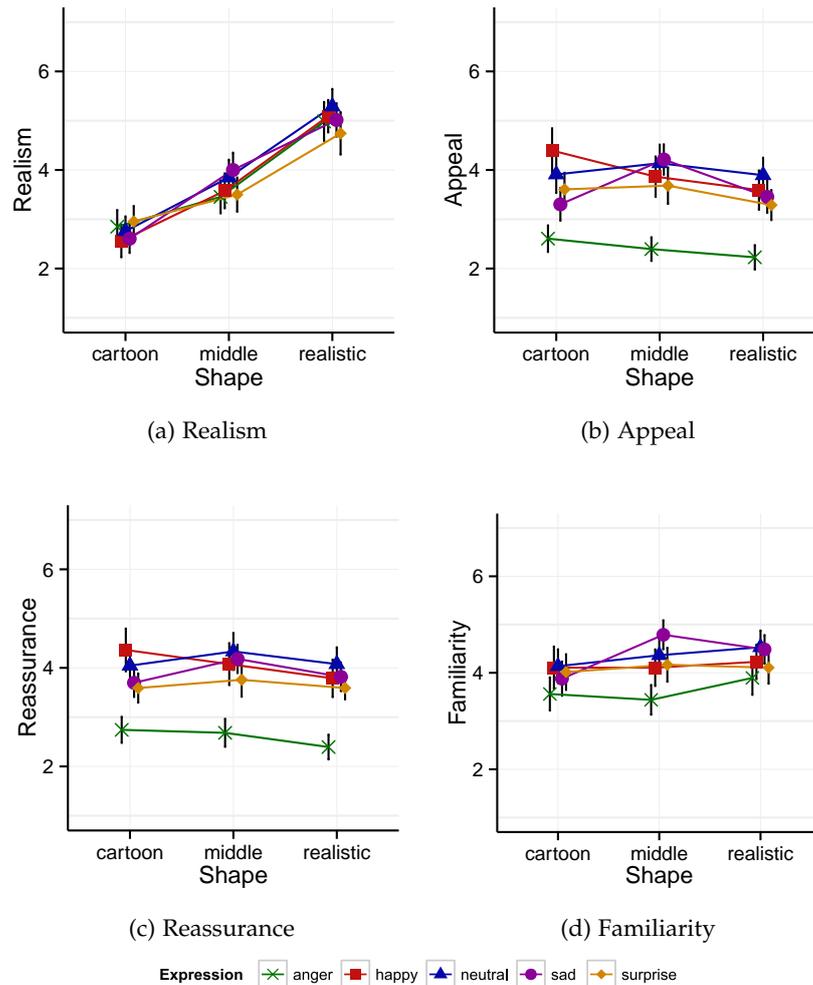


Figure 21: Results of Expressions in Experiment 1a: While emotions do not differ in realism, the anger expression was perceived as more eerie and unappealing for all stylization levels.

**REALISM** A main effect of expression for realism was found in the shading and lighting experiment (Exp. 1b:  $F(4, 76) = 3.78$ ,  $p = 0.007$ ), but not in the shape and material experiment. The effect could mainly be attributed to the sad expression, which has been rated slightly

more realistic ( $p = 0.005$ ) than others. Because the means ( $3.90 \pm 0.15$ ) of all groups are within a small range, we classify this effect as noise and omit similar examples for the rest of this section.

**APPEAL AND REASSURANCE** A main effect was found for appeal in both experiments (Exp. 1a:  $F^*(2.57, 54.01) = 33.14$ ,  $p < 0.0001$ ,  $\epsilon = 0.643$ ; Exp. 1b:  $F^*(1.53, 29.10) = 22.22$ ,  $p < 0.0001$ ,  $\epsilon = 0.383$ ), which is primarily caused by the anger expression ( $p < 0.0001$ ). Similarly, there is a main effect of expression for reassurance in both experiments (Exp. 1a:  $F^*(2.61, 54.77) = 24.61$ ,  $p < 0.0001$ ,  $\epsilon = 0.652$ ; Exp. 1b:  $F^*(1.68, 31.92) = 18.61$ ,  $p < 0.0001$ ,  $\epsilon = 0.420$ ), again mainly caused by the anger expression ( $p < 0.0001$ ).

**FAMILIARITY** A similar main effect is obtained for familiarity (Exp. 1a:  $F(4, 84) = 8.80$ ,  $p < 0.0001$ ; Exp. 1b:  $F^*(2.04, 38.84) = 5.15$ ,  $p = 0.001$ ,  $\epsilon = 0.511$ ). In this case anger is the only reason for the significant differences of the means. But the anger expression is only significantly different from the happy ( $p = 0.003$ ) and neutral expressions ( $p = 0.001$ ).

2.10 APPENDIX B. ADDITIONAL DIAGRAMS

2.10.1 Experiment 1a: Shape and Material

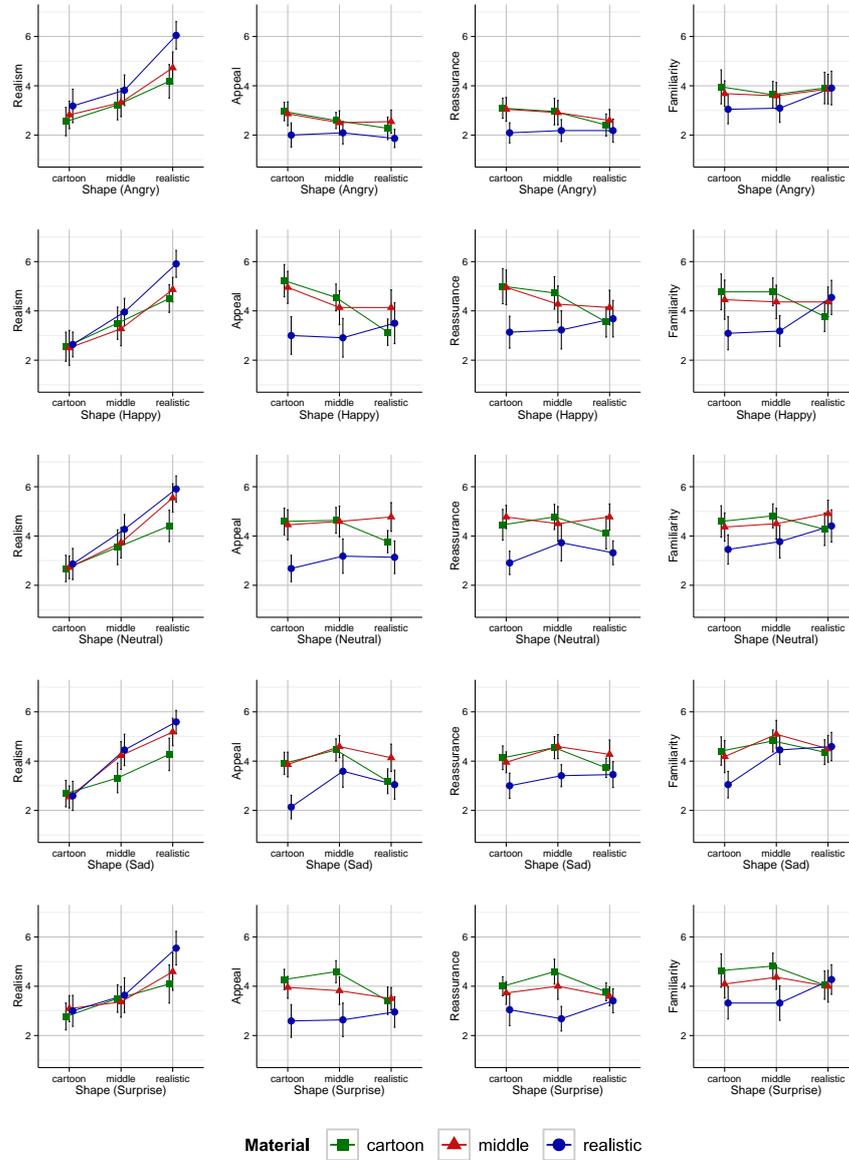


Figure 22: Results for Experiment 1a: Ratings for perceived realism, appeal, reassurance and familiarity, separated by expression.

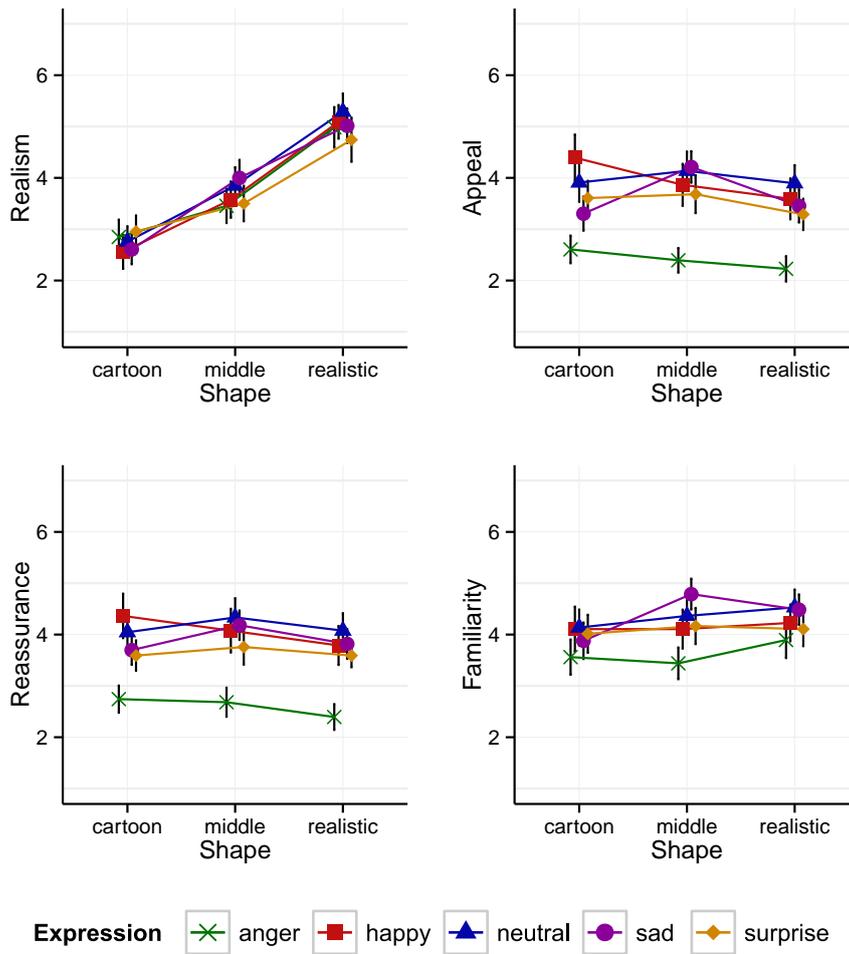


Figure 23: Results for Experiment 1a: Ratings for perceived realism, appeal, reassurance and familiarity averaged over all materials and grouped by shape and expression.

## 2.10.2 Experiment 1b: Shading and Lighting

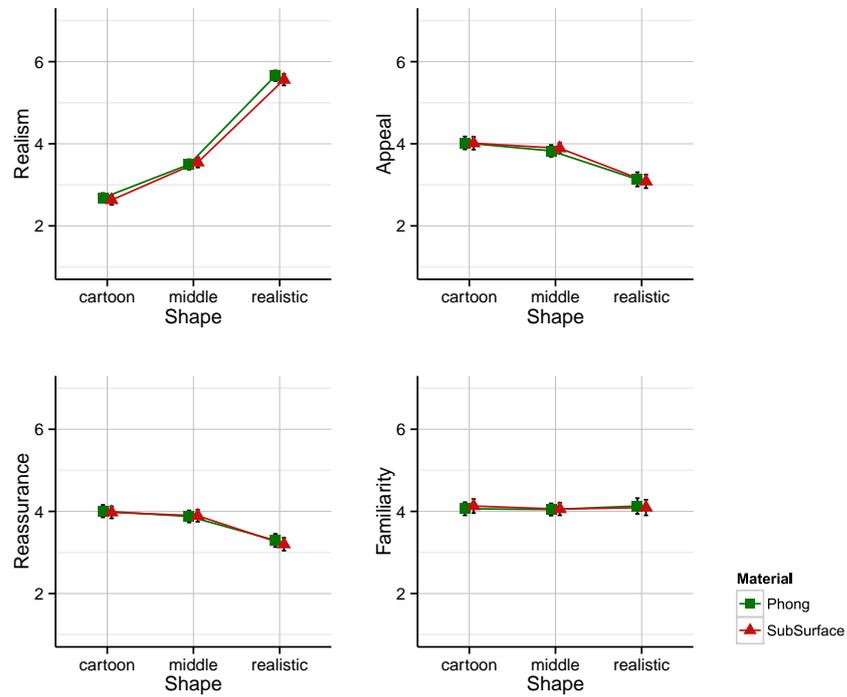


Figure 24: Results for Experiment 1b: Ratings for perceived realism, appeal, reassurance and familiarity grouped by stylization level and shading.

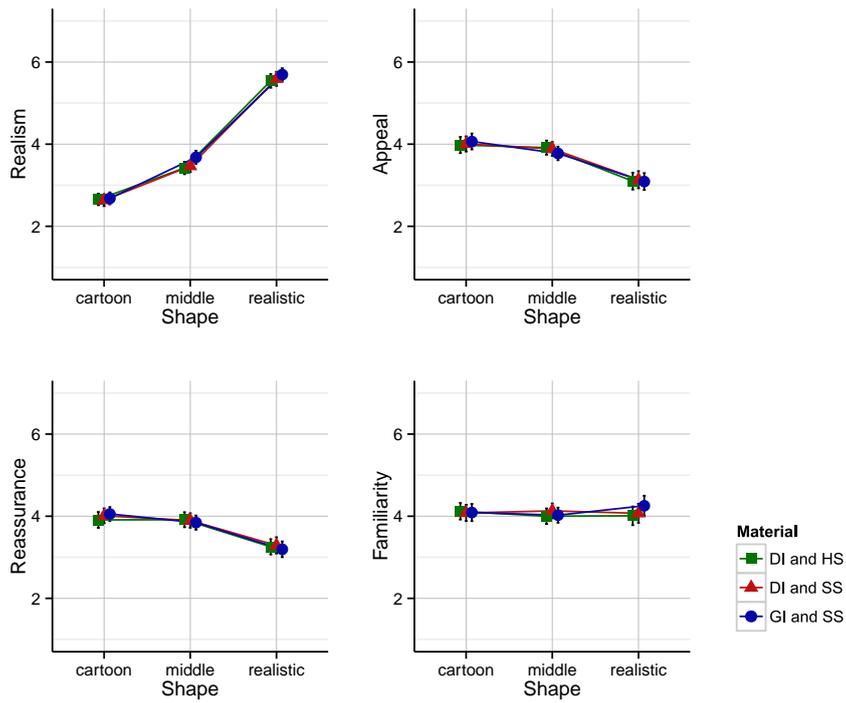


Figure 25: Results for Experiment 1b: Ratings for perceived realism, appeal, reassurance and familiarity grouped by stylization level and lighting.

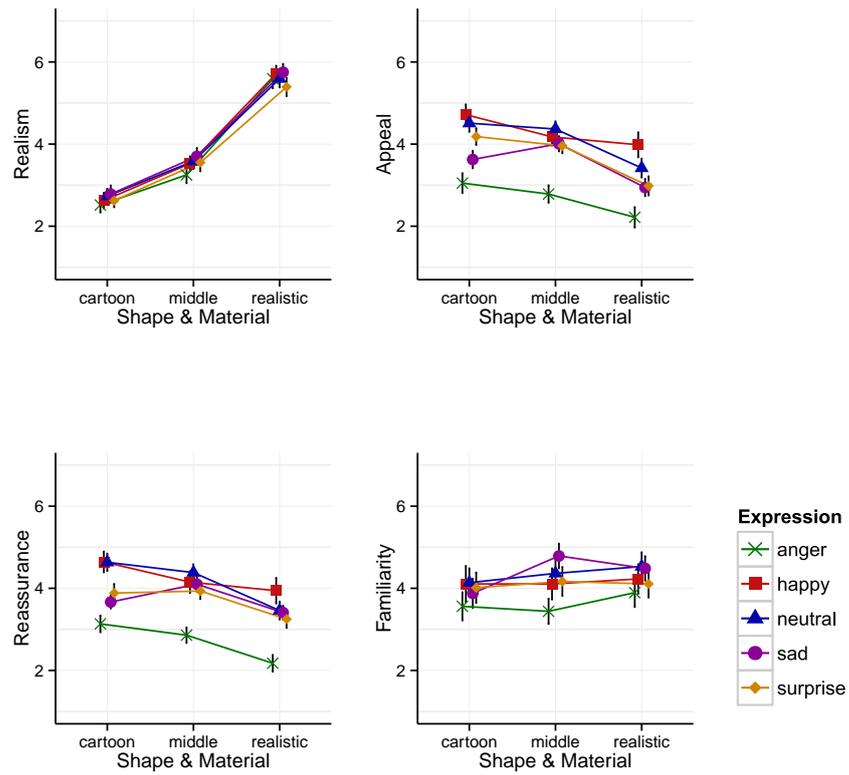


Figure 26: Results for Experiment 1b: Ratings for perceived realism, appeal, reassurance and familiarity grouped by stylization level and expression.

2.10.3 Experiment 2: Further Investigation of Shape and Material

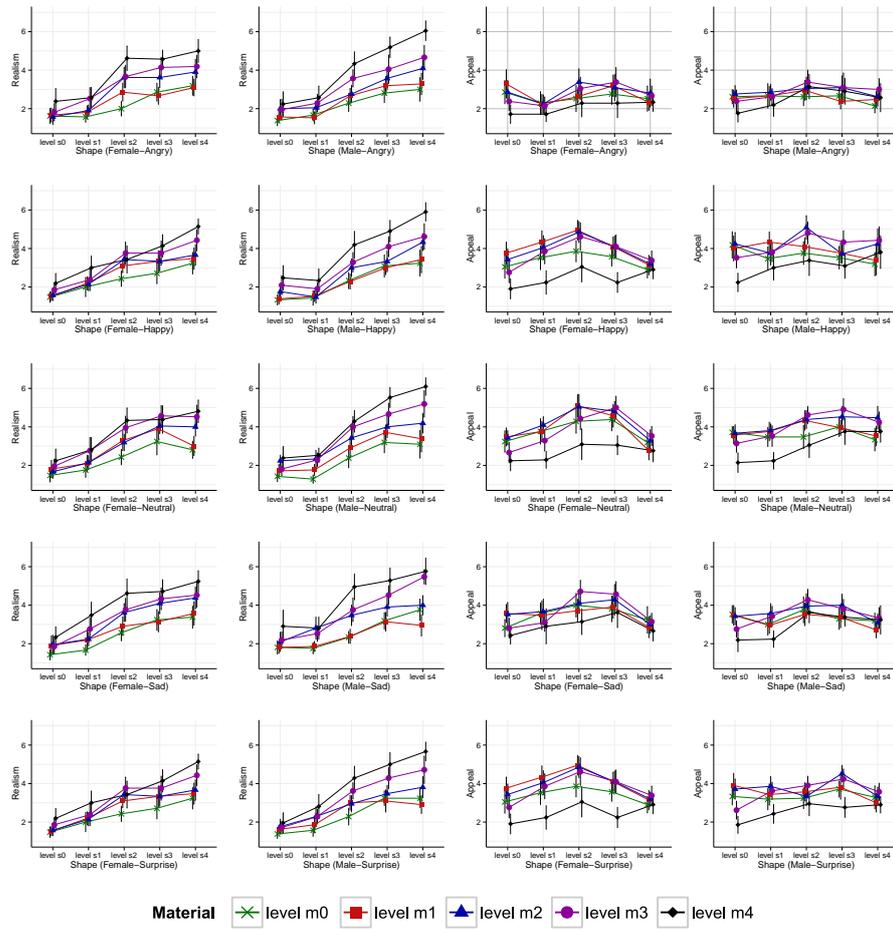


Figure 27: Results for Experiment 2: Ratings for perceived realism and appeal, separated by expression, male and female.

## 2.10.4 Experiment 3: Recognition and Intensity

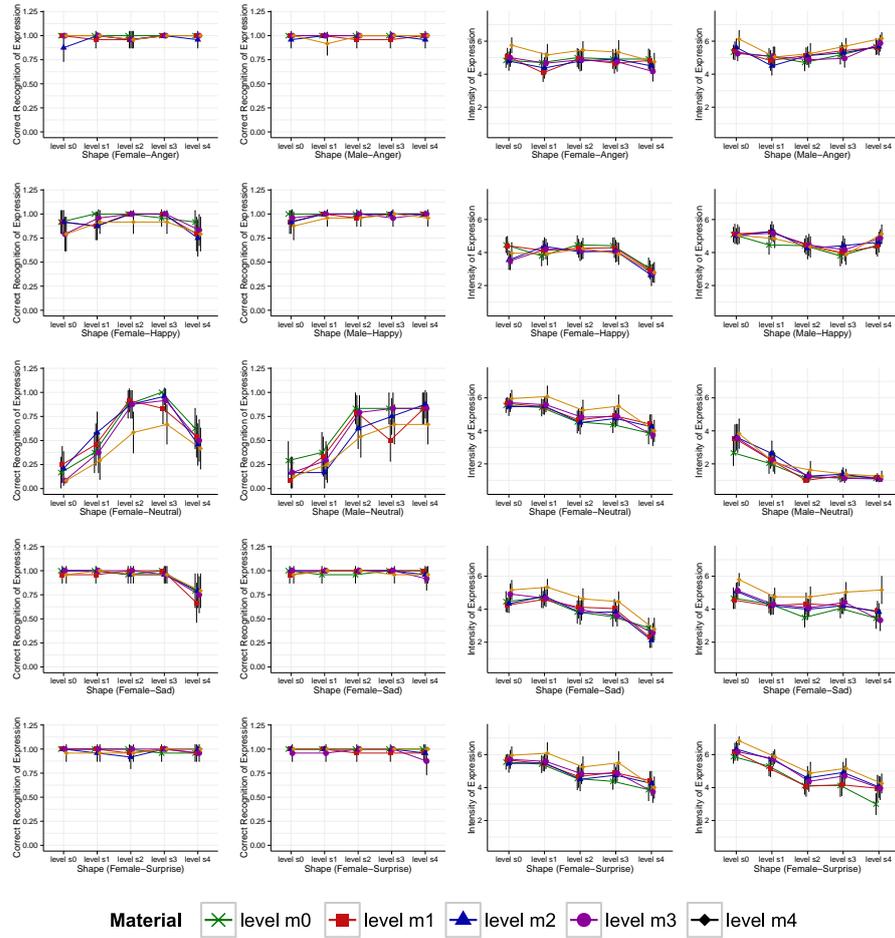


Figure 28: Results for Experiment 3: Ratings for correct recognition and intensity, separated by expression male and female.

## A BIOPHYSICALLY-BASED MODEL FOR SKIN AGING

This chapter presents a time-varying, multi-layered biophysically-based model of the optical properties of human skin, suitable for simulating appearance changes due to aging. We have identified the key aspects that cause such changes, both in terms of the structure of skin and its chromophore concentrations, and rely on the extensive medical and optical tissue literature for accurate data. Our model can be expressed in terms of biophysical parameters, optical parameters commonly used in graphics and rendering (such as spectral absorption and scattering coefficients), or more intuitively with higher-level parameters such as age, gender, skin care or skin type. It can be used with any rendering algorithm that uses diffusion profiles, and it allows to automatically simulate different types of skin at different stages of aging, avoiding the need for artistic input or costly capture processes. Moreover, we provide a simplified version of our model as well as tabulated sum-of-Gaussians profiles for gaming applications.



Figure 29: Examples for 30- and 80-year-olds, with mediterranean skin (type III) for the male, and caucasian (type I) for the female. The small inset shows the very different scattering profiles for each one: A and B correspond to the old and young males respectively, while C and D correspond to the old and young females. All the rendered images in this work are best viewed in the digital version..

### 3.1 INTRODUCTION

Human skin changes significantly with age. It becomes thinner and more dry, while the concentration of chromophores (the main skin

pigments) diminishes and becomes more irregular [13, 69, 168]. All these changes affect its optical properties, which in turn yield visible changes in its color, translucency, apparent roughness, and consequently its final appearance. Current techniques to simulate aged skin either rely on artistic skills, or require capture sessions with subjects of the appropriate age. However, simulating the *changes* in appearance for one single subject remains an open problem.

We present a comprehensive biophysically-based model of human skin suitable for simulating skin aging. We rely on existing in-vivo and ex-vivo measurements from the extensive medical and tissue optics literature. Our multi-layered model identifies and takes into account the key structural changes (such as the slimming of the dermis and epidermis, the flattening of the dermoepidermal junction or the deterioration of the dermal collagen fibers), as well as changes in the concentration of its main chromophores (such as melanin, hemoglobin or water). This naturally leads to changes in the scattering and absorption properties of skin, which translate into considerable appearance variations.

In particular, we reproduce the changes in the base color, translucency and skin surface reflectance. Current models (e.g. [135, 18, 59, 60]) are either not suitable to simulate aging, or assume many simplifications that ignore some key aspects of aging. In our work, we focus on optical properties, and do not simulate coarser structural changes like the production of wrinkles. Our model can be formulated in terms commonly used in graphics, such as absorption or scattering spectral coefficients, or even intuitive descriptors like age, gender, skin care or skin type, and can be easily integrated into existing rendering systems capable of handling diffusion profiles. Additionally, we provide sum-of-Gaussians tabulated results for gaming applications, as well as a simplified version of our model that still captures the main optical effects of skin aging, offering an attractive trade-off between accuracy and simplicity.

### 3.2 PREVIOUS WORK

There is a vast number of studies about human skin in the medical and tissue optics literature [222, 168, 70, 110]. We refer the reader to these excellent sources, and focus here on works more closely related to computer graphics.

**SIMULATING SKIN.** Some of the first rendered skin images used measured isotropic BRDFs [147]; multiple subsurface scattering (SSS) was later added by Stam [210]. One of biggest breakthroughs came with the introduction in computer graphics of the dipole diffusion approximation [121], a model that was quickly made more efficient in subsequent work [120]. Donner and Jensen [59] presented a BSS-

RDF model capable of dealing with layered materials, while D'Eon et al. [52] approximated the diffusion profiles of translucent layers by a sum of Gaussians. The approach was later transferred to screen-space by Jimenez et al. [122, 125], including transmittance [124]. Recently, the quantized diffusion model [50] and the subsequent photon beam diffusion [98] provide more accurate and still efficient solutions to the general problem of subsurface scattering, including skin. All these works focus on the simulation of light transport, while we focus on the optical properties of a biophysically accurate model.

Other works focus on acquiring the reflectance properties of human skin [147, 14]. In the case of human faces, Debevec and his team have developed increasingly sophisticated methods from the original light stage [49], where multiple images of a given subject are acquired under varying illumination conditions. These are then used to reconstruct the reflectance functions of the acquired model, allowing them to produce impressive results (e.g. [143, 86, 87]). Weyrich et al. [233] developed a model based on measurements of 149 subjects, whose parameters can then be edited. Recently, Graham et al [94] synthesize high-resolution facial surface microstructures based on captured patches of microgeometry.

Tsumura et al. [220, 221] developed an image-based technique for separating spatial distributions of melanin and hemoglobin in human skin by independent-component analysis of a skin color image. Donner and Jensen [60] introduced a two-layer skin model whose parameters controlled the amount of oil, melanin and hemoglobin in the skin. Later, a multi-layered model was presented, with spatially-varying absorption and scattering parameters [61]. Jimenez et al. [123] modeled changes in appearance due to varying concentrations of melanin and hemoglobin, caused by different emotional or physical states. Possibly the most detailed skin model in computer graphics nowadays is the work by Krishnaswamy and Baranoski [135, 18]. However, many key aspects that change with age and affect its optical properties are not taken into account.

**AGING IN HUMANS.** Not many works in graphics deal with aging in humans. Two notable exceptions are the works of Golovinskiy et al. [92], who presented a statistical model that allows to synthesize wrinkles and other fine *geometric* features due to aging, and Boissieux et al. [28], who proposed image and model based methods to simulate wrinkles. Their works are therefore complementary to ours.

### 3.3 ANATOMY AND OPTICAL PROPERTIES OF SKIN

We provide a brief overview of the anatomy and the optical properties influencing light transport in human skin, as well as the main effects of aging; in Section 3.4 we will present our biophysically-based

model for aging of human skin. Skin is composed of three main layers: epidermis, dermis and hypodermis, with additional sub-layers (see Figure 30). The optical properties of these layers are described by the reflection off its surface, as well as the absorption and scattering events produced in the randomly inhomogeneous distribution of blood, chromophores and other components. Practical measurements of these events are expressed in the form of an absorption coefficient  $\mu_a$  and a reduced scattering coefficient  $\mu'_s$ . Absorption and scattering coefficients represent how much light is absorbed or scattered by the medium per incremental path length traveled. In the case of reduced scattering coefficients, we assume a diffuse light measurement, having then an isotropic phase function ( $g = 0$ ). Spectral reflectance curves are commonly used to describe skin appearance.

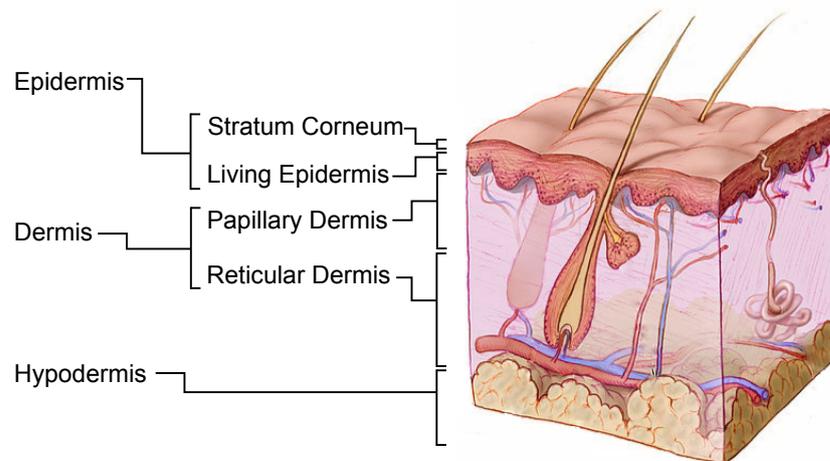


Figure 30: Skin anatomy showing the main layers present in skin (credits: Don Bliss, National Institutes of Health, US).

The **epidermis** is the outermost layer of the skin, and is divided in two sub-layers: The *stratum corneum*, composed of many flattened and packed dead cells, is relatively high on lipids and low on water content. Light reflection off this layer is determined by surface topography and sebum production [244, 180]. The second layer is the *living epidermis*, which primarily contains living cells like the *melanosomes*, responsible for the storage, transport and synthesis of melanin. This is the main light absorber of the epidermis (together with, to a much lesser extent, dietary carotenoids), and therefore its main colouring pigment. Melanosomes and melanin exhibit forward and isotropic scattering respectively [19].

The **dermis** is generally described as a dense irregular connective tissue. It has an inhomogeneous distribution of blood vessels in its two-main sub-layers: The *papillary dermis*, which is a thin layer of connective tissue that interdigitates with ridges of the epidermis and contains many capillary loops and elastin fibers; and the *reticular dermis*,

containing thick bundles of collagen, more elastin fibers, arteries and veins. The main absorbers in the dermis are blood hemoglobin, in both oxygen saturated (*oxy*) and desaturated forms (*deoxy*), carotene, bilirubin and water. The scattering properties of the dermal layers are defined by collagen fibril distributions, packed in bundles. Quantitative studies of dermal collagen provide morphometrical measurements showing histograms of fiber distributions per dermal area [32]. Given its larger thickness, scattering in skin is dominated by the reticular dermis. The interface between the epidermis and the dermis is called the *dermoepidermal junction*.

Last, the **hypodermis** is adipose tissue which may be up to a few centimeters thick or completely absent, depending on its location. Absorption in the hypodermis is mainly determined by hemoglobin, lipids and water. The main scatterers in this layer are spherical droplets of lipids.

### 3.3.1 Skin aging

The aging process in human skin causes variations in both its structure and its main constituents. In particular:

All this leads to visible changes in the optical properties of skin, thus changing its appearance (see Figure 31). Previous skin models [60, 135] simplify the structure of the skin, allowing to simulate only a subset of the changes enumerated above. Accurately simulating *all* these changes requires a finer structural and biophysical granularity. We thus first develop a complete model including accurate descriptions of the hypodermis, water concentration, the dermoepidermal junction, the production of sebum in the epidermis or the distribution of collagen fibers.

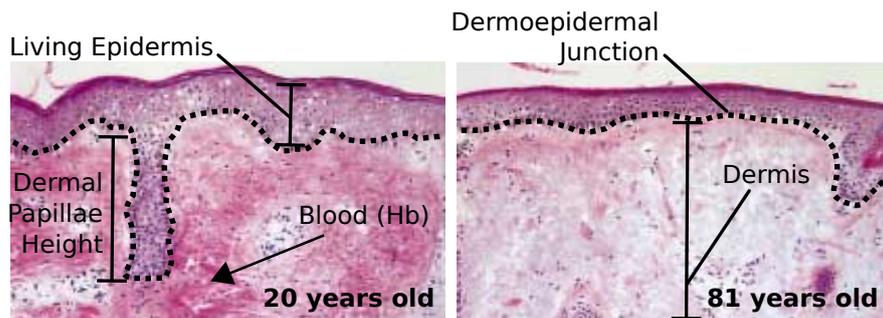


Figure 31: Skin close-ups (images augmented  $\times 10$ ) for healthy subjects, ages 20 (left) and 81 (right). Epidermal thinning, flattening of the epidermal junction and the decrease of hemoglobin are clearly visible (credits: image adapted from [40]).

### 3.4 OUR SKIN AGING MODEL

Our model can take as input specific values from accurate measured data (i.e. tissue optics studies), or roughly approximated based on the subject's age, gender, skin type, body location and even life habits (for instance modeling the effect of excessive exposure to UV rays). While obvious differences exist in skin according to gender or skin types, the aging process is rather similar in most cases [71]. Therefore, unless explicitly mentioned, our aging model can be applied to both genders and all skin types.

#### 3.4.1 Layered structure

Aging selectively affects the structural organization of skin. Capturing variations in its optical properties and therefore its color and translucency, requires a model that includes a fine detailed description of its layered structure. Our model is made up of the five main representative layers introduced in Section 3.3: stratum corneum (SC), living epidermis (LE), papillary dermis (PD), reticular dermis (RD) and hypodermis (HD). While the living epidermis, the papillary and reticular dermis are usually included in existing skin models, the stratum corneum and the hypodermis are commonly ignored or highly simplified. However, the former is the outmost layer, and therefore has a large impact in the final appearance, which also depends on its location and exposure to external agents (see Figure 42). The latter has been at best considered as a purely reflectance layer [135] or merged with other internal tissues [59]; we instead include it as a fully participating medium, which leads to significant changes in the appearance of the skin, as shown in Figure 32. Our model considers each layer as an optically thick and semi-infinite (lateral extent) medium, described by its thickness  $d$ , index of refraction  $\eta$ , as well as absorption and reduced scattering coefficients  $\mu_a$  and  $\mu'_s$  (see Subsections 3.4.3 and 3.4.4 respectively).

As skin ages, its rate of renewal slows down and the thickness of its layers is non-uniformly reduced, ranging from a 6.4% reduction in the living epidermis to up to 50% in the hypodermis. This large reduction in the hypodermis emphasizes the importance of this layer in a skin aging model. We model thickness atrophy by linearly reducing each layer using the values summarized in Table 1.

Additionally, the *dermoepidermal junction* becomes progressively flatter with age (see Figure 31). We model its interdigitated nature, as proposed by Meglinsky [155], using a periodic surface vertically centered at the mean depth of the layer as:

$$z_{(x,y)} = \bar{z} + A_x \sin(\omega_x x + \phi_x) \cdot A_y \sin(\omega_y y + \phi_y) \quad (1)$$

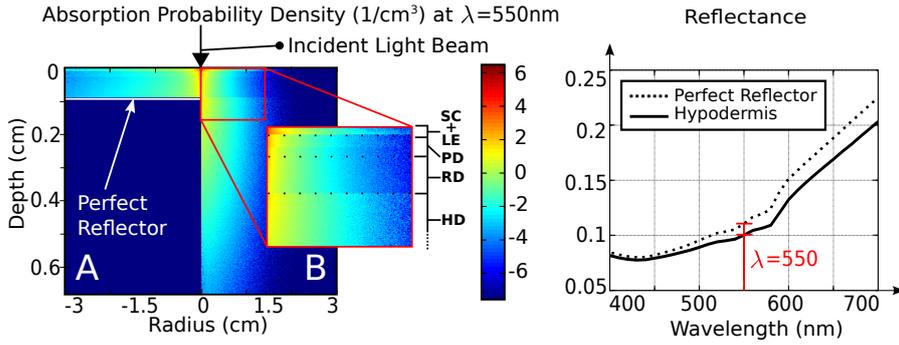


Figure 32: The influence of the hypodermis. **Left:** Absorption probability density for  $\lambda=550\text{nm}$  a skin sample. In A the hypodermis is considered to be a perfect reflector [135]; in B we model it as a participating layer ( $d = 5.9\text{mm}$ ,  $\eta = 1.44$ ), with absorption (Subsection 3.4.3) and scattering events (Subsection 3.4.4). The perfect reflector model in A increases absorption in the dermis and epidermis as a consequence of light being reflected off the hypodermis. **Right:** Spectral reflectance for both hypodermis models: accurately modeling the hypodermis reduces the amount of total outgoing energy, specially at longer wavelengths

Skin layer	Structure			
	$\eta$	$d$ (mm)	Decrease (p.d.)	Reference
SC	1.53	0.02	(n.a.)	[141]
LE	1.34	0.08	6.4%	[226]
PD	1.395	0.18	6% - 16%	[226]
RD	1.39	1.82	6% - 16%	[226]
HD	1.44	5.9	30% - 50%	[30]

Table 1: Typical thickness of the skin layers considered in our model (base values for  $t = 30$ ) and their range of variation per decade (p.d.). These values give a mean plausible range corresponding to average skin [155].

where  $x, y$  are coordinates lying on the plane and  $z$  is depth,  $\phi_{x,y}$  and  $\omega_{x,y}$  are the phase offsets and frequency modeling the surface roughness, and  $A_{x,y}$  is the peak amplitude of the rete ridge surface, which decreases linearly with age. The reported measurements by Giangreco et al. [88] range from  $100\mu\text{m}$  to almost  $0\mu\text{m}$ , for ages 20 through 80.

### 3.4.2 Surface reflection

On average, about 5% of the incident light is reflected off the surface of the stratum corneum, although the actual percentage varies with hydration and sebum production. Its effect on the appearance of skin depends mainly on its roughness and oiliness [180]. Tissue

optics studies commonly relate roughness variations related with aging with parameterizations of a Beckmann distribution function [244]. We model the skin's outmost surface using the Torrance-Sparrow microfacet BRDF model [218], using the average angle  $\alpha_{sr}$  of the microfacets as the Beckmann's distribution roughness. We model  $\alpha_{sr} = \arctan(R_z/l_e)$ , with  $R_z$  the average peak-to-valley height from a set of measurements, and  $l_e = 960\mu\text{m}$  the skin profile length of each measurement [64]. Values of  $R_z$  increase with age, and the values used in our model are shown in Figure 33 (left).



Figure 33: **Left:** Evolution of surface roughness  $R_z$ . Measured mean values vary from  $83\mu\text{m}$  in young individuals up to  $151\mu\text{m}$  in the elderly [64]; **Right:** Evolution of sebum production, measured in mg of lipids produced in  $10\text{cm}^2$  after 3 h. Surface roughness increases, while sebum production decreases in the elderly, being always greater in men than women [180].

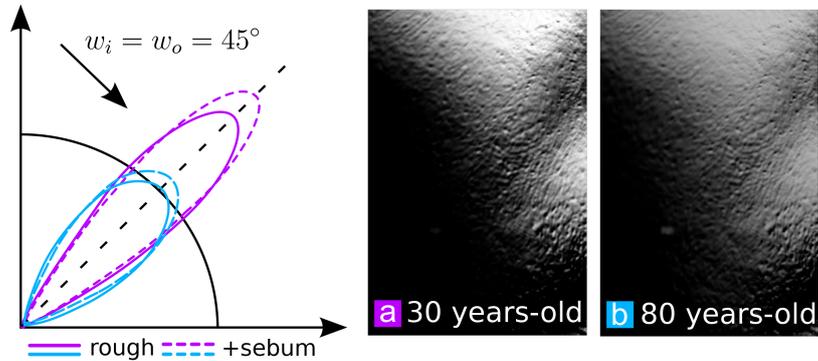


Figure 34: Effect of roughness variation with age on the specular reflection. **Left:** The roughness of the Beckmann distribution increases with age as a combination of both the sebum production and the skin surface roughness variations. **Center:** Reflectance in young skin (30 years old). **Right:** Reflectance in aged skin (80 years). In both cases, both the change of roughness and sebum is modeled (Figure 33 shows the values of a and b used)

Moreover, the production of sebum creates a thin film that smoothens the skin roughness. As humans age, sebaceous glands become less ac-

tive: the thickness  $s$  of the sebum film can vary from  $3\mu\text{m}$  in young skin to  $1.69\mu\text{m}$  in the elderly [180]. To simulate this, we assume that the average slope of the microfacets is reduced due to the sebum layer, so that  $\alpha'_{sr} = \psi_s \alpha_{sr}$ , where  $\psi_s \in [0..1]$  is a modulation factor dependent on the average sebum volume filling the microfacets (see Appendix 3.8 for details). We make an assumption about the pyramidal shape of the microfacets, but we only use that to derive the *statistical* roughness reduction, not the real per-facet slope variation. Figure 34 compares the specularity of skin with age, both with and without including sebum in the skin.

### 3.4.3 Absorption

We define the absorption coefficient  $\mu_a$  of a skin layer as the sum of contributions for all absorbing chromophores present in that layer, which is the common way to define it in multi-layered tissue optics [115]:

$$\mu_a = \ln(10) \sum_i c_i \epsilon_i \quad (2)$$

where  $\epsilon_i$  is the chromophore extinction coefficient<sup>1</sup> and  $c_i$  its concentration. Our model includes a wide range of chromophores identified in tissue optics literature as dominant light absorbers in human skin [115]. Similarly to previous biophysical models for computer graphics [135], we include melanin (*eumelanin* and *pheomelanin*), *haemoglobin* (Hb) found in blood (oxygen-saturated  $\text{HbO}_2$  and desaturated Hb), as well as *bilirubin* and *carotene*. In addition, and different from previous models, we include *water* as the seventh chromophore, using spectral absorption measurements given by Pope and Fry [181]. Although water is not a highly-absorbent component itself, it has a great influence in appearance, since it defines dryness and dilutes the concentration of other chromophores (see Figure 36). Water concentrations decrease with age, causing a progressive thinning of the layers (Subsection 3.4.1).

The concentration of chromophores is specified for each layer in anatomical volume fractions of tissue occupied by their principal container (melanosomes in the case of melanin  $\vartheta_m$  or whole blood in the case of hemoglobin  $\vartheta_{\text{Hb}}$ ). Water is expressed directly as the volume fraction of water content  $f_{\text{H}_2\text{O}}$  while carotene  $c_{\text{car}}$  and bilirubin  $c_{\text{bil}}$  are concentrations expressed in g/L. Table 2 shows per-layer chromophore concentrations for an adult skin sample. Last, we also estimate a baseline skin absorption, modeling depigmentation due to small-scale tissues. Our baseline is approximated from data measured by Saidi [192].

<sup>1</sup> The term  $\ln(10)$  stands in Equation 2 because first spectrometers recorded  $\epsilon$  in base 10 [115].

Skin layer	Chromophore concentrations				
	$\vartheta_m$	$\vartheta_{Hb}$	$c_{bil}$	$c_{car}$	$f_{H_2O}$
SC	-	-	-	$2.1e-4$	0.05
LE	[0-35]%	-	-	$2.1e-4$	0.2
PD	-	[0.1-20]%	[0.05-3]	$7.0e-5$	0.5
RD	-	[0.1-20]%	[0.05-3]	$7.0e-5$	0.7
HD	-	5%	-	-	0.7

Table 2: Chromophore concentrations for a typical adult. Values inside brackets indicate plausible range values [155].

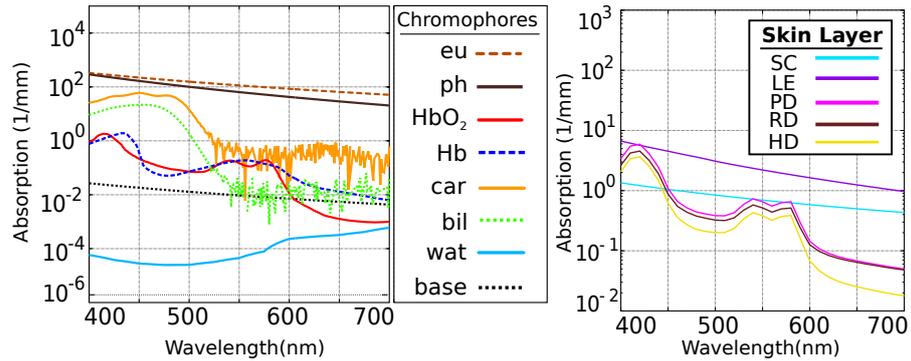


Figure 35: **Left:** Spectral absorption curves for the seven chromophores, plus the baseline absorption used in our model. We refer to the supplementary material for equations used to model absorption coefficients for each chromophore. **Right:** A representative example of the spectral absorption coefficients for the five layers considered in our model, for a 30 years old woman with skin type II  $\vartheta_m = 3\%$  and  $\vartheta_{Hb} = 1\%$ .

The final absorption values are determined by applying Equation 2 to each of the five layers considered in our model. Figure 35 (left) shows the absorption spectral curves for all the chromophores in our model, including our skin baseline absorption. Figure 35 (right) shows an example of the spectral absorption curves resulting from our model. Notice how the spectral signature of each chromophore (Figure 35, left) drives the absorption of the layers in which their concentration is higher (e.g. melanins in epidermal layers or hemoglobin in the dermal layers). Figure 36 shows the influence of water in skin, visibly reducing its spectral absorption in every layer, while increasing its reflectance.

**DECREASE OF MELANIN AND HEMOGLOBIN.** As mentioned before, the two main reasons for the decrease of chromophore concentration with age are the thinning of skin layers, and the loss of water in the hypodermis. Melanin and hemoglobin, however, suffer an additional decrease that strongly influences the appearance of aged

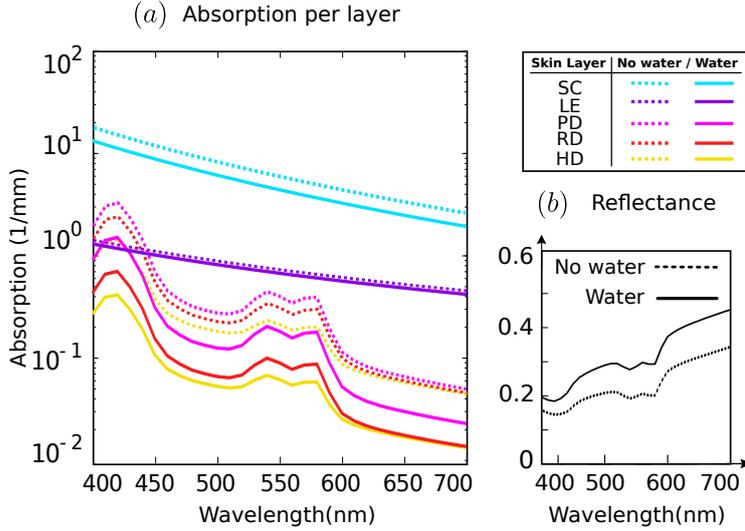


Figure 36: Influence of water on the spectral absorption of skin. Despite its low absorbance, water plays an important role in the state of skin. Notice how, by considering water concentrations, almost all layers suffer a considerable reduction in their spectral absorption, while the reflectance increases significantly.

skin [69, 40]. In the case of melanin, its rate of decline  $r_{\vartheta_m}$  has been measured at about 8% per decade, with no significant differences with respect to sun exposition or gender [89]. Given a concentration of whole blood  $\vartheta_m$  at age  $t$  in decades (Table 2) and a time interval  $\Delta t$ , we obtain the new concentration  $\vartheta'_m$  at  $t + \Delta t$  as:

$$\vartheta'_m = \vartheta_m - (\Delta t \cdot r_{\vartheta_m})\vartheta_m. \quad (3)$$

For hemoglobin, we need to consider both intrinsic and extrinsic factors. We thus model the new concentration  $\vartheta'_{Hb}$  at time  $t + \Delta t$  as:

$$\vartheta'_{Hb} = \vartheta_{Hb} - \Delta t(r_{\vartheta_{Hbi}} + \xi r_{\vartheta_{Hbe}})\vartheta_{Hb}, \quad (4)$$

where  $r_{\vartheta_{Hbi}} = 6\%$  and  $r_{\vartheta_{Hbe}} = 10\%$  are the intrinsic and extrinsic decrease rates respectively for Hb concentrations [40], and  $\vartheta_{Hb}$  is the hemoglobin concentration at time  $t$  (Table 2). Extrinsic factors like exposition to UV rays lead to a further decreased concentration of hemoglobin; they are modeled by an empirical value  $\xi \in (0, 1)$  representing the exposure of skin to aging by external agents. Figures 37 and 38 show the effect in appearance caused by melanin and hemoglobin reduction, respectively. Parameters like surface roughness and other chromophore concentrations are kept fixed to values shown in Tables 3 and 4.

#### 3.4.4 Scattering

Contrary to absorption, there is a lack of accurate spectral scattering data for human skin across a wide range of ages. To overcome this,

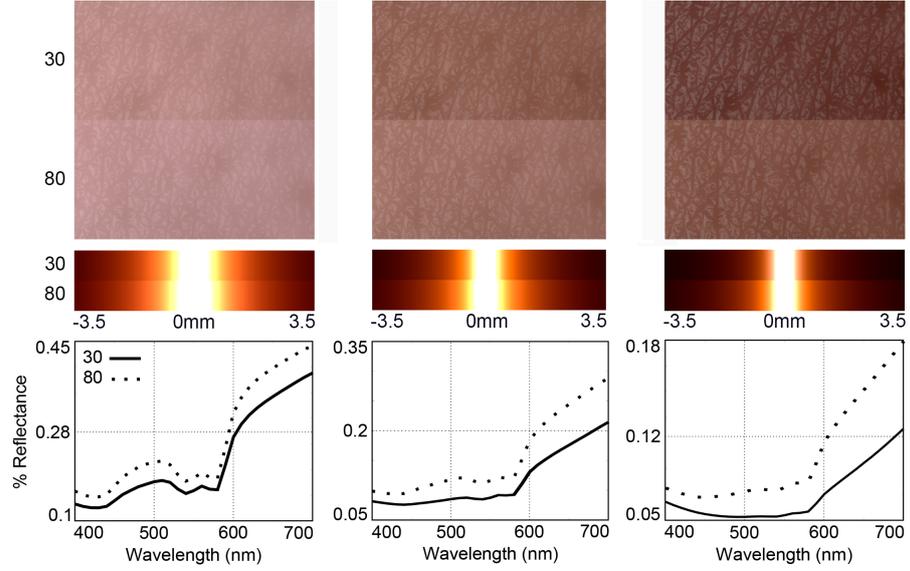


Figure 37: Influence of the decrease of melanin with age. Each column shows a patch of skin for a 30- (top) and a 80-year old (bottom), along with their corresponding scattering profiles and spectral reflectance curves. From left to right, we show light caucasian skin with 2% of melanin concentration in the epidermis, mediterranean skin (7%), and a moderately dark skin (15%). Melanin concentration in the 80-year-old samples is computed using Equation 3, while other model parameters are kept fixed (see Tables 3 and 4).

the spectral scattering behavior of skin can be instead characterized by wavelength-dependent mathematical power law functions fitted from existing sparse measured data [115, 19]. This wavelength dependence can be expressed in terms of the separate contributions of Rayleigh ( $\mu_s'^R$ ) and Mie ( $\mu_s'^M$ ) scattering as:

$$\mu_s'(\lambda) = \rho \cdot \mu_s'^R(\lambda) + (1 - \rho) \cdot \mu_s'^M(\lambda) \quad (5)$$

where  $\rho \in [0..1]$  indicates the proportion among Rayleigh and Mie scattering, and:

$$\mu_s'^R(\lambda) = \mu_s'(\lambda_r) \cdot (\lambda/\lambda_r)^{-4} \quad (6)$$

$$\mu_s'^M(\lambda) = \mu_s'(\lambda_r) \cdot (\lambda/\lambda_r)^{-\gamma} \quad (7)$$

The scaling factor  $\mu_s'(\lambda_r)$  depends on the reference wavelength  $\lambda_r = 500\text{nm}$ , and the scattering factor  $\gamma$  characterizes the wavelength dependence of the Mie scattering component. In the following, we describe the scattering for each layer independently.

**STRATUM CORNEUM AND LIVING EPIDERMIS.** The scattering of the epidermal layers is characterized by considering both the stratum corneum and the living epidermis as a whole, given the strong forward scattering characteristics of both layers and the thinness of the

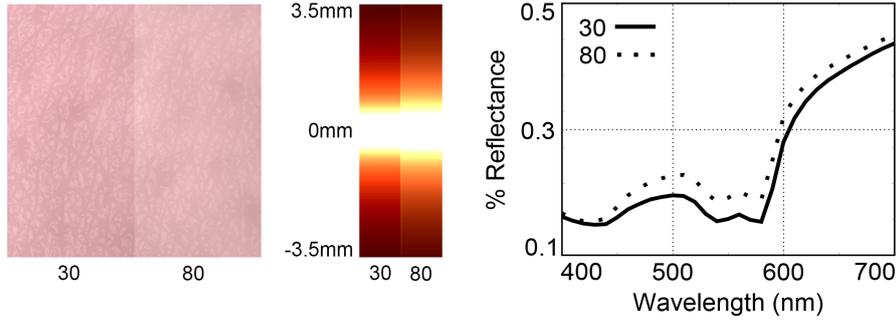


Figure 38: Influence of the decrease of hemoglobin with age. Hemoglobin concentration in the 80-year-old sample is computed using Equation 4, taking into account both intrinsic and extrinsic effects. The remaining model parameters are fixed (see Tables 3 and 4). **Left:** Patch of caucasian skin for a 30- (top) and a 80-year old (bottom), showing a decrease in hemoglobin concentration from 8.4% to 5.9% in the PD and from 7.4% to 5.2% in the RD. **Middle:** Scattering profiles. **Right:** Spectral reflectance curves.

stratum corneum [178]. Variations in melanin with age are not significant. We thus model the scattering of both layers using Equation 5 using  $\rho = 0.29$ ,  $\mu'_s(\lambda_r) = 6.67\text{mm}^{-1}$  and  $\gamma = 0.689$  [115].

PAPILLARY AND RETICULAR DERMIS. For scattering in the dermal layers, previous models [60] assume that the reduced scattering coefficient is halved with respect to the epidermal regions. While this assumption may be valid for stationary models, it falls short for skin aging. Light scattering in the dermis changes according to specific redistributions of collagen fibers, which are mainly responsible of Mie scattering. Their amount and mean diameter size differ between the papillary and reticular dermis, as shown by Branchet [32]. We thus keep Rayleigh scattering constant in our model, with  $\mu'_s(\lambda_r) = 4.36\text{mm}^{-1}$  for  $\lambda_r = 500\text{nm}$  [115], the measured fraction  $\rho = 0.41$  [193], and vary only the  $\gamma$  parameter.

To compute  $\gamma$ , we assume that the distribution of diameters  $d$  of collagen fibers follows a skewed normal distribution with mean  $\bar{d}$  and standard deviation  $\sigma$ , with an skew factor  $\zeta$ . This allows us to model the asymmetric distributions of diameters that appear with age [257, 159]. Therefore, we obtain a probability density function  $h(d)$  for a fiber diameter  $d$  as:

$$h(d) = 2 \cdot \phi(d, \bar{d}, \sigma) \cdot \Phi(\zeta d, \bar{d}, \sigma), \quad (8)$$

where  $\phi$  is a Gaussian distribution and  $\Phi$  is the cumulative distribution function of  $\phi$ . Note that this is a continuous function, while in morphometric analysis in tissue optics  $h(d)$  is generally reported using a discretized histogram. We then model  $\gamma$  for a given  $d$  as an

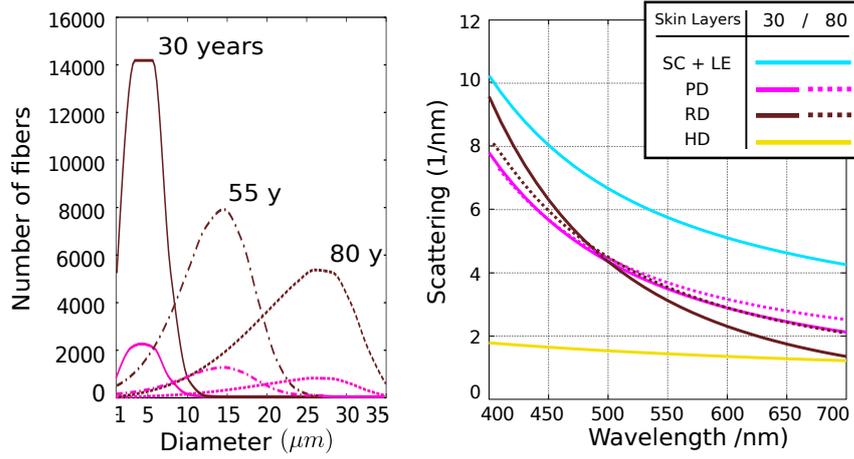


Figure 39: **Left:** Distributions of collagen fibers diameters obtained for papillary dermis skin samples, for ages 30, 55 and 80-years old. **Right:** Corresponding variations of the spectral scattering coefficient for the different skin layers.

inverse linear approximation [151] in the range  $(\gamma_{\min}, \gamma_{\max})$ , with  $\gamma_{\min} = 0.2$  and  $\gamma_{\max} = 4$  [257], with  $d \in [1.15, 35] \mu\text{m}$ .

Given a distribution  $h(d)$  of fiber diameters and their corresponding  $\gamma(d)$  values, we obtain  $\mu_s'^M(\lambda)$  from Equation 7. We finally compute  $\mu_s'^M(\lambda)$  for both the papillary and reticular dermis as:

$$\mu_s'^M(\lambda) = \frac{\sum_d \mu_s'^{M(\gamma(d))}(\lambda) \cdot h(d)}{\sum_d h(d)}. \quad (9)$$

Aging and redistribution of the collagen fibers is modeled by adjusting the parameters of the diameter distribution function  $h(d)$ . As the mean thickness of the fibers increase with age, the model adjusts  $h(d)$  around greater thickness values by proportionally adjusting  $\bar{d}$ . The redistribution of the fibers with increasing heterogeneity is simulated by progressively increasing  $\sigma$  in the Gaussian distribution. The physical meaning of varying  $\zeta$  corresponds to simulating more asymmetric distributions of fibers with time (see Figure 39). For the age range of 30-80 years, mean, variance and skew vary linearly [32].

**HYPODERMIS.** The main aging characteristic of the hypodermis is the loss of fat tissue. The spectral dependence of the reduced scattering coefficient for the hypodermis  $\mu_s'^{hd}(\lambda)$  can be approximated from existing sparse measurements by fitting to the power law [20]:

$$\mu_s'^{hd}(\lambda) = 105.06\lambda^{-0.68} \quad [\text{mm}^{-1}] \quad (10)$$

### 3.5 USING OUR MODEL

Our time-varying skin model can be plugged in any multilayered rendering framework for translucent materials (e.g. [229, 59, 50]), by com-

puting for each layer its width  $d$ , index of refraction  $\eta$ , and absorption  $\mu_a$  and reduced scattering  $\mu'_s$  coefficients for a given age. This requires to set the appropriate parameters to define the type of skin (mainly melanin concentration  $\vartheta_m$ ) and blood profusion (hemoglobin concentration  $\vartheta_{\text{Hb}}$ ), in addition to the gender of the subject, skin care  $\xi$  and the body location of the skin patch, using the formulas provided in this model. Note that, although in our work we fix other parameters, such as bilirubin, carotene and water concentration, these can also be tuned to simulate effects due to sickness or dehydration.

Our proposed model offers a high degree of control over its biophysical constituents, which is desirable for many fields such as medical applications. However, we also provide a database with the diffusion profiles for the six Fitzpatrick's skin types under a combination of different ages  $t$  and levels of skin care  $\xi$ , to simplify its use (see the supplemental material). These profiles have been computed using the MCML package [5, 229], modified to include non-planar layer interfaces (Equation 1). We calculate such diffusion profiles for an incoming pencil beam perpendicular to the surface, for wavelengths in the 400 to 700 nm range, sampled every 10 nm. Each diffusion profile is then transformed from spectral XYZ tri-stimulus values to sRGB color space for rendering. These profiles are given in tabulated form, and in the form of a sum of six Gaussians (following D'Eon et al. [51]), making them ready to use even in real-time rendering applications [52, 122].

### 3.6 RESULTS

We now show results using plausible data ranges reported on previous studies in the tissue optics literature (see Table 2). All the data used for each image is reported in Tables 3 and 4. We demonstrate the versatility of our skin aging model generating results for a wide range of skin samples and types. All the rendering results in this work have been obtained using Jensen and Buhler's hierarchical sub-surface scattering [120] implemented in PBRT [179]. We use tabulated RGB diffusion profiles obtained from multilayered simulations from the MCML package as explained in Section 3.5.

Figure 40 shows how our skin model is able to reproduce appearance changes due to aging for different types of skin. We show the subtle but noticeable changes in color due to intrinsic and extrinsic age-related changes in the tissue structure and composition. Additionally, the change of shape of the specular lobes gives a more glossy appearance to young skin. Figure 29 shows two additional examples of the change in appearance of skin due to aging, for a male and female of skin types III and I respectively. Heterogeneities are modeled using melanin and hemoglobin concentration maps [61, 123]. Close-ups for a skin type II can be seen in Figure 41. Notice the increased

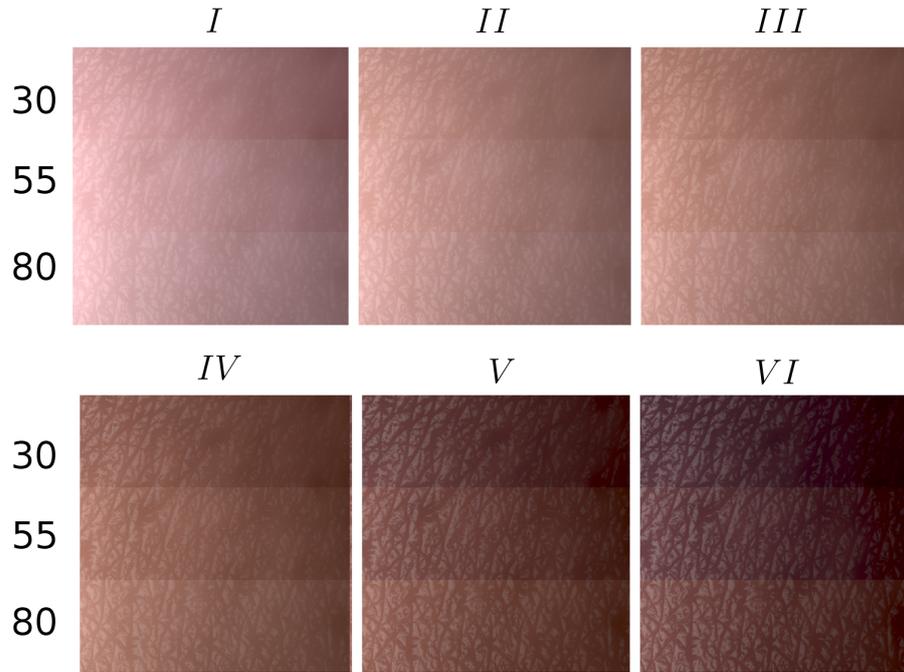


Figure 40: Results of aging a skin patch for the six types of skin according to the Fitzpatrick scale [79]. The skin type is mostly defined by the concentration of melanin  $\vartheta_m$ , as: type I,  $\vartheta_m < 3\%$ ; type II,  $\vartheta_m \in [3 - 5)\%$ ; type III,  $\vartheta_m \in [5 - 15)\%$ ; type IV,  $\vartheta_m \in [15 - 25)\%$ ; type V,  $\vartheta_m \in [25 - 35)\%$ ; and type VI,  $\vartheta_m \geq 35\%$ .

translucency in the ears, as well as the overall loss of pigment due to the lower chromophore concentration.

Last, Figure 42 shows the appearance difference for skin of the same subject belonging to different body locations: the palm and the back of the hand. It can be seen how changes are much obvious in the back: this is due both to its different skin structure and chromophore concentrations, and to its higher sensibility to extrinsic factors: the palm suffers less impact since its outermost layers are 4 – 5 times thicker.

### 3.6.1 A simplified model

The completeness and associated complexity of our model is important for certain application domains where the accuracy of the simulation plays an important role, such as medical and dermatology studies, or the cosmetic industry. Aging is the result of a complex combination of parallel processes causing changes on both the structure and composition of the skin's constitutive tissues (see Section 3.3); thus, a finer granularity than previous approaches is needed. However, less demanding applications, such as games, may capture the main optical effects of aging using a simplified version of our model. While the results will not be as accurate, this flexible tradeoff between simplicity and accuracy increases the applicability of our model.

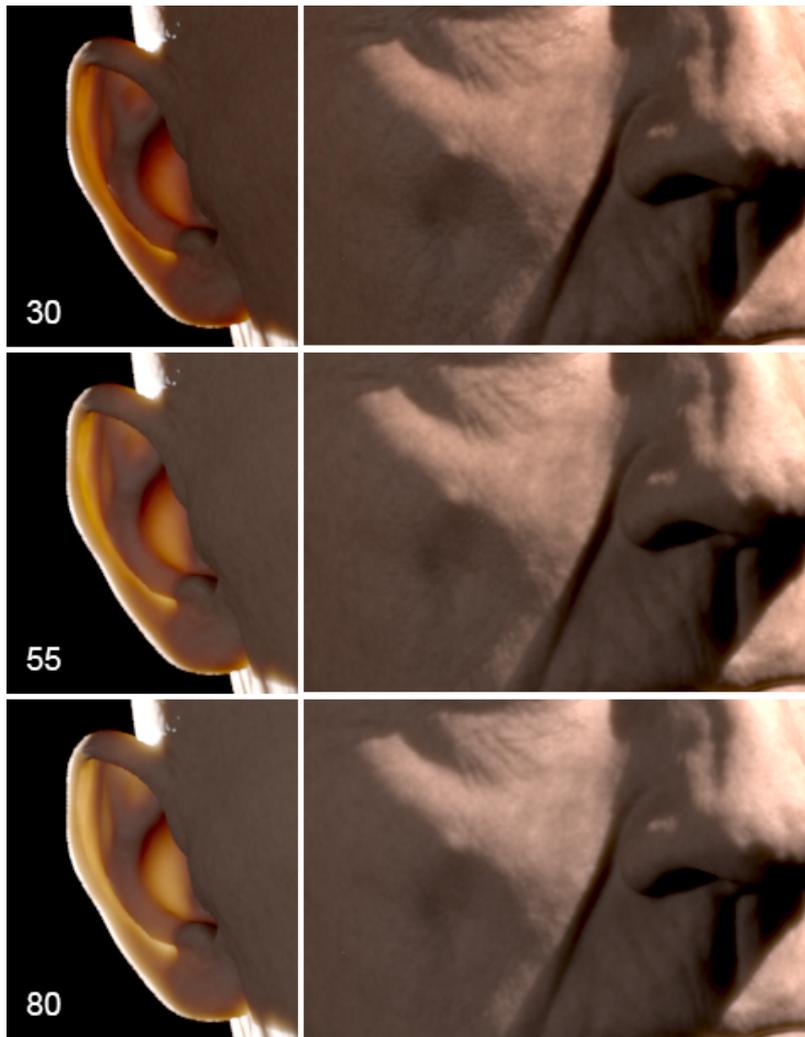


Figure 41: Differences on skin translucency and overall color between 30, 55 and 80 years old, for a subject with skin type II (note that we do not change the geometry for a more direct comparison).

We propose one such simplified version, which only uses three layers. For the epidermis, we consider a single layer containing the same chromophores and concentrations of the living-epidermis, but occupying the aggregated thickness of the living-epidermis and the stratum-corneum together. A similar approach is followed for the dermis, adding the thickness of both the papillary and the reticular dermis, assuming the characteristics of the latter. To further simplify our model, the dermoepidermal junction is considered to be flat with a constant collagen fiber density. Figure 43 (left) shows a patch of mediterranean skin (type III) for ages 30 and 80, using this simplified three-layered version, and our complete model. Since tone mapping tends to equalize the images, the false color version to the right, as well as the scattering profiles in Figure 43 (right), highlight the differences between the two; we can see the differences between the two

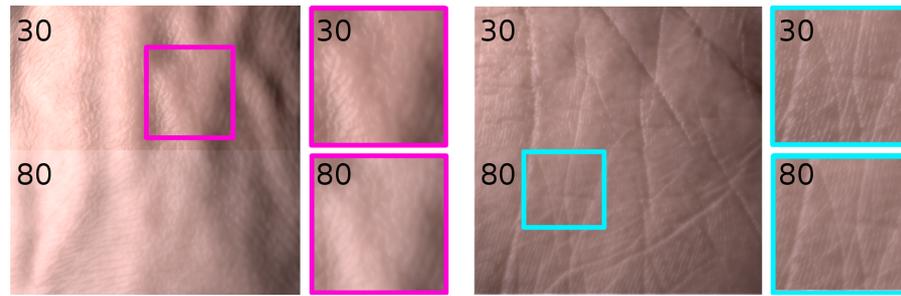


Figure 42: Aging effects for the same skin type, 30 and 80 years, but from two different body locations: back of the hand (left), and palm (right). As in other aging examples, the skin becomes less glossy and more pale. However, the palm is less sensible to extrinsic factors, due to its thick outer layers [65, 234].

models, and how the errors in the predicted scattering are larger for the simplified aged version of skin.

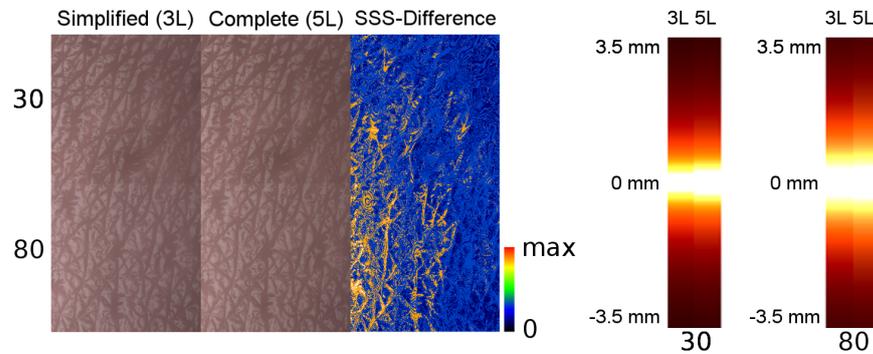


Figure 43: **Left:** Differences when using a simplified, three-layered version (3L) of our original model (5L). **Right:** Comparison of the scattering profiles pairs showing 30 and 80-year-old skin patches, using the simplified version and our complete model. Again, the differences are noticeable.

### 3.7 CONCLUSIONS AND FUTURE WORK

We have presented a comprehensive, biophysically-based skin model capable to simulate the effects of aging, without the need for skilled artistic input or subject-specific capture sessions. It is based on a detailed multilayered structure (including the hypodermis and the dermoepidermal junction), the most dominant chromophores, and scattering profiles fitted from sparse measured data (including novel aspects such as the distribution of collagen fibers). Our model can simulate the appearance changes due to intrinsic and extrinsic skin aging, and can be expressed in terms of low-level biophysical parameters or, more intuitively, with high-level parameters such as age, gender or skin type. Our tabulated diffusion profiles allow it to be easily integrated in real-time applications such as games or medical simulations.

Figure	Input Model Parameters				
	Age	Gender	Body Location	Skin Type	$\xi$ (External Agents)
29 - a	80	Male	Face	III	0.5
29 - b	30	Male	Face	III	1
29 - c	80	Female	Face	I	0.5
29 - d	30	Female	Face	I	1
37	30 (80)	Female	Forearm	I, III, IV	1
38	30 (80)	Female	Forearm	I	0
40	30 (55, 80)	Female	Forearm	I,II,III,IV,V,VI	1
41	30 (55, 80)	Male	Face	III	1
42-Left	30 (80)	Male	Hand - back	III	1
42 Right	30 (80)	Male	Hand - palm	III	1
43 (5L)	30 (80)	Female	Forearm	III	1
43 (3L)	30 (80)	Female	Forearm	III	1

Table 3: Input parameters for the images shown along the chapter. Skin type is specified by the Fitzpatrick scale. Skin care  $\xi$  ranges from 0 to 1. Aged samples are represented within parenthesis.

Although our skin model requires more structural complexity than previous approaches, it can be used to obtain a diffusion profile similar to previous works, so this required complexity does not affect the rendering stage: As described, our model can be directly plugged into any existing rendering system capable of handling diffusion-based subsurface scattering without further modifications.

Our model offers a potential range of applications beyond graphics. For instance, its predictive nature can be useful in fields such as cosmetics, dermatology, or tissue optics. We are currently bounded by existing measurements and collected data. For instance, there is no existing bio-physical data about heterogeneities, which could nevertheless be added to our model using textures modulating the spatially-varying distribution of chromophores or the albedo of skin, similar to previous works (e.g. [123, 61]). An interesting avenue of future work would be modelling the time-varying nature of these heterogeneities based on a data-driven approach. We hope that our work helps stimulate the creation of more complete datasets, although this will be a long-term effort.

Figure

Intermediate and Output Model Parameters													
	SC	LE	Thickness (mm)				Main Chromophores Concentrations (%)				R <sub>z</sub> (μm)	Surface Reflection Sebum prod. ( $\frac{\text{mg lipid}}{10\text{cm}^2/3\text{h}}$ )	m
			PD	RD	HD	$\phi_m$ (LE)	$\phi_{\text{Hb}}$ (PD)	$\phi_{\text{Hb}}$ (RD)	$\phi_{\text{Hb}}$ (RD)				
29 - a	.02	.051	.13	1.28	5.9	7	2.4	1.8	151	1.69	0.143		
29 - b	.02	.08	.18	1.82	5.9	7	2.4	1.8	78	2.4	0.068		
29 - c	.02	.06	.13	1.27	5.9	2	1.2	0.9	151	0.85	0.148		
29 - d	.02	.08	.18	1.82	5.9	2	1.2	0.9	78	2	0.069		
37	.02 (c)	.08 (c)	.18 (c)	1.82 (c)	5.9 (c)	2 (1.2), 7 (4.2), 15 (9)	6 (c)	4.5 (c)	78 (c)	2.4 (c)	0.068 (c)		
38	.02 (c)	.08 (c)	.18 (c)	1.82 (c)	5.9 (c)	1 (c)	8.4 (5.8)	7.4 (5.216)	78 (c)	2.4 (c)	0.068 (c)		
40	.02	.08 (.07, .06)	.18 (.15, .13)	1.82 (1.55, 1.27)	5.9	1 (0.8, 0.6), 3 (2.4, 1.8) 5 (4.3), 10 (8, 6) 20 (16, 12), 30 (24, 18)	6 (3.6, 1.2)	4.5 (2.73, .9)	78 (103, 151)	2.4 (2.4, 1.7)	.0768 (.093, .143)		
41	.02	.08 (.07, .06)	.18 (.15, .13)	1.82 (1.55, 1.27)	5.9	7 (5.6, 4.2)	2.4 (1.4, .5)	1.82 (1.1, .4)	78 (103, 151)	2.4 (2.43, 1.69)	.068 (.093, .143)		
42-Left	.029	.08 (.06)	.06 (.04)	.65 (.45)	5.9	5 (4.3)	6 (3.6, 1.2)	4.5 (2.7, .9)	78 (151)	2.4 (1.69)	.068 (.143)		
42-Right	.173	.48 (.35)	.08 (.05)	.80 (.55)	5.9	5 (4.3)	6 (3.6, 1.2)	4.5 (2.7, .9)	78 (151)	2.4 (1.69)	.068 (.143)		
43 (5L)	.02	.08 (.06)	.18 (.13)	1.82 (1.27)	5.9	7 (4.2)	6 (1.2)	4.5 (.9)	78 (c)	2.4 (c)	.068 (c)		
43 (3L)	-	.10 (.07)	-	2 (1.4)	5.9	7 (4.2)	-	4.5 (.9)	78 (c)	2.4 (c)	.068 (c)		

Table 4: Intermediate and output model parameters (aged samples are represented within parenthesis). The amplitude of the dermoepidermal junction is  $A_{x,y} = 100 \mu\text{m}$  for all the 30-year old samples, and  $A_{x,y} = 0 \mu\text{m}$  for 80-year old ones. Some values are the same for all images and not reported in the table. These are: Eumelanin concentration in melanosomes is 80 g/L. Pheomelanin concentration in melanosomes 12 g/L. Hb (oxy+deoxyb) in blood is 150 g/L.  $\phi_{\text{Hb}}$  is always set to 5% in HD. Bilirubin in blood is 0.05 g/L.  $\beta$ -carotene in blood is  $7.0e^{-5}$  g/L; in LE and SC, it is  $2.1e^{-4}$  g/L. The oxy-deoxy ratio of hemoglobin is 75 %. Water content, given in gram water per gram wet skin, is 5, 20, 50, 70, and 70 for SC, LE, PD, RD and HD respectively. Values labelled as c are kept constant for the purpose of isolating the influence of the studied parameter.

## 3.8 APPENDIX A. MODELLING OF SEBUM

We include the effect of skin surface sebum in the specular reflection, by assuming that the skin is composed of consecutive pyramidal holes whose edges' length and slope are defined by the microfacet properties derived for the surface roughness. Then we fill the holes with a certain volume of sebum. As humans age, sebaceous glands become less active. The thickness of the sebum film can vary between up to 3  $\mu\text{m}$  in young skin to 1.69  $\mu\text{m}$  in the elderly [180]. Thus, the new value for  $\alpha$  is computed as

$$\alpha' = \psi_s \alpha_{sr} \quad (11)$$

with  $\psi_s = 1 - a_s/a_t$ ,  $a_t$  is the projected area of a facet of the pyramidal hole, and  $a_s$  is the portion of that area covered by sebum. Using trigonometry, we get:

$$\frac{a_s}{a_t} = \frac{l_s^2}{l^2} \quad (12)$$

with  $l$  the length of the skin profile scan, and

$$l_s = \sqrt[3]{\frac{3V_s}{2 \tan \alpha_{sr}}} \quad (13)$$

where  $V_s$  is the volume of sebum computed for an specific age, based on the data of Pochi and colleagues [180].

Chromophore	Absorption coefficient
Eumelanin	$\mu_a^{em}(\lambda) = 6.6 \times 10^{10} \times \lambda^{-3.33}$
Pheomelanin	$\mu_a^{pm}(\lambda) = 2.9 \times 10^{14} \times \lambda^{-4.75}$
Oxy-Hb	$\mu_a^{HbO_2}(\lambda) = \frac{\epsilon_{HbO_2}(\lambda)}{66500(\text{g/mole})}$
Deoxy-Hb	$\mu_a^{Hb}(\lambda) = \frac{\epsilon_{Hb}(\lambda)}{66500(\text{g/mole})}$
Carotene	$\mu_a^{car}(\lambda) = \frac{\epsilon_{car}(\lambda)}{537(\text{g/mole})}$
Bilirubin	$\mu_a^{bil}(\lambda) = \frac{\epsilon_{bil}(\lambda)}{585(\text{g/mole})}$
Water	$\mu_a^{H_2O}(\lambda) = \epsilon_{H_2O}(\lambda)$
Baseline	$\mu_a^{base}(\lambda) = 7.84 \cdot 10^7 \cdot \lambda^{-3.255}$

Table 5: The partial contribution of each chromophore to the final spectral absorption coefficient  $\mu_a^l(\lambda)$  of a given layer, is obtained by multiplying the spectral extinction coefficient  $\epsilon(\lambda)$  of the chromophore, generally measured in ( $\text{cm}^{-1}\text{M}^{-1}$ ), by its concentration  $c_i$  in that layer. Chromophore absorptions are obtained from both measured data and previous empirical fits. For  $\mu_a^{em}$  and  $\mu_a^{pm}$ , we use empirical fits from [60]. For  $\mu_a^{HbO_2}$ ,  $\mu_a^{Hb}$ ,  $\mu_a^{car}$  and  $\mu_a^{bil}$  we use equations from the BioSpec model [135]. Last, we employ measured values for  $\mu_a^{H_2O}$  taken from [181] and baseline absorption for depigmented skin from Saidi [192].

Skin layer absorption coefficient	Reformulated absorption equation
$\mu_a^{sc}(\lambda)$	$[(0.1 - 0.3 \times 10^{-4} \lambda) + 0.125 \frac{\lambda}{10} \times \mu_a^{base}(\lambda)] \times (1 - f_{H_2O}) + \mu_a^{H_2O}(\lambda)$
$\mu_a^{le}(\lambda)$	$[v_m \times (\mu_a^{eu}(\lambda) + \mu_a^{ph}(\lambda)) + (1 - v_m) \times (\mu_a^{car}(\lambda) + (1 - c_{car}) \times \mu_a^{base}(\lambda))] \times (1 - f_{H_2O}) + \mu_a^{H_2O}(\lambda)$
$\mu_a^{pd}(\lambda)$	$[v_{Hb}^{pd} \times (\mu_a^{oxy-Hb}(\lambda) + \mu_a^{deoxy-Hb}(\lambda) + \mu_a^{bil}(\lambda) + \mu_a^{car}(\lambda)) + (1 - v_{Hb}) \times \mu_a^{base}(\lambda)] \times (1 - f_{H_2O}) + \mu_a^{H_2O}(\lambda)$
$\mu_a^{rd}(\lambda)$	$[v_{Hb}^{rd} \times (\mu_a^{oxy-Hb}(\lambda) + \mu_a^{deoxy-Hb}(\lambda) + \mu_a^{bil}(\lambda) + \mu_a^{car}(\lambda)) + (1 - v_{Hb}) \times \mu_a^{base}(\lambda)] \times (1 - f_{H_2O}) + \mu_a^{H_2O}(\lambda)$
$\mu_a^{hd}(\lambda)$	$\mu_a^{H_2O}(\lambda)$

Table 6: Wavelength-dependent absorption equations for each layer of our skin model. All equations fit data absorption in  $\text{mm}^{-1}$ . Different volume fractions  $v_p$  and  $v_r$  of blood present in tissue are estimated for the papillary and reticular dermis respectively (adapted from [155, 135]).

PERCEPTION OF COMPUTER GENERATED CLOTH

---

In this chapter we conduct a set of perceptual experiments to explore some of the main factors that contribute to the perception of realistic cloth. Plausible cloth simulation is of great importance for animating compelling and realistic computer generated scenes. To this end, both complex dynamics simulations and rendering algorithms are required. However, many real cloth fabrics exhibit very complex interaction with light and their mechanical properties can be difficult to specify.



Figure 44: Different fabrics have both different visual appearance and mechanical properties. We create replicas of several common woven fabrics, like the cotton or silk shown in the image, covering a wide range of movements in a set of video stimuli.

This often leads to cumbersome processes in which manual parameter tweaking by skilled technical artists is needed, together with time-consuming computations. Our goal is to provide insights into how efficiency can be maximized without sacrificing plausibility. Using real video footage of several fabrics that span a wide range of visual appearances and dynamic behaviors, and their simulated counterparts, we explore their visual attributes and the interplay of visual

appearance and dynamics in the perception of computer generated cloth.

#### 4.1 INTRODUCTION

Recent advances in physical simulation and rendering allow for the creation of increasingly realistic and convincing shots. Most productions now include scene elements that require costly simulations and rendering processes. Furthermore, parameters are often hard to specify and tune, which requires time consuming manual input. However, many elements of these scenes may go unnoticed by the viewer, so a better understanding of the perceptual factors would facilitate more efficient use of available resources.

In this work, we investigate the level of fidelity needed in order to achieve plausible animations of cloth. Do both appearance and dynamics need to be simulated with a high degree of accuracy? By studying how the interplay of visual appearance and dynamics can affect the perception the viewer, we determine whether a fully accurate shader combined with a simplified dynamics simulation might be sufficiently realistic for most practical purposes, and vice versa. Depending on the particular mechanical properties and visual attributes of the types of fabric being depicted, efficiencies may thus be gained.

As described in [8] and discussed in more depth in this article, videos of seven different real cloth samples made of different fabrics that span a wide range of visual appearances and dynamic behaviors have been recorded. Photo-realistic synthetic versions have also been created that emulate the real cloth sequences as closely as possible (see Figure 44) From these seven ground-truth animations, all possible combinations of mismatching appearance and dynamics are rendered, yielding a 7x7 stimulus matrix with the correct solutions along the diagonal. For this, we rely on methods and techniques used in a real production environment. Many different options exist to generate the replicas. Cloth is a very active research area with many open problems in rendering, simulation, or the combination of both. There is a great heterogeneity in the algorithms, with some simulation frameworks designed at yarn level [42], and many different representations used for rendering (e.g. thin layers [112, 191]; volumes and / or explicit fibers [200, 248, 202, 132, 252]), in addition to many scattering models [149, 256, 54, 132]. A thorough comparison of all existing methods and their combination would be of great interest for the community but is beyond the scope of this work.

Three sets of perceptual experiments are presented that explore the perception of the cloth videos and animations. The first two experiments were initially presented in [8] and we provide more details and in-depth analysis of the results in this article. The aim is to study how well participants can identify different types of fabric, and how sen-

sitive they are to mismatches in the visual attributes and mechanical properties. We find that appearance is the dominant factor for most fabrics in both cases. Finally, in a new experiment, we study the effects of dynamics and appearance in isolation in order to determine how certain key attributes related to each factor (e.g. perceptual stiffness for the dynamics or glossiness for the appearance) are contributing to the dominance of appearance. We find that participants are better able to discriminate between different appearance attributes and are again less sensitive to the dynamics.

Our studies are the first to explore the relative importance of appearance and dynamics on the perception of photo-realistic animated cloth. Our results may be useful to guide a better distribution of resources when planning shots involving cloth simulations, or could affect how shot approvals are done. For example, our results indicate that the perception of a given fabric is influenced by its visual appearance to a much stronger extent than by its dynamics. Therefore, viewing the simulation without a reasonable depiction of the final shader to be employed would not be sufficient to predict the plausibility of the final result. Conversely, if the rendering of the cloth is sufficiently realistic, there can be a higher level of confidence that the final animation will also be so.

#### 4.2 RELATED WORK

In this work, we continue the studies of Aliaga et al. [8] to evaluate the contribution of appearance vs. dynamics for a range of different fabrics and viewing distances. Closely related to this goal, McDonnell et al. [153] evaluated the perceptual impact of different geometric and image-based LOD representations of animated cloth, and guidelines for developing crowd systems with realistic clothed humans were presented. Sigal and colleagues [208] developed a control space for garment simulation by mapping the complex parameters of any physical simulator to several intuitive predictable perceptual parameters, learned from a set of experiments. Also focused in dynamics, a recent paper [26] explores whether humans have an invariant representation of the mechanical properties (mass, stiffness) of fabrics under varying external forces in dynamic scenes (videos of oscillating wind with different strength).

*Appearance:* Many approaches focus on generating visually plausible materials. Pellacini et al. [176], Westlund and Meyer [232] and Ferwerda et al. [74] developed psychophysically-based models for gloss perception. Wills et al. [236] performed similar experiments to derive a perceptual space of measured BRDFs. Vangorp et al. [223] evaluated the influence of shape and illumination on surface gloss perception, showing how objects with smooth bumps provide more cues than simpler ones like spheres. Other studies include translucency

and subsurface scattering [81, 91], or surface texture and reflectance [47, 75, 119]. From a generic perspective, Serrano et al. [204] recently proposed an intuitive control space for material appearance, mapping perceptual attributes to an underlying PCA-based representation of BRDFs.

Fleming and colleagues [82, 83] conducted reflectance matching experiments to demonstrate that people can recognize material properties more accurately under natural illumination than under artificial lights. Other examples focus on perceptually guided global illumination [162, 211]. Ramanarayanan and colleagues [182, 183] evaluated the effects of changes in environment lighting over different shapes and materials. Through several transformations in the illumination maps, such as warping or blurring, they found that many objects had the same appearance (they are visually equivalent) when illuminated by both transformed and original maps. Similar studies evaluated the effect of approximations in illumination on the perception of complex animated scenes [118] or materials [136]. Other studies evaluate not only approximations but the effect of intended artistic stylizations in shape, material or rendering style over human faces [154, 246].

*Dynamics:* Some studies have evaluated the effects of degrading or distorting physically-based simulations on the perceived plausibility of animations, e.g., [173, 243, 100]. Similar studies have also been conducted in the context of cartoons [84]. Other works focus on collisions; O’Sullivan et al. [172] developed a model of collision perception for real-time animation, while Dingliana and O’Sullivan [57, 171] examined the perception of detail simplifications for LOD rigid-body physically-based animation. Also, several recent articles in perception of liquid viscosity explore the roles of optical properties and motion cues [3, 2, 129].

Some other works evaluate the perception of dynamics on animated characters. Reitsma et al. [185] studied the visual tolerance of ballistic motion for character animation, finding that horizontal velocity errors are more detectable than vertical. Vicovaro et al. [224] evaluated the plausibility of altered throwing motions. Finally, Hoyet et al. [108] conducted several psychophysical experiments to measure the perceived realism of pushing interactions, evaluating the influence of timing errors or force mismatches.

### 4.3 STIMULI CREATION

In order to cover a reasonable range of different fabric appearances and dynamics, we chose seven commonly used woven cloths. In approximate order of more to less stiff, the selected fabrics are: Burlap (also commonly known as Sackcloth), Canvas, Denim, Linen, Cotton, Polyester satin and sheer Silk. We acquired real samples of all

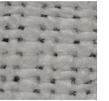
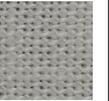
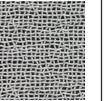
	Burlap	Canvas	Bull Denim	Linen Solid	Cotton	Polyester Charmeuse	Iridescent Silk
a)							
b)	100% Jute	100% Cotton	100% Cotton	100% Linen	100% Kona® Cotton	100% polyester	100% silk
c)	206.82	345.84	361.44	308.54	145.79	139.01	27.12
d)	0.69	0.56	0.76	0.43	0.24	0.36	0.1
e)	Plain	Plain 2x1 basket	Twill 1 warp x 3 weft, 1 warp offset	Plain	Plain	Satin 4 warp x 1 weft, 2 warp offset	Crepe de chine
f)	4, 4	30, 11	26, 20	17, 12	27, 23	60, 60	42, 42
g)	1, 1	0.33, 0.33	0.38, 0.5	0.3, 0.5	0.24, 0.24	0.17, 0.17	0.08, 0.08

Table 7: Manufacturer details for the real fabrics used in the experiments, which are used to simulate the CG replicas. a) 5x5mm closeup; b) Composition; c) Weight ( $\text{g}/\text{m}^2$ ); d) Thickness (mm); e) Weaving Pattern; f) Threads/cm: warp, weft; g) Thread diameter: warp, weft (mm).

of them, cut into squares of 1x1 meters. They all are of roughly the same albedo, in order to avoid color being a confounding factor for the experiments (see Figure 45).

We then recorded videos of all the fabrics in a studio with diffuse black walls, floor, and roof, using two spot lights placed at about 45 degrees from the focal plane (Figure 46, left). Every piece of cloth was recorded while draping over a flat swivel stool which then spins, in order to show as many mechanical and dynamic properties of the fabric as possible (e.g., shape of the folds, angle of swing). View-dependent appearance features for each fabric are also visible in this way. We ensured that the movement was as similar as possible for each fabric.

To create computer generated replicas of the reference fabrics, we needed to emulate both the appearance and the dynamics. Note that appearance refers to the spatially varying reflected radiance of the cloths, which depends on several factors such as the texture pattern or the optical properties of the fabrics (e.g.: albedo or surface scattering). Thus, all pieces of cloth were rendered using path tracing with deferred shading [66], simulating rough dielectric materials with dif-

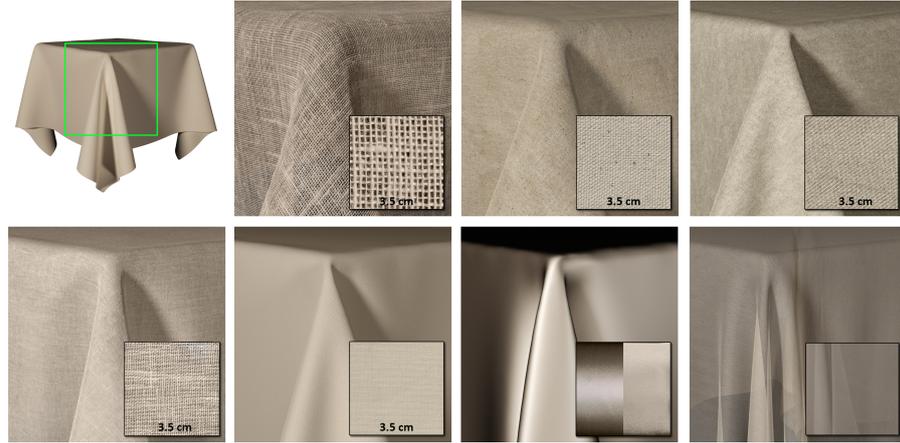


Figure 45: Comparison between the real fabrics and the CG replicas (from l-r): burlap, canvas, denim, linen, cotton, polyester satin and sheer silk. Images are renderings and the insets show close-ups of the real fabrics in the first five cases. For polyester satin and sheer silk, the weaving pattern is too small to notice at normal viewing distances. For polyester satin the inset shows the fabric wrapping a cylinder along the warp and weft directions to show the viewing and lighting dependent anisotropic highlights. For sheer silk, the inset shows the real fabric draping the swivel stool.

fuse transmittance, together with albedo, bump and opacity textures. For these, a set of close-up pictures perpendicular to the fabrics was taken to generate tileable seamless textures representing patches of 30x30 cm. The only exception was polyester satin; given its more anisotropic reflectance and color shifts, we relied on the empirical microcylinder model of Sadeghi and colleagues [191]. This approach represents the state of the art for rendering specular woven cloths when they are represented by a triangle mesh as a thin layer (Irawan and Marschner [112] only handle specular reflection). Other options such as the works of Khungurn et al. and Zhao et al. [132, 248] really represent the state of the art for realistic cloth appearance, but were discarded by several reasons: the need of a computer tomography scan for capturing the micro volumes, the need of a real lighting set up to fit the volumetric rendering parameters of the model, and the fact that such optimization to derive the scattering parameters has only been proved for static images. The appearance of the final CG replicas can be found in Figure 45.

Our choice of methods for the simulation part are also currently used in production. Works like Cirio et al. [42] are able to yield yarn-scale detail, but have not been designed for higher-scale simulations (like our draping garment), and the parameters that need to be set are difficult to obtain from real pieces of cloth, to create accurate replicas. Thus, the dynamics of the different fabrics were simulated by modeling the cloth as a triangular mesh, along with proximity forces to prevent primitives near each other from colliding, as proposed by

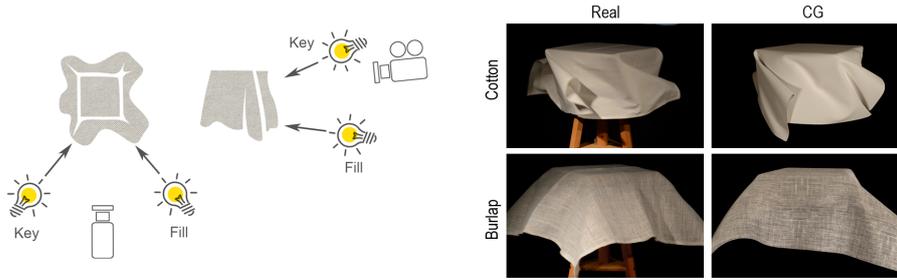


Figure 46: *Left*: Lighting studio setup for capturing the the real cloth videos, (bottom and side views). *Right*: Comparison between the movements of the real cloth samples and the CG replicas. The first row shows the cotton rotating at the maximum speed. The second row shows the burlap at the frame just before starting to stabilize. Note that the real and CG samples are rotating in the same direction in these images just for comparison, but do so in opposite directions during the experiments to avoid exact image matching. To emulate the cloth motion, we paid special attention to the number, size and shape of the folds created (both at static and dynamic frames), the amount of bouncing, the effect of air forces, and the maximum height and width reached when rotating.

Baraff and Witkin [17]. Similarly, we use additional constraints for cloth-object collisions. If continuous time collisions remain after the initial solve, we rely on the robust collision algorithm from Bridson et al. [33], augmented by a fail-safe that cancels impact while maintaining sliding motion [103]. We relied on physical parameters given by the manufacturer<sup>1</sup> when available, such as density and thickness (e.g., burlap weighs  $207\text{g}/\text{m}^2$  with  $0.69\text{mm}$  thickness, while the values for silk are  $207\text{g}/\text{m}^2$  and  $0.69\text{mm}$ ; see Table 7 for more details); all the remaining parameters were manually adjusted to obtain a result as close as possible to the real cloth properties (see Figure 46, right). We considered using an automatic method for estimating the material properties from real cloth, as Bouman et al. [31] do for videos. The problem is that the high level parameters estimated, like stiffness or density, do not have a direct connection with the simulation models employed in practice. In fact, this is a hard problem that Sigal and colleagues [208] attempted to solve in a concurrent work, where a model was built to map perceptual parameters to low level parameters specific in each simulation engine. In addition, the range of movements, friction and self collisions exhibited in our videos made us discard any automatic method, opting to have full control with the goal of achieving simulations as close as possible to the real replicas, using affordable tools and resources in a production pipeline. The swivel stool was not included due to limitations in our rendering system at the time when the tests were being carried out. While this may affect the perception of the most transparent cloth (silk), many other

<sup>1</sup> Michael Levine, Inc., <http://www.lowpricefabric.com/>

cues for transparency are already obtained through the overlapping moving folds of the cloth itself.

We then rendered all possible combinations of appearance and dynamics, yielding  $7 \times 7 = 49$  videos (six seconds each) replicating the movement in the recorded video (Figure 47). Thus for each row (column) of the matrix, only one rendered video matches the appearance with the correct dynamics. In addition, to study the effect of viewing distance on the perception of mismatched properties, we rendered all of the stimuli at three different camera distances, resulting in resolutions of  $1728 \times 1123$ ,  $1000 \times 650$  and  $520 \times 338$  from close to far viewing distances respectively. These resolutions were set empirically with the goal of having a perceptual level of detail, also taking into account the latter arrangement in the screens within the experiment setup and the screens' size and resolution.

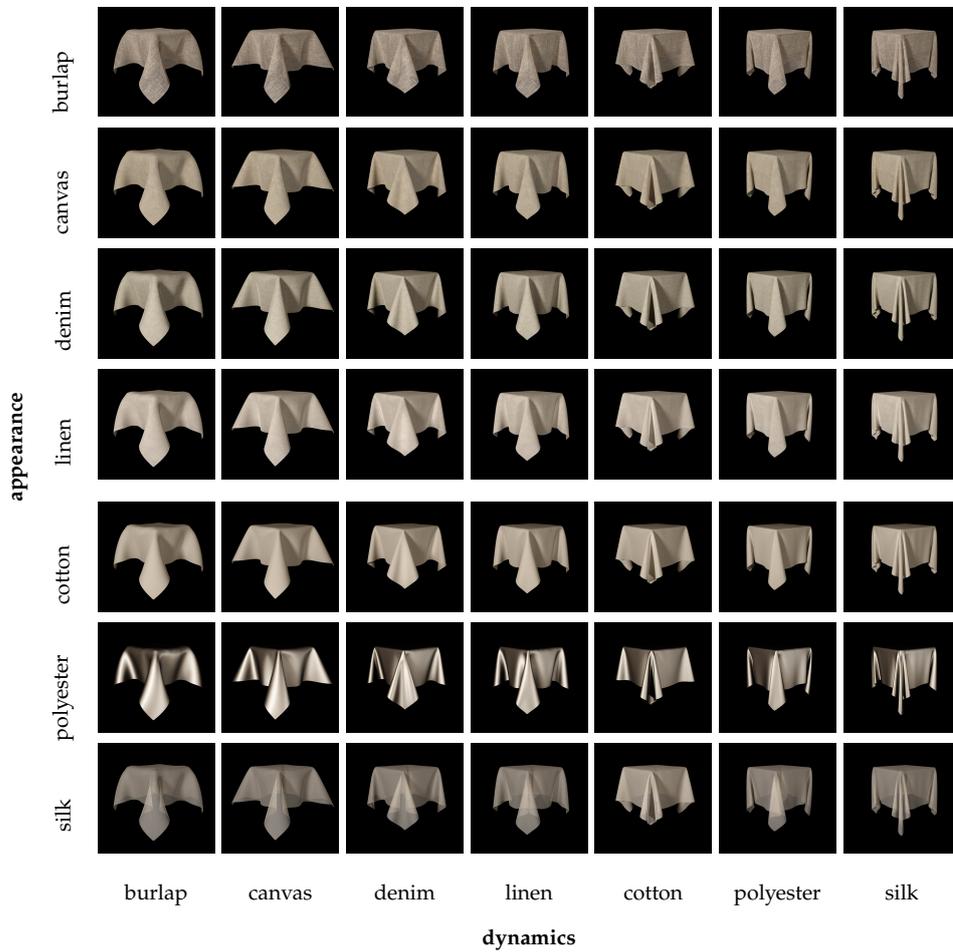


Figure 47: All CG stimuli from Experiments One and Two. All appearance (vertical axis) and dynamics (horizontal axis) combinations for our seven fabrics are shown at a fixed viewing distance. In this figure, we chose a frame of the video where all the fabrics are stabilized so that the folds and drape can be better appreciated without motion. Note that we rendered all videos with the swivel stool rotating in the opposite direction from the real videos, to avoid participants basing their judgments on exact visual matching.



Figure 48: Two screen layout for Experiment One. On the left, the navigation screen with the seven real (ground truth) reference fabrics. Each thumbnail has a radio button for selection and a replay button. On the right, the CG cloth that is currently being displayed.

#### 4.4 EXPERIMENTS

In the introduction, we asked whether both appearance and dynamics need to be simulated with a high degree of accuracy for plausible cloth simulation. To answer this question we conducted a set of perception experiments. The first two experiments were both carried out in counterbalanced order with 63 naive participants (34F/29M, aged 18–27) who had varying levels of experience in computer graphics. We found that participants were more sensitive to the appearance of the fabrics than to the dynamics, so we then carried out a third experiment with 15 naive participants (9M/6F) to further explore some perceptual attributes of the various fabrics.

##### 4.4.1 Experiment One: Ground Truth comparison

The goal of the first experiment is to answer two questions: (i) How effective are our cloth simulations at emulating the appearance and dynamics of the real stimuli; and (ii) are dynamics and appearance equally important when comparing photo-realistic cloth animations with videos of the real fabrics?

**Method:** Two equally calibrated screens of the same model were used for the experiment (Dell U2311H IPS FullHD 23”) arranged in parallel at about 50 cm far from the participants. At the start, all participants familiarized themselves with all real stimuli on the left screen, by clicking on a thumb-nail reference to view the corresponding six-second video, which they could repeat as many times as needed. These real stimuli remained available for viewing throughout the experiment. On the right screen, one of the 147 rendered videos is shown, and the participant is asked the question: ‘Which of the reference cloths on the left best matches the one on the right?’. The participant can answer by choosing any of the seven reference cloths shown in thumbnails on the left (Figure 48), which may be replayed as often as needed (with no time limit). Each time a reference video

is replayed at full resolution on the left, the current CG replica that is being evaluated is played on the right for comparison purposes. Both videos are synchronized, but the cloths rotate in opposite directions to discourage exact visual pattern matching.

In order to avoid fatigue, each participant viewed a subset of the stimuli, distributed so that each video is seen by 45 different people, and each person sees 105 different samples of the total set of 147. The experiment took between 25 and 45 minutes, including a 5-min break. An option for showing each participant the whole set of stimuli could have been to limit the amount of replays and/or the viewing time per trial. However, the average number of repetitions was close to 0 except for the first  $\approx 10$  trials (typically 1-2 repetitions for those cases), and the viewing time was the minimum to have enough dynamics cues to judge (the piece draping and rotating). Then, increasing the total length of the test in about a 40 percent would have drawn participants attention.

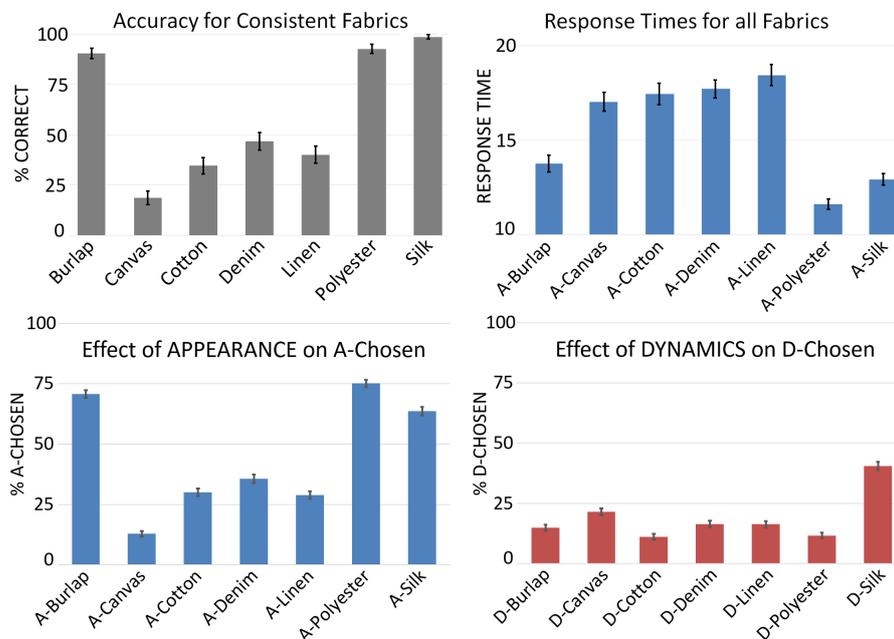


Figure 49: Experiment 1 results (means with standard error bars). Top Left: correct responses for fabrics with consistent appearance and dynamics, i.e., how often the correct real cloth was matched by the corresponding consistent simulation; Top Right: response times (RT) for fabrics with different Appearance (A); Bottom: results for fabrics with mis-matching Appearance (A) and Dynamics (D), showing how often the real cloth chosen was determined by either the corresponding A or D of the mis-matched simulation.

**Results:** We first summarized all participant responses in a *Multi-way Frequency Table*, with frequency counts for each combination of: *Distance* (close, medium, far), *Appearance* (A-Burlap, A-Canvas, A-Cotton, A-Denim, A-Linen, A-Polyester, A-Silk), *Dynamics* (D-Burlap

– D-Silk) and *Response*  $\times$  7 (R-Burlap – R-Silk). We then analyzed these data using Log-Linear Analysis and found that the best model that fits the observed data is (Appearance, Response), (Dynamics, Response), with a likelihood ratio of  $\chi^2(216) = 218.57, p = 0.4382$ . This means that there was a main effect of both Appearance and Dynamics on which real fabric was chosen. However, the distance from the camera appears not to have a significant effect on the responses.

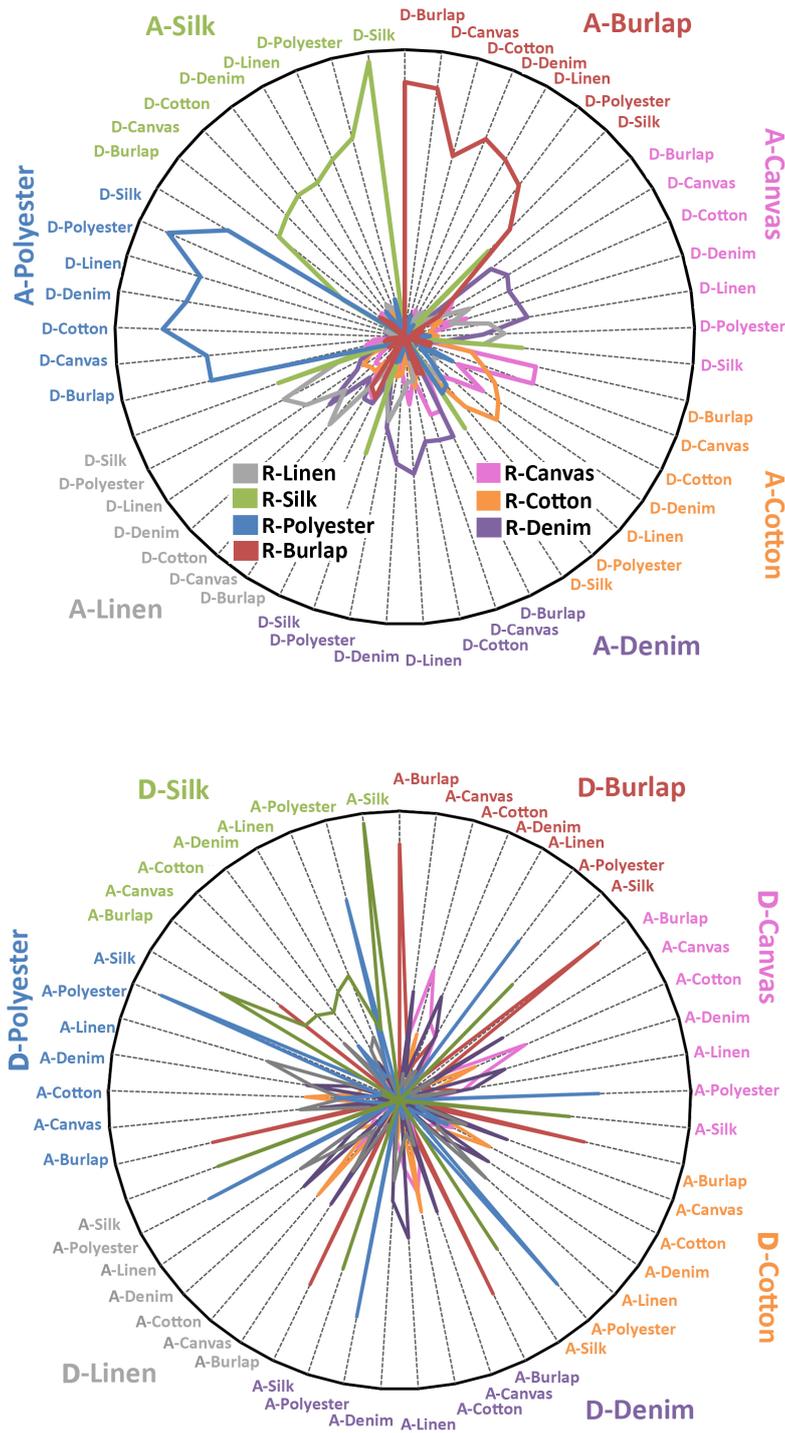


Figure 50: Experiment 1 results, sorted by Appearance (top) and Dynamics (bottom), summarized as a radar graph and collapsed over distance. The colored areas in the graph represent how often each Response was given for the Appearance/Dynamics combinations depicted on the perimeter.

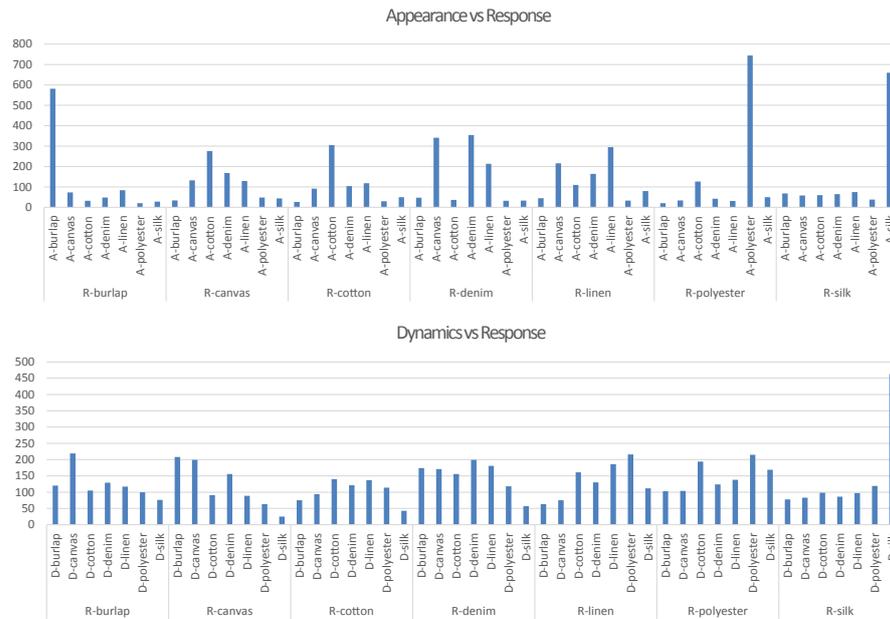


Figure 51: Experiment 1 results, showing how often each Response ‘R-fabric’ was given for the different appearances (top) and dynamics (bottom). It is interesting how appearance is very well matched in average, only getting more confusion for cotton-made fabrics (Canvas is wrongly identified as Cotton, and both Denim or Linen are perceptually close to Canvas). For the case of dynamics, even if Silk, Polyester, Denim and Cotton are well matched, the deviation is higher than in the case of appearances, with Burlap being identified as Canvas and vice versa, or Linen wrongly identified as Polyester.

To explore these effects further, we ran several unbalanced ANOVAs using the Variance Estimation and Precision function in STATISTICA™. This was necessary as there were unequal numbers of observations for each combination of variables. As described in [157], we report ANOVA results (univariate tests of significance) using Type 1 sums of squares for the unbalanced scores. In each case we tested for main and interaction effects of fixed independent variables *Distance*, *Appearance* and *Dynamics* (Fixed) and random independent variable *Participant* on the dependent variable of interest.

In Figure 49 (top right) we see the main effect of Appearance ( $F(6, 372) = 25.74, p < 0.001$ ) on the Response Time (RT) dependent variable. Fabrics with the appearance of Burlap, Polyester and Silk were recognized significantly faster than the other materials (post-hoc LSD tests confirmed these results). We found no main effect of either distance or dynamics on response time. As we included synthetic fabrics in which the dynamics and appearance were consistent, we tested these results separately from the mis-matched stimuli. In Figure 49 (top left) we see the main effect of independent variable *Fabric* on dependent variable %Correct ( $F(6, 367) = 88.87, p \approx 0$ ). In this case, Burlap,

Polyester and Silk were correctly matched with their real counterpart significantly more often than the other fabrics. Again, we found no effects of distance or dynamics. These results together already suggest that appearance is having a more dominant effect on recognition than dynamics. It also indicates that some fabrics are being confused with each other quite often, with recognition of Canvas well below chance level.

In Figure 49 (bottom) we see the main effects of dynamics and appearance on the real cloth chosen. In this case the dependent variables were %A-Chosen and %D-Chosen, indicating the percentage of times Appearance or Dynamics respectively determined the choice of the corresponding real fabric. We found no effect of distance in either case. For A-Chosen, we found a main effect of Appearance ( $F(6, 372) = 98.04k, p \approx 0$ ), where Burlap, Polyester and Silk drove the choice of the real fabric at above chance levels (Figure 49: bottom left). An effect of Dynamics on A-Chosen was found (not shown,  $F(6, 372) = 19.21, p \approx 0$ ), as when the dynamics of the synthetic stimulus was Burlap, Polyester or Silk, the percentage of times that the Appearance dominated was reduced, in particular for silk. For D-Chosen (Figure 49: bottom right), we can see that the effect of Dynamics on the choice of the real fabric was much lower than for Appearance. However, the dynamics of Silk had a significantly higher influence on recognition than any of the other fabrics ( $F(6, 372) = 22.26, p \approx 0$ ). Appearance also had an effect on D-Chosen ( $F(6, 372) = 43.32, p \approx 0$ ), again with Burlap, Polyester and Silk appearances reducing the number of times that dynamics determined the choice of real cloth. Since the same effect was seen on dynamics over appearance, the results suggest that when either the appearance or the dynamics are very particular, they tend to introduce some confusion on the participants responses even if the other factor is the correct match. This effect is more prominent when the appearance is very distinct, as Figure 49 bottom right shows.

In order to determine whether there was an effect of participant sex on performance, we ran a mixed repeated measures ANOVA with categorical predictor Sex and independent variable Chosen (A or D). The ANOVA was balanced in this case due to the way in which we distributed the stimulus combinations. As expected, we found that Appearance was chosen more often than Dynamics ( $F(1, 61) = 47.69, p \approx 0$ ). However, there was also a marginally significant Chosen\*Sex interaction ( $F(1, 61) = 3.87, p = 0.054$ ), where men chose appearance significantly more often than women, and vice versa for dynamics.

All results are summarized in Figure 50, where we can see more details on when Appearance dominated the responses for the three fabrics: Burlap, Silk and Polyester. There was more confusion between the other materials. We can also better see when Dynamics affected the choices, and the only material where dynamics was in any way

influential was for Silk, where the green line in the figure shows how the response was always silk when the dynamics were silk, and silk was also often picked when the appearance was a different material. For an additional view to the data, responses based on Appearance and Dynamics are shown in Figure 51.

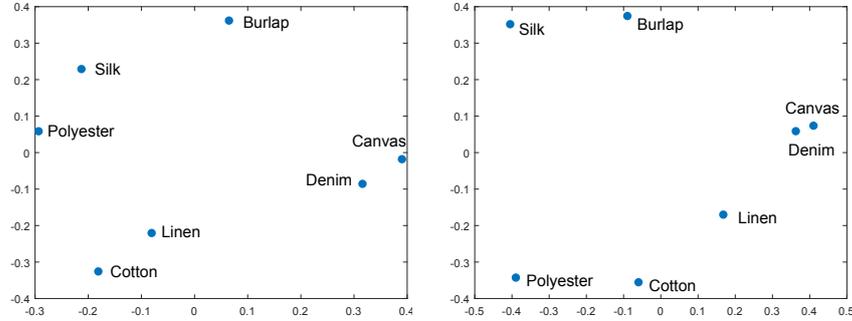


Figure 52: *Left*: perceptual distance of the fabrics by using the dissimilarity matrix including all kinds of mismatching (appearance wrong; dynamics wrong; appearance and dynamics wrong), using  $\kappa = 1$  in Equation 4.4.1. We can observe how the fabrics are grouped in a way that Burlap is pretty isolated, and groups of two are formed by Silk with Polyester, Cotton with Linen, and Denim with Canvas. The last one is expected, since both are made of cotton with tight weaving patterns that result in very similar dynamics and appearance. Polyester and Silk have very characteristic dynamics and appearance, being lightest fabrics. Thus, they are at a very far distance from Denim and Canvas, the stiffest ones. Burlap lies at a similar distance from every group, since it is very stiff but not as heavy as Denim or Canvas. Its transparency can make its position to be closer to Silk. *Right*: perceptual map just using the total mismatched weights (appearance and dynamics wrong) to set the distances ( $\kappa = 0$ ). This way we can better see how fabrics were identified completely wrong. Canvas and Denim are the most confused, a fact that was observed in reality due to their similar look and dynamic behavior. On the other hand, Burlap and Silk also get closer, which could mean there is indeed an effect of transparency in these cases. The increased distance between Silk and Polyester also reveals that their very characteristic appearance and movements prevent users from confusing them whenever their attributes are set right.

In addition, we use multidimensional scaling to better visualize the similarity of the set of fabrics studied. To do so, we establish a  $7 \times 7$  distance matrix  $D$  using a metric such that:

$$D_{i,j} = f_{i,i,j} + f_{j,j,i} + \kappa (f_{i,j,i} + f_{j,i,i} + f_{i,j,j} + f_{j,i,j}) \quad (14)$$

where  $f_{i,j,k}$  is the frequency the fabric  $k$  was matched as response given a fabric with the appearance  $i$  and the dynamics  $j$ , being  $i, j, k \in$

{Burlap, Canvas, Denim, Linen, Cotton, Polyester, Silk}. The term  $\kappa$  weights the effect of the perfect mismatches vs the mismatches where either appearance or dynamics are correct (e.g.: silk appearance / silk dynamics / burlap response vs silk appearance / burlap dynamics / burlap response). Results are shown in Figure 52.

**Discussion:** We may now attempt to answer the questions we asked above. (i) *How effective are our cloth simulations at emulating the appearance and dynamics of the real stimuli?* It is difficult to provide a conclusive answer to this question at this point. On the one hand, the accuracy results for some of the consistent simulations was very high for Burlap, Polyester and Silk; yet for the other four fabrics there was a lot more confusion. (ii) *Are dynamics and appearance equally important when comparing photo-realistic cloth animations with videos of the real fabrics?* The results for the mismatched simulations indicate that there were far more accurate matches based on the appearance of the fabrics, than on the dynamics. In the latter case, only Silk and Canvas were accurately matched at above chance levels (i.e.,  $> 1/7 \approx 14\%$ ). We discuss this further in the conclusions.

#### 4.4.2 Experiment Two: Identifying Mismatches

In Experiment Two we further explore the sensitivity of participants to mismatches between the appearance and dynamics of photo-realistic cloth animations.

**Method:** The screens and controlled settings are as in Experiment One, and again participants first familiarize themselves with the seven real videos. Thereafter, one of the real videos is shown on the left screen while two CG videos are shown side-by-side on the right screen (Figure 53). One CG video is the corresponding replica of the real video shown, with matching appearance and dynamics, whereas the other is mismatched with either the appearance or the dynamics from a different cloth. Order is randomized throughout. This leads to 252 combinations in total: 7 fabrics  $\times$  12 mismatched options (6 each for appearance and dynamics)  $\times$  3 viewing distances. The participant is asked which of the two simulated cloths on the right is most similar to the ground-truth cloth video shown on the left. There is no time limit, and the participant is allowed to replay the videos as often as necessary. As before, each participant only watches a subset of the stimuli, where each stimulus pair seen by 45 different people, and each person sees 180 different samples of the total set of 252. This experiment lasted between 50 and 70 minutes, including a 5-min break.

**Results:** We cross-tabulated all participant data in a multi-way frequency table, with frequency counts (i.e., %correct responses) for each combination of: *Distance* (close, medium, far), *Fabric* (Burlap – Silk), *WrongFabric* (Burlap – Silk) and *Mismatch Type* (Appearance, Dynam-



Figure 53: Two screen layout for Experiment Two. On the left, the current reference real fabric is shown. On the right, the CG replicas being evaluated are displayed synchronously. Here, both dynamics are the same, but only one appearance matches.

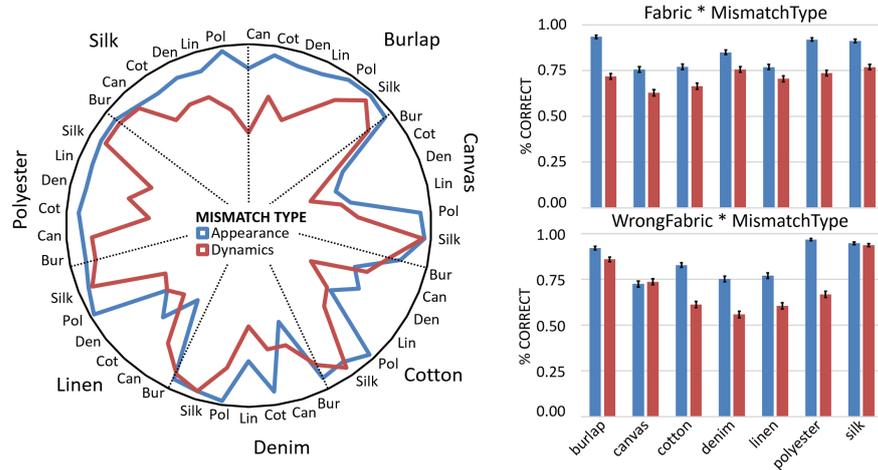


Figure 54: Experiment Two results. Left: All results summarized as a radar graph and collapsed over distance. The outermost labels on the perimeter indicate the “correct” fabric, while the innermost ones show the mismatched one. Right: Interaction effects between the mismatch type (Appearance or Dynamics incorrect) and the correct or mismatched fabric type.

ics). Log-Linear Analysis shows that the best fit model is (Fabric, WrongFabric, MismatchType), with a likelihood ratio of  $\chi^2(196) = 68.89, p = 1.00$ . Therefore, these three factors interact to influence the choice of the correct synthetic cloth, but again distance appears to have no effect.

As before, we report ANOVA results (univariate tests of significance) using Type 1 sums of squares for the unbalanced scores. We tested for main and interaction effects of fixed independent variables *Distance*, *Fabric*, *WrongFabric* and *MismatchType* (Fixed) and random independent variable *Participant*, on the dependent variable *%Correct*. (We also tested for dependent variable Time but found no significant effects). We found main effects of Fabric ( $F(6, 372) = 22.57, p \approx 0$ ), WrongFabric ( $F(6, 372) = 121.96, p \approx 0$ ) and MismatchType ( $F(1, 62) = 77.87, p \approx 0$ ). Performance was above 66% correct for all fabrics, but was significantly higher for Burlap, Denim, Polyester and Silk. When Silk was the WrongFabric, performance was highest, followed by Burlap and then Polyester. Appearance mismatches were more

easily identified than Dynamics (84% vs 71%), again indicating that appearance dominates.

The results are shown in Figure 54 (Left). We found no main or interaction effects for Distance, as before. These results can be further understood by exploring the two significant interaction effects: Fabric\*MismatchType ( $F(6, 372) = 4.94, p < 0.000001$ ) and Wrong-Fabric\*MismatchType ( $F(6, 372) = 29.72, p \approx 0$ ) as shown in Figure 54 (Right). In particular, we can see that both mismatching appearance and dynamics were equally salient for Silk and Canvas, and close for Burlap. Instead, Polyester is where appearance dominates most over dynamics, followed by Cotton, Denim, Linen and Burlap. In this experiment, there was no effect of participant sex.

#### 4.4.3 Experiment Three: Perceived Cloth Features

From Experiments One and Two, we may conclude that the Appearance of a fabric is consistently more salient than its Dynamics. Furthermore, this effect appears to be independent of viewing distance. In order to explore this finding further, we picked the most poorly matched fabrics for appearance and dynamics respectively attempting to isolate the effect of each factor. Log-Linear Analysis found an effect of appearance and dynamics independently on the responses, with a perfect fit to the observed data in both cases  $\chi^2(0) = 0, p = 1$  (summary can be seen in Figure 55). In addition, we tested the saliency of several key dynamics and appearance features of the fabrics, which is the bulk of the Experiment 3, detailed in the following.

**Method:** In order to explore the features related to the Dynamics factor, an identical and neutral appearance is needed for every cloth, as neutral as possible. We therefore rendered all seven cloth dynamics with the same gray diffuse material, free of texture, and displayed them at the medium viewing distance. For the case of Appearance factor, we displayed cropped images showing the fabrics with the same folds under homogeneous lighting. We chose cropped images in this case to reduce the impact of any intrinsic information about the dynamics properties of the cloth, thus allowing participants to focus on the appearance. We chose a 3AFC design, where three samples are displayed with each other on one screen. We distributed these combinations of three fabrics in such a way that each one appears 6 times on the left, center and right of the screen, leading to 42 possible combinations. This design allowed all of the fabrics to be compared with all of the others in the most efficient manner. Some examples can be seen in Figure 56.

Fifteen participants volunteered for this experiment (6F/9M, aged 18-35), and the entire experiment took approximately 30 minutes. Each participant viewed both conditions in counterbalanced order, and in each case all 42 combinations were shown in randomized or-

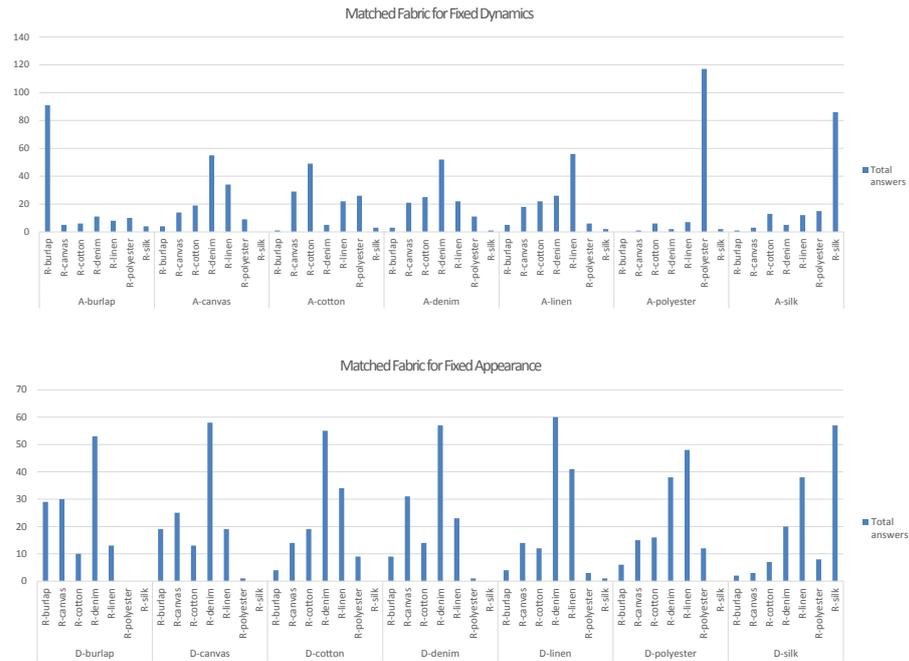


Figure 55: Frequency of chosen fabrics when dynamics (top) and appearance (bottom) are fixed. Attempting to study the effect in isolation of each factor, we chose the most poorly matched fabric in each case: cotton dynamics and canvas appearance (see Figure 49). We can see how appearance is very well recognized in many cases, excelling for Burlap, Polyester and Silk, getting more confused for fabrics with similar middle-roughness look, weaving patterns and materials (cotton). Instead, canvas appearance is mostly identified as denim except for very particular dynamics like Silk or Linen, in agreement with the multidimensional scaling shown in Figure 52

der. In the Dynamics condition they are asked to indicate which of the fabrics is most *Stiff*, and which one is most *Heavy* (Figure 56, left). For the Appearance condition, they are asked which fabric looks most *Glossy* and which one looks most *Rough* (see Figure 56, right).

**Results:** We performed repeated measures single factor ANOVAs on each set of ratings and post-hoc analysis using Newman-Keuls pair-wise comparisons of means. The results are shown in Figure 57. For dynamics attribute Stiffness ( $F(6, 84) = 255.51, p \approx 0$ ), the ranking was as follows (from most to least stiff): Burlap > Canvas > Linen > (Canvas = Denim) > (Silk = Polyester). For Heaviness ( $F(6, 84) = 2.66, p < 0.05$ ), all choices were at chance level (i.e., approx 33%) except for cloth, which was only chosen as most heavy 14% of the time. For appearance Glossiness ( $F(6, 84) = 21.40, p \approx 0$ ), polyester was chosen as most glossy 90% of the time, with burlap and linen only being chosen about 10% of the time each; the others were in between these extremes but in overlapping homogeneous groups

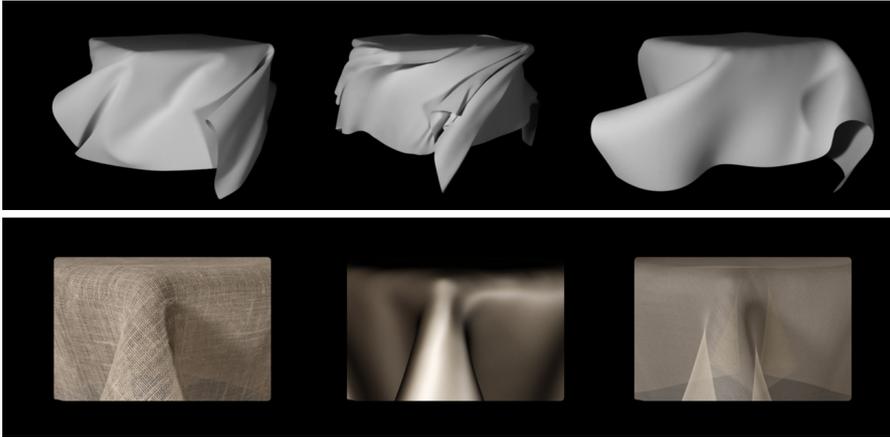


Figure 56: One screen layouts for Experiment Three. *Left*: three videos of neutral appearance and different dynamics are played simultaneously, with the swivel stool rotating in the same direction; *Right*: three cropped images of the fabrics draping the stool with the same folds but different appearances are shown.

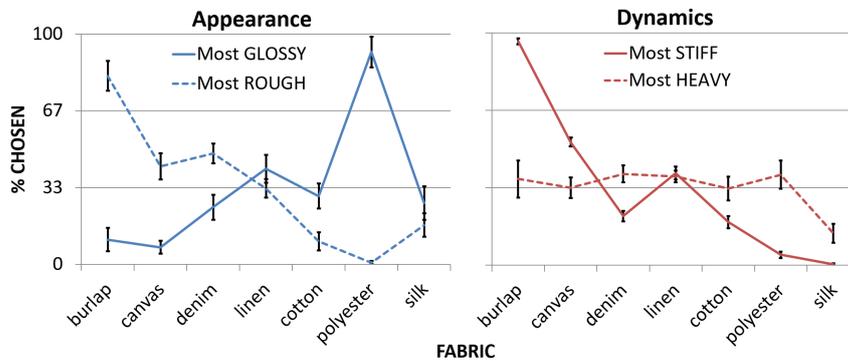


Figure 57: Experiment 3 results, showing the mean percentage of times each fabric was chosen as being the most Glossy, Rough, Stiff or Heavy, along with standard error bars.

based on the post-hoc tests. The ranking for appearance Roughness ( $F(6, 84) = 29.95, p \approx 0$ ) was as follows: Burlap > (Denim = Canvas = Linen) > Polyester, Cotton and Silk. Again, we found no effect of participant sex on the results.

We can see from these results that participants appear to discriminate better between the appearance attributes of the cloth, with more significant differences between the choices of glossiness and roughness for all fabrics. The stiffness of a fabric appears also to be a salient feature, but participants appeared to be unable to distinguish between the weights of the different fabrics, which is surprising, as this is an intrinsic property of its dynamics. These results further emphasize the findings of the previous experiments that the Appearance of a cloth is the most robust indicator of its identity.

#### 4.5 CONCLUSIONS

In this chapter, we have presented the results of three perceptual experiments where we explored the interactions of the Appearance and Dynamics of seven common woven fabrics and also explored the saliency of some of their key attributes (stiffness and heaviness for dynamics; glossiness and roughness for appearance). We demonstrate how appearance dominates over dynamics, except in the case where where dynamics is very characteristic, i.e., for silk. We also found that these effects are robust across different viewing distances. We found that there is a significant tendency for men to rely on appearance more than women do when identifying fabrics. Perhaps this is because women often have more experience with wearing a larger variety of fabrics and flowing garments and hence may be more familiar with their dynamics.

We have explored only a small range of different fabrics in this work to keep the problem tractable, and a wider choice with more varied dynamics would be needed to draw any definitive conclusions. In fact, it would be interesting to further explore the no effect of distance on the results, which was unexpected beforehand. This should be addressed with a wider sampling of the relationship between the weaving pattern sizes, their projection onto the image plane, and the image resolution itself. Our results indicate that, in general, appearance dominates over dynamics in our study. We speculate that this may be due to the fact that we are more accustomed to understanding appearance information; when looking at different fabrics, we always see their appearance, whereas watching them in motion occurs more rarely. This may not be the same in other appearance vs. dynamics studies involving different elements other than fabrics. As in all studies of similar nature, our conclusions are only strictly valid for our chosen stimuli. Our particular choice of fabrics is intended to provide an ample enough range of appearances and dynamics, while keeping the experiments tractable; it could be that using different visual stimuli, such as non-shiny, opaque fabrics, or adjusting the thickness to be similar across fabrics, the result would vary somehow. From the results we can observe that when either the dynamics or the appearance are particularly dominant, the other factor has perceptually less contribution. Without the cue of appearance, we have the intuition that fabrics with characteristic movements (Polyester, Silk, Burlap or Canvas) would get high accuracy ratios. Among fabrics with similar behavior, like cotton, denim or canvas for instance, the threshold to distinguish between them would remain low, since the materials chosen for those fabrics are very neutral.

We should also be wary of drawing any general conclusions regarding perception of appearance and dynamics in general, as it is plausible that some of our results could be affected by the fact that

rendering techniques are far more advanced than dynamics simulation, having been studied for longer. It is therefore possible that appearance has been synthesized with more fidelity to the real fabric than the dynamics have been. Furthermore, performing a similar study with animated characters wearing clothes made from these fabrics would allow us to confirm our findings in more ecologically valid and familiar scenarios.

It would be interesting to consider some other factors that may have an effect on the perception of moving cloth (e.g., different illumination conditions such as environment lighting), or to explore more deeply the influence of the most important factors of cloth simulation considered here (e.g. BRDF and spatial frequency of the textures in the case of the appearance, dynamics parameters in the case of motion synthesis).



## AN APPEARANCE MODEL FOR TEXTILE FIBERS

Accurately modeling how light interacts with cloth is challenging, due to the volumetric nature of cloth appearance and its multiscale structure, where microstructures play a major role in the overall appearance at higher scales. Recently, significant effort has been put on developing better microscopic models for cloth structure, which have allowed rendering fabrics with unprecedented fidelity. However, these highly-detailed representations still make severe simplifications on the scattering by individual fibers forming the cloth, ignoring the impact of fibers' shape, and avoiding to establish connections between the fibers' appearance and their optical and fabrication parameters.

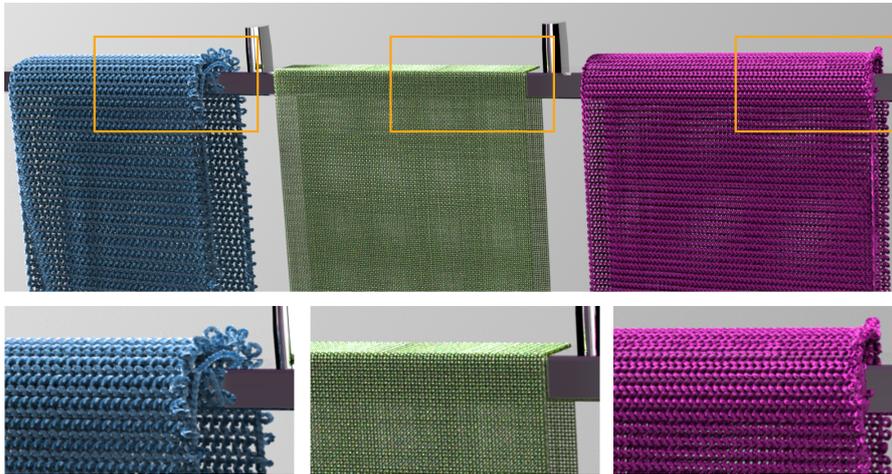


Figure 58: Volumetric rendering of three fabric samples using our highly-detailed physically-based BCSDFs models. Our model is based on the optical and structural parameters particular of each cloth fiber. From left to right, a 2-ply knitted garter (7 twist/cm, 120 fibers/yarn) of cotton fibers, a single-ply woven satin made of silk (1 twist/cm, 30 fibers/yarn) and a 2-ply knitted garter (4 twist/cm, 160 fibers/yarn) made of polyester fibers. Silk and cotton pieces are colored with reactive dyes, and a disperse dye is used for polyester.

In this work we put our focus in the scattering of individual cloth fibers; we introduce a physically-based scattering model for fibers based on their low-level optical and geometric properties, relying on the extensive textile literature for accurate data. We demonstrate that scattering from cloth fibers exhibits much more complexity than current fiber models, showing important differences between cloth type, even in averaged conditions due to longer views. Our model can be

plugged in any framework for cloth rendering, matches scattering measurements from real yarns, and is based on actual parameters used in the textile industry, allowing predictive bottom-up definition of cloth appearance.

## 5.1 INTRODUCTION

Rendering realistic fabrics is an active research field in computer graphics, with many applications in other areas like entertainment or textile design. Accurately reproducing the appearance of cloth remains challenging, due to the micro-structures found at fiber level, and the complex light scattering patterns exhibited at such scales. This affects the overall look of cloth both in close-ups, and at longer viewing distances, with effects like anisotropic specular highlights, or anisotropic multiple scattering.

Given this intrinsic complexity, and the effect of microscale appearance at coarser scales, it is necessary to accurately represent appearance at very small scales, to capture the subtleties and rich optical phenomena of cloth fibers. Recent approaches have made significant advances in this direction, either by modeling the arrangement of the fibers [199], or capturing the structure of small pieces of cloth using macro photographs [203] or Computed Tomography scanners (CT) [247, 131, 251]. However, these works focus on reconstructing the geometry of the fibers, while adopting several convenient simplifications to simulate the scattering of light, such as using parametric, simplified fiber scattering models. While these models improve over general volumetric phase functions such as microflakes when rendering cloth [131], they 1) still rely on very simplified models of fibers, and 2) lack the actual fabrication parameters used when physically manufacturing the fiber. These two points limit the applicability of current models for predictive rendering, based on optimization-based approaches.

In this work we introduce a novel appearance model for cloth, focusing on high-quality scattering functions, which takes into account the optical and structural properties of real-world cloth fibers, from a ray-optics perspective. We leverage the wealth of measured, real-world data available from the textile research community, and build digital replicas of different types of fibers (polyester, wool, silk and cotton). We then rely on brute force simulations of light scattering of such fibers to obtain highly detailed tabulated *Bidirectional Curve Scattering Distribution Functions (BCSDF)*, without the most common simplifications assumed in current BCSDF models [255, 130] (e.g. circular or elliptical cross sections), which are unable to capture the rich visual features of real fibers. Moreover, since our model is based on low-level real-world fiber properties, it allows us to specify actual fabrication parameters used by fiber makers; this is very relevant in fields

such as textile or cloth design, where manual or optimization-based parameters fitting might not match the properties of actual existing fibers. To our knowledge, ours is the first scattering model based on such physical and optical measurements, including fibers' index of refraction, surface roughness and dye concentration.

We compare our model against state-of-the-art BCSDFs at different scales: at highly detailed close-up views where the scattering of individual fibers is visible, as well as at longer shots where the scattering of the individual fibers is averaged, and the macroscopic appearance of cloth becomes clear. Our BCSDFs can be directly used for rendering both volumetric-based (see e.g. Figure 44) and explicit fiber-based representations of cloth; moreover, since it allows for a bottom-up definition of cloth appearance, it bridges the gap between actual fiber fabrication parameters and simulated cloth appearance in a predictive manner.

## 5.2 RELATED WORK

We focus on detailed volumetric or explicit representations of cloth appearance, although other approaches using planar representations of cloth exist, based either in bidirectional texture functions (BTFs) [46, 194] or in parametric local scattering models [113, 190]. For more details we refer to the original papers, or the survey by Schröder and colleagues [201].

**VOLUMETRIC CLOTH MODELING** Two different approaches for volumetric cloth representation have been widely adopted in the field: 1) assuming cloth as an heterogeneous anisotropic volume [117] and posing the rendering problem as volumetric light transport, and 2) modeling the geometry of each yarn or fiber explicitly, similar to recent hair rendering approaches.

Zhao and colleagues [247] followed the former approach, leveraging Micro-CT scanners of cloth in order to obtain highly detailed volumes at micron resolution of small pieces of fabric. Micro-CTs have also been used as building blocks for larger garments with repeated patterns [249], and to take advantage of the repeated structure of cloth to precompute inter-blocks light transport to accelerate rendering [250]. While CT-based models allow for high-quality renders of cloth, they rely in complex and expensive capturing setups, and manual intervention or optimization [131] with respect to a target cloth appearance is needed to define the optical parameters of cloth garments. Instead of relying on captured data, Schröder and colleagues [199] procedurally generate cloth garments as a collection of individual yarns. These yarns are then transformed into a volumetric representation using Gaussian distributions of fiber orientation and density. Although such model can produce realistic results, it does not hold for

close-up views where a huge number of voxels needs to be computed. Similarly, Lopez-Moreno et al. [142] build a volumetric representation of cloth on the GPU. They adapt the classic *lumislice* algorithm [238] to interactively create high-quality volumetric representations of cloth at fiber scale.

As opposed to volumetric approaches, Schröder and colleagues [203] relied on a fully geometrical representation of individual yarns, each with a different scattering function. They extended their previous work [199] entirely avoiding the volumetric representation of cloth. Their fully procedural fiber generation method even simulates the small fibers protruding from the yarn (hairiness) and allows predictive reverse engineering of real cloth. Zhao et al. [251] automated the fitting of such yarn procedural models from physical measurements acquired using micro CT imaging. Recently, Wu and Yuksel [237] demonstrated real-time rendering of cloth modeled using explicit yarns.

Despite their differences, both approaches (volumetric and explicit geometry) are able to deliver comparable high-quality results [131]. Our work is orthogonal to the chosen cloth representation, and our realistic fiber-based scattering functions can be applied to both.

Finally, although not explicitly for cloth, we note that hybrid approaches considering volumetric and explicit representation of packed discrete media have been proposed, taking the best of each world by keeping the high-frequency details of explicit geometry in the first bounces, while accelerating rendering for higher-order scattering using a volumetric approximation [156, 160].

**SCATTERING MODELS** Most previous approaches have assumed general scattering models for fibers, ranging from microflakes [117, 105] for volumetric models, to explicit fiber scattering models similar to hair rendering e.g. [127, 148, 255, 242].

While both extremes can produce visually realistic results, and many approaches have successfully used microflakes for rendering cloth [247, 250, 142], it has been shown that microflakes cannot match the appearance of real-world cloth, while fiber scattering models are more suitable for this task [131]. In this context, Schröder and colleagues [199, 203] use the parametric BCSDf model proposed by Zinke and Weber [255]. Khungurn et al. [131], on the other hand, proposed a simple fiber-based model suitable for both rendering and appearance capture of real garments. The authors highlight the importance of the fiber-specific scattering model to achieve good results, and choose a simplified BCSDf that fits well within their optimization framework. However, their fiber scattering model has certain limitations from a physical point of view, since it assumes only direct reflection and transmission, ignoring longer scattering orders within the fiber, and allowing colored reflectance which is not predicted in

Fresnel’s equations. The recent paper of Zhao et al. [251] relies on the same scattering model.

A common limitation shared by previous physically-based fiber scattering models [148, 255, 53, 242] and artist oriented models [187, 39] is the lack of generality: they assume fibers are cylinders with a cross section that can be elliptical up to some degree. Ogaki et al. [169] proposed a similar approach to ours for fur modeling, where explicit geometry was used to obtain hair scattering functions from fibers with arbitrary geometry. Their work, however, does not account for realistic fiber geometries based on Scanning Electron Microscope (SEM), and does not consider real-world measurements of optical parameters. The recent model by Khungurn and Marschner [130] is able to handle arbitrary eccentricity for the first time. However, fabric fibers lie far from this assumption (see Figure 59), leading to scattering patterns much more anisotropic and difficult to represent through current models. In addition, these models do not provide any connection between fabrication parameters and optical properties of the fibers; as we will show, these fabrication parameters have a major impact on the scattered field.

**TEXTILE RESEARCH** Beyond computer graphics, several works in textile research have used simulation to predict the appearance of cloth. Most of these works use light simulation in yarns or fibers focusing on quality assessment of specific features such as luster [212, 15], and on determining the optical properties of cloth by means of inverse rendering [99, 184]. Closer to our work, Yamada et al. [239] compare through simulations the scattering functions of synthetic fibers with circular, triangular or rectangular cross sections. A similar approach is followed by Liu et al. [139] and Aslan et al. [16] for synthetic and cotton fibers respectively. None of these works use the computed scattering functions for rendering, nor propose a full bottom-up approach for defining cloth appearance. Finally, similar to other related works in graphics [247, 131], Grasso et al. [95] rely on simulation to study the effect of textile properties in the macroscopic appearance of cloth, while alternative approaches evaluate radiative transfer in fibers and textiles [241, 146, 240], and compare simulations and measurements of the propagation of polarized light through textile materials [177]. These approaches focus either on a few properties individually, or on particular types of fiber, and do not attempt to provide a general model suitable for computer graphics. Instead, our model is capable of reproducing the appearance of cloth, based on fiber fabrication and measured optical parameters, making it well suited for predictive applications.

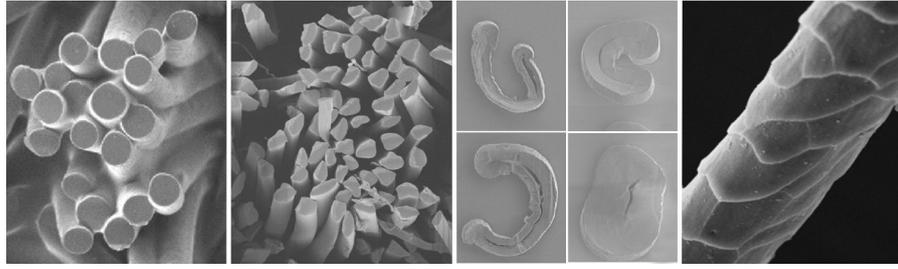


Figure 59: SEM images of cross sections and shapes for real polyester, silk, cotton and wool fibers. Images from [55], [207], [165] and [128], respectively.

### 5.3 LIGHT SCATTERING FROM TEXTILE FIBERS

Most cloth fibers are made of an absorbing dielectric medium. Although their shape has been traditionally considered as a cylinder extruded from circular or elliptical cross sections, their actual cross sections greatly differ from these simplified shapes, depending on the type of fiber, as well as its chemical or mechanical treatment. Moreover, the assumption of a smooth cylindrical shape does not hold for some natural fibers such as wool, which presents overlapping, tilted scales on its surface (similar to human hair [148]). This can be seen in Figure 59, showing SEM samples of real textile fibers.

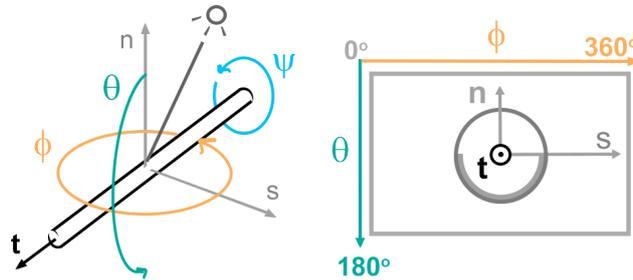


Figure 60: *Left*: spherical coordinates and local frame defined for the fiber, aligned along the  $t$  axis. *Right*: illustration of coordinate system and shape of lobes observed in the scattering plots; circle in the middle is the reflection cone under illumination angle  $\theta_i = -45^\circ$ .

This leads to visually important reflectance features such as self occlusions or caustics, which cannot be simulated with current BCSDf [255]. Figure 8 shows several examples of 2D slices of the four-dimensional BCSDf for different incoming directions  $\omega_i$  (the coordinate system is described in Figure 60), computed for a set of fibers of polyester, silk, cotton and wool, each with different fiber cross sections and optical properties. While the BCSDf of polyester, a synthetic radially-symmetric fiber, is rotation-invariant and exhibits uniform high-frequency lobes due to sharp reflection and caustics on the cone of reflection, the rest of the natural fibers yield heterogeneous and highly anisotropic scattering profiles, dependent on the cross section and the incident

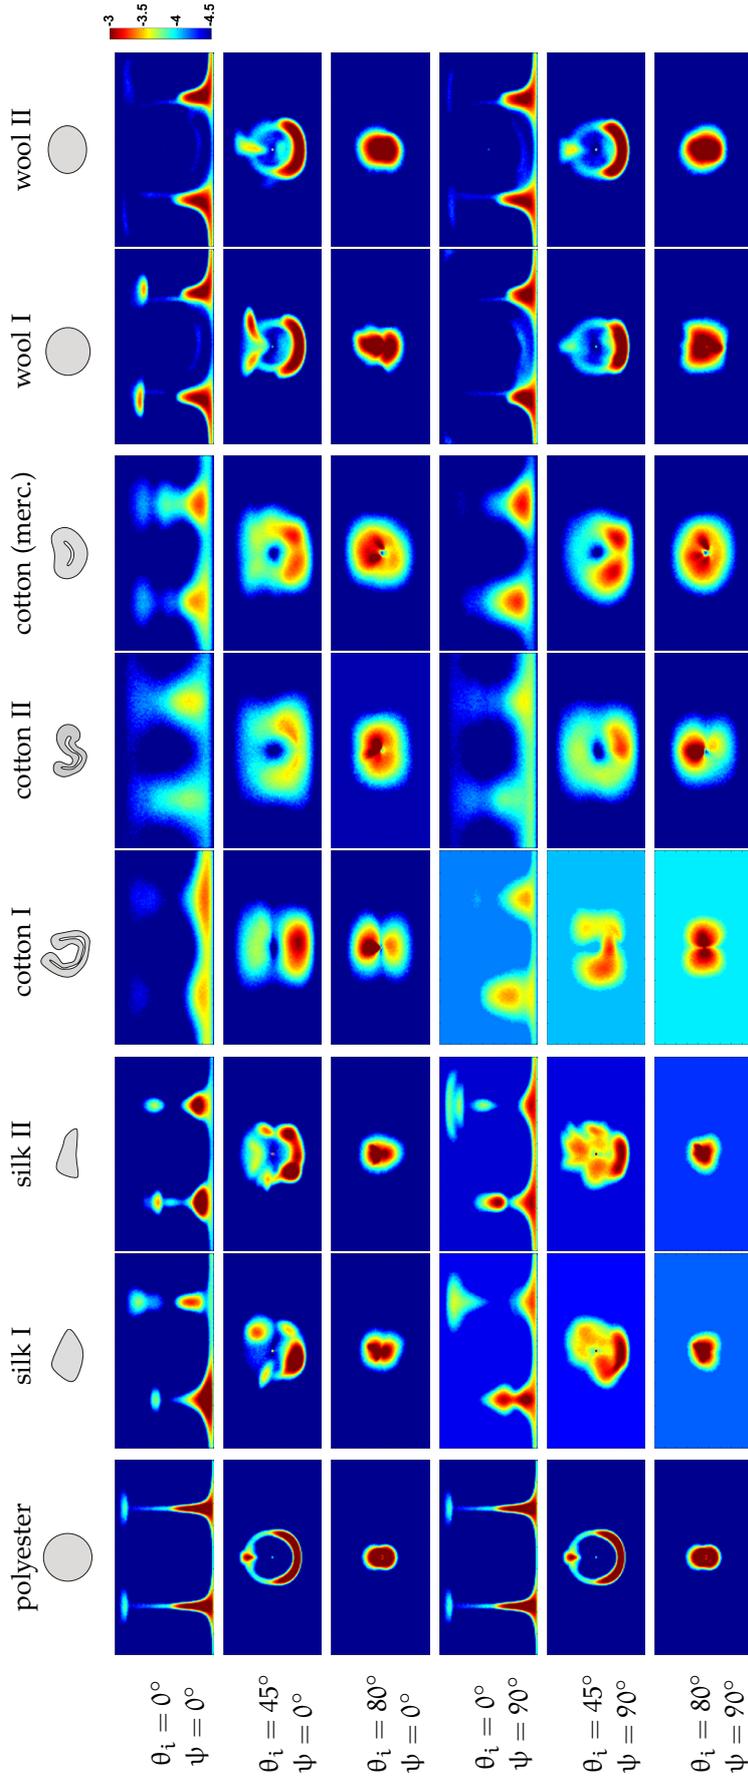


Table 8: Outgoing field of the BCSDf for a set of cloth fibers under varying illumination setups, with false color depicting radiance in logarithmic space. From left to right: polyester, silk, cotton and wool, each with several different cross-sections (top), equalized in diameter for visualization purposes (real relative sizes can be found in Table 9). Each fiber is softly dyed ( $\zeta = 0.1\%$ ). Note that the cross section of mercerised cotton, which is a typical treatment of this fiber to gain luster, is very different from untreated cotton. From top to bottom, the incident light varies from perpendicular direction to near grazing illumination angle  $\theta_i = 0^\circ$ ,  $-45^\circ$ ,  $-80^\circ$  respectively, and  $\phi_i = 0^\circ$ , and the fiber is rotated along its tangent direction by  $\psi = 0^\circ$  and  $90^\circ$ ; see Figure 60 for the coordinates in the BCSDf.

light direction  $\omega_i$ . Silk for instance exhibits sharp, highly anisotropic caustics due to its flat irregular shape. Cotton, on the other hand, presents wider lobes due to its higher surface roughness, its inside hole, and the multiple self-interreflections within the fiber boundaries, specially in the case of non-mercerized fibers. Finally, the reflectance of wool is more similar to perfect cylindrical fibers (e.g. polyester) due to the low eccentricity of its cross section, but presents a set of high-frequency anisotropies due to the irregular cuticle variation along the longitudinal axis.

In addition to structural properties, light scattering in a cloth fiber also depends on its optical properties, such as its index of refraction or the absorption, which in turn is highly dependent on the concentration of the dye used to color the fiber. Most previous work have attempted to model these parameters using a top-down approach, either by directly defining the fiber color [117, 199], or matching a specific appearance through optimization [131]. While this leads to visually plausible results, these approaches may lead to physically invalid parameters, such as absorption values out of the range of actual fabrics, or artifacts like colored specular reflections. Instead, we opt for a bottom-up approach, defining fiber appearance from real-world low-level structural and chemical data, including the specific dyes used with different types of cloth.

Given this complexity, it is not feasible to derive an analytical BCSDF to simulate fiber appearance without major simplifications; instead, we follow previous works on modeling complex reflectance and appearance [189, 156, 160], and build a tabulated BCSDF from a physically accurate fiber model, which takes into account precise structural and optical properties. This is described in detail in the following subsections, and compared against state-of-the-art fiber scattering models. Later we will show how to use these complex BCSDFs for rendering volumetric models of cloth.

**CROSS SECTION AND LONGITUDINAL STRUCTURE** To take into account how the different cross sections affect the optical properties of cloth, we analyzed SEM samples of real fibers, and manually modeled a representative set of different cross sections for each type. While some fibers such as polyester fit well into the assumption of cylindrical cross section, natural fibers present significant variability in their shapes, as we have seen. In addition, we also take into account the absolute size of each type of fiber; for instance, wool can be up to eight times larger in diameter than silk (see Table 9). Of all the fibers analyzed, wool presents an additional challenge: its surface is made up of overlapping cuticles (not present in other fibers, see Figure 59), which significantly affect light scattering. We approximate this effect by modeling shape variations along the longitudinal axis.

Type	Diameter $\mu\text{m}$	Density $\text{g}/\text{cm}^3$	IOR $\eta_{\parallel}, \eta_{\perp}$	$R_a$ (l) nm	$\beta$ degrees
Polyester	10	1.39	1.73, 1.54	2.33-5 (30)	2.7-3.5
Silk	5-10	1.34	1.591, 1.538	8-9 (30)	6-7
Cotton	17-20	1.52	1.578, 1.532	12.5-15.8 (50)	14-17
Wool	24-40	1.31	1.553, 1.542	6 (50)	5

Table 9: Measurements of the physical and optical properties of four of the most common fabrics, acquired from the textile literature: fiber diameters from Trotman [219]; fiber densities (in a standard atmosphere of 65% relative humidity, 20°C) and index of refraction from Hearle and Morton [104]; surface roughness  $R_a$  and  $l$  from [90, 137].

### 5.3.1 Fiber properties

**SURFACE ROUGHNESS** The dielectric boundaries of the fibers present roughness at nanoscopic scale, which prevents us from using directly the Fresnel equations (see index of refraction for all fibers in Table 9; note that the wavelength dependency on  $\eta$  is neglectable [104]). As in recent work [58, 164], we model surface roughness statistically, following a microfacet-based approach, where the Normal Distribution Function (NDF) is modeled as a Beckmann distribution with average orientation  $\beta$ , following a V-cavity model [44, 170]. The average normal orientation is thus modeled as  $\beta = \arctan(R_a/l)$ , where  $R_a$  is the average peak-to-valley height, and  $l$  is the profile length (both in nm). Note that this models roughness at the nanometer scale only; coarser features are captured by the geometry of the fibers and our specific cross sections.

**ABSORPTION** Absorption inside the fiber plays a crucial role on its appearance [104]. The main source of absorption is the dye used to color the fiber. Thus, we compute the absorption coefficient  $\mu_a$  [ $\text{m}^{-1}$ ] based on the amount of dye and its particular absorption as  $\mu_a = \kappa \varepsilon$ , where  $\kappa$  is the dye concentration [ $\text{g l}^{-1}$ ], and  $\varepsilon$  is the extinction per gram [ $\text{l g}^{-1} \text{m}^{-1}$ ]. The latter is given by  $\varepsilon = \varepsilon_m w_m^{-1}$ , being  $\varepsilon_m$  the molar extinction coefficient in [ $\text{l mol}^{-1} \text{m}^{-1}$ ], and  $w_m$  the molar weight of the dye [ $\text{g mol}^{-1}$ ].

We model the dye concentration  $\kappa$  as a function of the fibers' density  $\rho$  [ $\text{g l}^{-1}$ ] (see Table 9), and the *Depth of Shade* (DoS)  $\zeta$ , as  $\kappa = \rho\zeta$ . The DoS  $\zeta$  is a quantity used in industry for controlling the saturation of dyed cloth, which is the ratio of grams of dye to grams of fiber (ranging from 0.1% for pale shades to 4% for deep shades [166]). Note that at such small dye concentrations the impact of the dye on other optical properties of the fiber is minimal [104]. We choose two of the most common dyes suitable for a very wide range of commercial fabrics (reactive and disperse). Depending on their chemical and

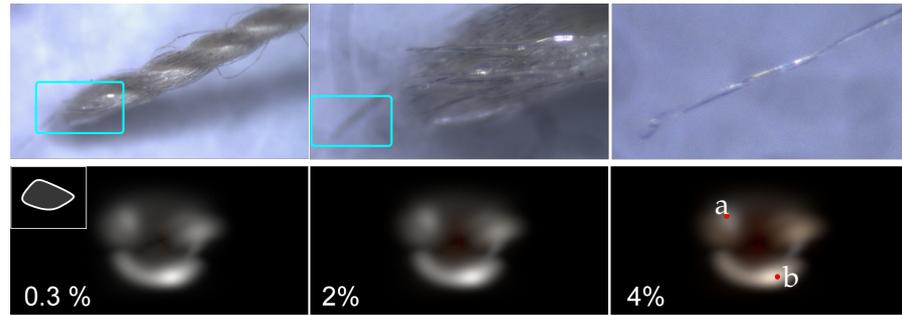


Figure 61: Light absorption in fibers and threads. First row shows images under the microscope of a real silk thread of 3 plies, each composed by 90 fibers. The rightmost image shows how a single fiber has very low absorption, but the cumulative effect of each of the 90 fibers per ply give the thread its overall yellowish color. Second row shows the effect in our model of the DoS  $\zeta$  used to dye the fibers, in particular, the scattering of a silk fiber (see cross section on the inset on the left) lit from  $\omega_i = \langle 0^\circ, 45^\circ \rangle$  with an orange dye with a different  $\zeta$  parameter: from left to right, it varies from very pale to very deep DoS ( $\mu_a = 5.08, 40.46$  and  $84.68 \text{ mm}^{-1}$ ). While the reflection cone remains uncolored (specular reflection, (a) in the rightmost figure), it can be seen how the colored lobes increase in saturation when the amount of dye increases (b). Since the silk fibers are thin, note also that the refraction lobe remains softly colored even for deeply dyed fibers (right).

physical properties, some dyes are more suitable to be used for different kinds of fabrics: reactive dyes, which are mainly used for cotton, silk and wool, have an extinction coefficient  $\epsilon$  ranging from 0.005 to 0.0158 [ $\text{l mg}^{-1} \text{ cm}^{-1}$ ]; disperse dyes, often used for polyester and other synthetic fabrics, have much higher extinction values, ranging from 0.045 to 0.246 [174, 102]. Figure 61 shows the effect of the DoS on a fiber of silk; it does not have a large effect on the color of an individual fiber, however its effect accumulates and becomes visible as it interacts with all the fibers in the cloth.

### 5.3.2 Obtaining the BCSDFs

Since no existing analytical BCSDF model can represent the complex high-frequency, anisotropic features of light transport in realistic fibers, we opt for precomputing light transport under different view and light directions, and tabulate this information to be used directly as a BCSDF.

We place a single straight fiber lit by a beam of light whose width equals the projected area of the cross section in the incoming direction. We consider an infinitesimal longitudinal patch, and adopt a far-field approximation of the BCSDF, a suitable approximation that has been shown to introduce a very low error [255]. This setup allows us to simplify the eight dimensional BCSDF  $f(\mathbf{x}_i, \omega_i, \mathbf{x}_o, \omega_o)$ , with

Type	Average Cost	Resolution ( $\omega_i, \omega_o$ )	Storage BCSDf	Total Storage
Polyester	7 min	$45 \times 1, 90 \times 180$	2.7 MB	5.56 MB
Silk	360 min	$45 \times 180, 90 \times 180$	500.5 MB	1001 MB
Cotton	420 min	$45 \times 180, 90 \times 180$	500.5 MB	1001 MB
Wool	900 min	$90 \times 180, 90 \times 180$	1001 MB	2002 MB

Table 10: Average computation and storage cost for each of the four types of fiber considered in this work. The computation cost is affected both by the complexity of the geometry modeling the fibers, and the number of angular measurements. For rotationally symmetric fibers, such as polyester, very few angular samples are needed to characterize the full BCSDf. Fibers with longitudinal symmetry (silk, cotton) require half the measurements than wool. The total size of the BCSDf (in 32 bits floats for a monochromatic BCSDf) is directly proportional to the total resolution of the BCSDf. Finally, note that in addition to the BCSDf we also store the tabulated CDF for sampling, and the directionally-resolved albedo  $\Lambda(\omega_o)$  and fiber’s projected area  $A(\omega_o)$ .

$\mathbf{x}_i$  and  $\omega_i$  ( $\mathbf{x}_o$  and  $\omega_o$ ) the incoming (outgoing) point and direction, respectively, into a four-dimensional BCSDf  $f(\omega_i, \omega_o)$ .

We discretize the angular domain in  $2^\circ$  intervals, which our tests showed to be enough to capture the sharpest illumination features due to direct reflection and caustics (present in polyester and silk). This results into an angular resolution of  $180 \times 90$  in  $\langle \phi, \theta \rangle$ , for a total of  $(180 \times 90)^2$  combinations of  $\omega_i$  and  $\omega_o$ . For each light direction, we shoot sixteen million photons from the light, and simulate their random walk as they interact with the fiber. These photons are collected in an infinite sphere bounding the fiber, as a function of  $\omega_o = \langle \phi_o, \theta_o \rangle$ , using the irradiance meter sensor in Mitsuba [116]. We compute a two-dimensional slice of our 4D function for each incoming light direction. Our final BCSDf  $f(\omega_i, \omega_o)$  is characterized by the fiber’s parameters (cross section, size, index of refraction, surface roughness, dye used, and depth of shade), and is stored as a 4D table which is accessed directly in rendering time. We use the same approach for computing the directionally resolved fiber projected area  $A(\omega_o)$ , which will be used later for rendering (see Section 5.5). The average precomputation time and storage cost for each type of BCSDf is listed in Table 10. Our BCSDfs are publicly available at online<sup>1</sup>.

## 5.4 ANALYSIS & DISCUSSION

**FAR-FIELD ASSUMPTION** We assume that light interactions occur at the same differential point  $\mathbf{x}$  in the fiber, which allows us to reduce the dimensionality of the BCSDfs to four dimensions. We also leverage this assumption to consider that the fiber is *locally* straight,

<sup>1</sup> <http://giga.cps.unizar.es/~ajarabo/pubs/clothEGSR17/data/>

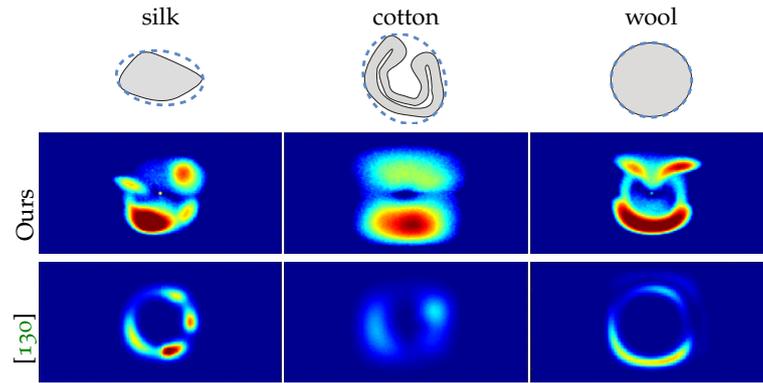


Figure 62: Comparison between our BCSDFs (top) and the results predicted by Khungurn and Marschner’s elliptical fiber model [130] (bottom), for fibers of silk, cotton and wool, illuminated by a beam of light with incoming direction  $\theta_i = 45^\circ$ . We use the best physical fit of the elliptical model, by setting the same or closest fiber parameters - roughness, bounding ellipse and diameter, cuticle slope, etc.- to match the fiber’s real parameters (bounding ellipses are rotated accordingly). Elliptical cross sections are unable to capture the complexity of realistic fiber reflectances.

homogeneous, and has the same cross section in all cases but wool, similar to most previous works on hair rendering [148, 53, 242]. This is a reasonable simplification given the small diameter of the fibers, in general much lower than the fibers’ curvature, and assuming that the fiber’s cross section and twisting varies slowly along the longitudinal axis with respect to the path length.

**COMPARISON TO ELLIPTICAL FIBER MODELS** Most previous fiber reflectance models have assumed circular, or moderately elliptical, cross sections [255, 53, 242, 131]. Recently, Khungurn and Marschner [130] presented a fiber model supporting elliptical cross sections of arbitrary eccentricity. We compare against this model in Figure 62, by adjusting their elliptical cross section to fit ours as closely as possible and using real optical parameters. As the figure shows, even an advanced elliptical fiber model is unable to represent the rich and complex outgoing radiance field of natural fibers. Moreover, even in fibers with an actual elliptical cross section, such as wool, the complex patterns in the reflectance due to the tilted cuticles are better approximated with our model. Finally, the complexity of the azimuthal scattering component requires costly numerical integrations that cannot be performed on the fly during rendering; therefore we rely on pre-computed tabulated data, similar to most current models, to produce highly accurate results.

**GEOMETRIC OPTICS ASSUMPTION** Our computation setup used to generate the BCSDFs (Section 5.3.2) is based on geometric optics.

This means that wave effects such as polarization are not taken into account. However, given that even for polyester the surface is not a perfectly smooth dielectric, the degree of polarization introduced by Fresnel reflection is small, and its impact on the final image would be minimal. In addition, most cloth fibers present a small amount of birefringence [230], specially polyester (see Table 9). We assume isotropic dielectrics, and set an average index of refraction. Our model does not account for diffraction, which will probably appear due to the small fiber's cross section and geometric features. It is unclear how much diffraction would impact the BCSDf, although its effect will probably be slightly masked by the surface roughness and the medium absorption. Finally, some dyes may present fluorescent appearance. We assume only elastic scattering (i.e. with no energy transfer between wavelengths [97]), although a more sophisticated model could include a bispectral BCSDf, similar to the bispectral BRDFs [109]. Although our results show improved accuracy over previous existing models, a deeper exploration on wave-related phenomena remains as future work.

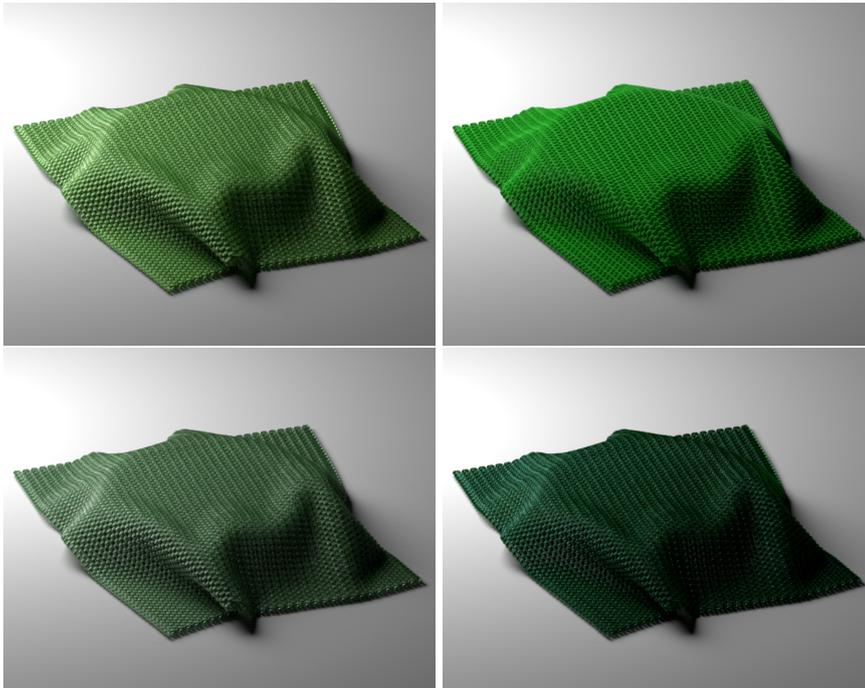


Figure 63: Volumetric renders of a knitted stockinette fabric, where each yarn has one ply and 60 fibers, and a twist of five turns per cm, with fiber types (from left to right): cotton, polyester, silk and merino wool. All fibers are rendered with the same DoS. The almost perfectly dielectric polyester fiber produces a highly saturated fabric since very few light is directly reflected out, while the silk cloth has a clear white specular reflection that previous approaches [131] are unable to capture. Given that wool fibers are significantly wider than the rest, for the same DoS wool fibers absorb a larger amount of light.

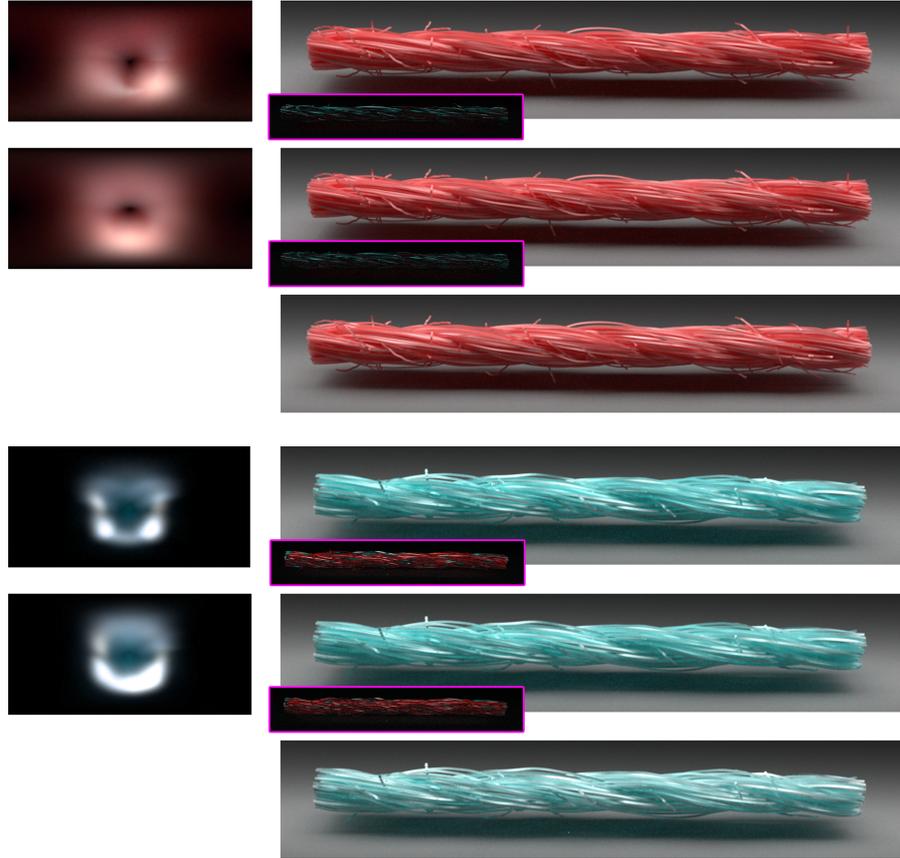


Figure 64: Renders of cotton (red) and silk (green) yarns, captured by Zhao et al. [251]. From top to bottom, using 1) only one of our BCSDFs for every fiber, with varying cross section orientation per fiber; 2) a precomputed average of 15 and 8 BSDFs of cotton and silk, respectively; and 3) five of our BCSDFs randomly distributed among the fibers. Differences between 3), and 1) and 2) are shown in the insets (3x scaled for a better visualization).

## 5.5 RENDERING

Our BCSDFs from Section 5.3.2 can be plugged in both volumetric-based and fiber-based representations of cloth [131], with a few minor modifications. The main difference is the need of keeping track of the the full frame of the fibers, not only their direction  $\omega_f$ , given the lost rotational symmetry of the cross section. This affects both the phase function  $f_r$ , as well as the fiber’s projected area  $A_i$ . We define the frame centered around  $\omega_f$  using the rotation angle  $\psi_f$  which defines the angle between the up-vector of the fibers frame with respect to the plane defined by  $\omega_f$  and the y-axis<sup>2</sup> When rendering, each fiber is assigned with an initial random rotation, as well as a BCSDF from

<sup>2</sup> When  $\omega_f$  is aligned to the y-axis, we compute the rotation based the plane defined by  $\omega_f$  and the x-axis.

our dataset. Both values are uniformly sampled from  $[0, 2\pi)$  and our set of BCSDFs, respectively.

Including fiber’s rotation  $\psi_f$  is straight forward when the fibers are represented by explicit geometry in the form of curves. For the volumetric approach, we build upon the anisotropic radiative formulation by Jakob et al. [117], where the incoming radiance at  $\mathbf{x}$  in direction  $\omega_o$  is:

$$L(\mathbf{x}, \omega_o) = \int_0^s \Lambda(\mathbf{x}_t, \omega_o) \mu_t(\mathbf{x}_t, \omega_o) T_r(\mathbf{x}, \mathbf{x}_t) L_i(\mathbf{x}_t, \omega_o) dt, \quad (15)$$

where  $\mathbf{x}_t = \mathbf{x} - \omega_o t$ ,  $\Lambda(\mathbf{x}_t, \omega_o)$  and  $\mu_t(\mathbf{x}_t, \omega_o)$  are respectively the directionally-resolved albedo and extinction, and  $T_r(\mathbf{x}, \mathbf{x}_t)$  is the transmittance between  $\mathbf{x}$  and  $\mathbf{x}_t$ .  $L_i(\mathbf{x}_t, \omega_o)$  is the inscattered radiance at  $\mathbf{x}_t$  in direction  $\omega_o$ :

$$L_i(\mathbf{x}_t, \omega_o) = \int_{S^2} f_r(\mathbf{x}_t, \omega_o, \omega_i) L(\mathbf{x}_t, \omega_i) d\omega_i, \quad (16)$$

with  $f_r$  being the normalized phase function at  $\mathbf{x}_t$ ,  $L(\mathbf{x}_t, \omega_i)$  represents the incoming radiance at point  $\mathbf{x}_t$  from direction  $\omega_i$ , and  $S^2$  is the spherical domain. We include the fiber’s rotation  $\psi_f$  into the extinction  $\mu_t(\mathbf{x}_t, \omega_o)$  by approximating it as:

$$\begin{aligned} \mu_t(\mathbf{x}, \omega_o) &\approx \frac{1}{V} \sum_{i=1}^N A_i(\omega'_o) \sin(\omega_o, \omega_{f,i}) \\ &\approx \frac{N}{V} \bar{A}(\omega'_o) \sin(\omega_o, \bar{\omega}_f), \end{aligned} \quad (17)$$

where  $N$  is the number of fibers falling into a voxel and  $V$  is the voxel’s volume,  $A_i(\omega'_o)$  and  $\bar{A}(\omega'_o)$  are the projected cross section of fiber  $i$  and the mean projected cross section of all  $N$  fibers in direction  $\omega'_o = R_f(\omega_o)$ , which is the outgoing direction transformed to match the frame of the fibers.  $\omega_{f,i}$  and  $\bar{\omega}_f$  are the direction of fiber  $i$  and the fibers’ mean direction. Note that we approximate both the distributions of fibers directions  $\omega_f$ , rotations  $\psi_f$  and projected cross sections  $A(\omega_o)$  using their respective mean value. The albedo  $\Lambda(\mathbf{x}_t, \omega_o)$  is computed analogously.

**IMPLEMENTATION DETAILS** We implemented our BCSDFs in the physically-based renderer Mitsuba [116] as specialized BSDF and phase functions, fed in both cases with the tangent frame defined by the fibers direction  $\omega_f$  and rotation  $\psi_f$ . We tabulate the BCSDFs, and their CDFs for efficient sampling; the BCSDFs are normalized to one, with a separated table coding the directionally-varying albedo. We also tabulate the projection of the fibers cross section  $A(\psi_f)$ . The total memory cost of each of these tabulations can be found in Table 10. We do not use any compression for the BCSDFs, although standard compression techniques for high-dimensional data (e.g. tensor decomposition [186]) could be applied to significantly reduce the memory cost.

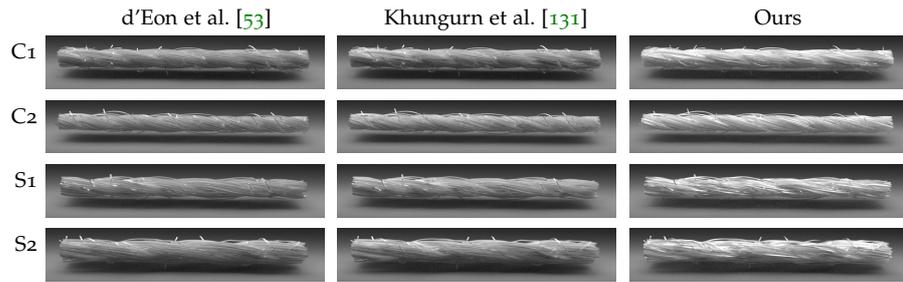


Figure 65: Comparison between renderings of the same base fibers using the models of d’Eon et al. [53] and Khungurn et al. [131] and our model. We fit these parametric models to match our cotton (C) and silk (S) BCSDFs as close as possible. The parameters of the fits, as well as our BCSDFs, can be found in Figure 68.

To account for inter-fiber variability we randomly assign to each fiber a fixed cross section from our database, which spins with the fiber’s local frame; we assume that this cross section does not vary along the fiber. Introducing geometrical variability along the fiber is an interesting avenue of future work.

## 5.6 RESULTS

We demonstrate the use of our BCSDFs with the two main types of representations for cloth: volumetric and explicit fibers geometry. We use the former to render cloth garments illustrating the effect of our BCSDFs on the macroscopic appearance of cloth, and explicit fiber representations for rendering small yarns showing very detailed, close shots. We generate the volumetric garments using the work of Lopez-Moreno et al. [142] over simulations at the yarn level [41, 43]. Figure 44 and Figure 63 show examples on the macroscopic effect on the cloth appearance of an accurate BCSDF. In particular, Figure 63 shows the exact same knitted pattern rendered using our four types of BCSDFs, keeping the rest of the parameters constant. Significant differences on the overall appearance of the garment can be found between each type of fiber.

**EFFECT OF AVERAGE BCSDFS** The effect of a single, detailed high-frequency scattering function can be downplayed due to the multiple bounces in the volume, as well as by the multiple individual fibers with different BCSDFs falling in a single pixel (or voxel in a volumetric representation). In Figure 64 we investigate the effect of having a single BCSDF for all fibers in a yarn, an average BCSDF, and different BCSDFs for each fiber in the yarn. As expected, the most realistic scenario (different BCSDFs for each fiber) introduces the largest variability in terms of specular reflections, leading to a more detailed appearance. This confirms that, for very detailed shots we need to

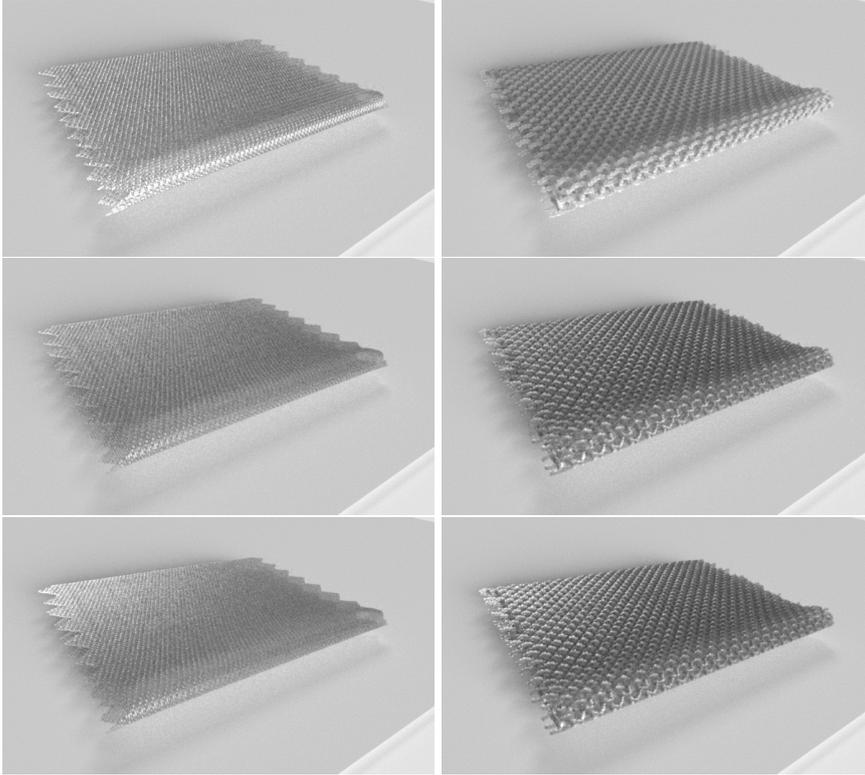


Figure 66: Comparison between our BCSDFs (top), and the scattering models of d’Eon et al. [53] and Khungurn et al. [131], fitted to match our BCSDFs (see Figure 68), for silk (left) and cotton (right) garments, rendered using a volumetric representation of cloth.

preserve the variability on the fibers BCSDFs to obtain optimal results. However, even the less detailed scenario (the average BCSDF) preserves the overall look of the BCSDFs, including an anisotropic look, which suggests that even in a convoluted volumetric rendering the final look will still vary when using detailed BCSDFs rather than parametric simplified models. In the following, we investigate this in more detail.

**COMPARISON WITH PARAMETRIC MODELS** Figure 65 compares the resulting appearance of rendering four high-quality yarns from [251] using our BCSDFs and two state-of-the-art BCSDF models [131, 53] fitted numerically to our BCSDFs (see resulting BCSDFs and fitted parameters in Figure 68). Given the non-linearity of volumetric scattering, small differences in the fit result in very different appearances: this is specially visible in the specular reflections, where current parametric methods are unable to match the anisotropic scattering found in real-world fibers, even in relatively diffuse fabrics such as cotton. This emphasizes that these models are very good for matching macroscopic appearance in top-down optimization frameworks [131], but might lack expressivity for bottom-up definitions of appearance.



Figure 67: Renders of virtual replicas of the *Gütermann* sewing yarns described in Table 11. From left to right, and top to bottom, the yarns are Skala360 (polyester), ORA 120 (polyester), CNe 50 (cotton), and S 303 (silk), with varying DoS.

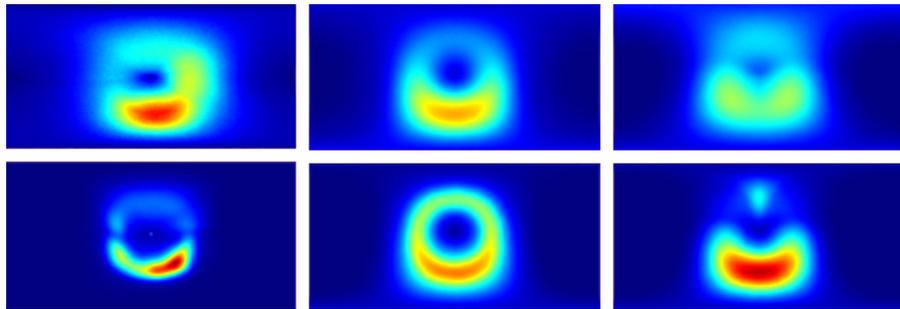


Figure 68: Comparison between our model (left) and the parametric fiber scattering models by Khungurn et al. [131] (middle), and d'Eon et al. [53] (right), fitted to match our BCSDFs. *Left*: slices of our cotton (top) and silk (bottom) BCSDFs, averaged over 15 and 8 different fiber's cross section, respectively. *Center*: numerical fit for the BCSDF of Khungurn et al. [131], with parameters (cotton, silk):  $C_R = 0.001, 0.1$ ,  $C_{TT} = (0.5, 0.25)$ ,  $\beta_R = (32^\circ, 8^\circ)$ ,  $\beta_{TT} = (16^\circ, 16^\circ)$ ,  $\gamma_{TT} = (57^\circ, 46^\circ)$ . *Right*: numerical fit for d'Eon et al.'s BCSDF [53], with parameters  $\beta = (20^\circ, 13^\circ)$ ,  $\text{IOR} = (1.55, 1.55)$ ,  $\mu = (0.1, 0.1)$ .

Figure 66 shows a similar comparison, this time at a macroscopic level on a volumetric representation of a full cloth garment, with fabrics made of silk and cotton, respectively.

### 5.6.1 Validation

We validate our model against sewing thread samples of 100% polyester, cotton and silk, measuring how each yarn scatters light, with different parameters to cover a reasonable range of yarn examples. We build digital replicas of these yarns (see Figure 67 for renders, and Table 11 for detailed specifications), and render them using the BCSDFs presented in Section 5.3.

<i>Sample</i>	Composition	Fiber Length	Twist	Fibers/ply	Plys
Skala 360 (*)	Polyester	continuous	5	24	1
ORA 120	Polyester	3 - 4	8.8	90	2
CNe 50 (*)	Cotton	3 - 4	12.8	70	3
S 303 (*)	Silk	5 - 8	9	90	3

Table 11: Specifications of the set of *Gütermann* sewing yarns rendered in Figure 67. For completeness we also include the *dtex* [g/10000m] measurement, which is a value typically used in industry, being from top to bottom: 74 dtex  $\times$  1; 140 dtex  $\times$  2; 90 dtex  $\times$  3; 90 dtex  $\times$  3. *Scala360*, *CNe 50* and *S 303* are used for validation of our model (Section 5.6.1).

**ACQUISITION SETUP** We capture the reflectance of the yarns using a controlled setup at the optics laboratory. The light source is a 633 nm He-Ne laser, with a spatial filter to produce a clean Gaussian beam. The spatial filter assembly consists of a microscope objective, a pinhole aperture, and a collimating/focusing lens. An iris diaphragm is placed before the photodetector to limit its aperture acceptance to  $\approx 10$  mrad. The photodetector is a silicon PIN-photodiode with conventional transimpedance amplifier circuit, mounted on a goniometer, providing a measurement range of  $360^\circ$ . Since the laser beam is constant, it needs to be modulated by an optical chopper. Also, a lock-in amplifier is referenced to the operating frequency of the modulator to discard ambient light and improve the SNR allowing to extract the signal in the noisy environment. This is done by a data acquisition module connected to a computer with a lock-in amplifier algorithm [34].

Measurement repeatability is over 0.5%. Figure 69, a shows photographs of the full measurements setup. Note that two different apparatus were needed for holding the yarn vertically (Figure 69, b) and horizontally (Figure 69, c). Yarns were carefully attached to the holders with a soft tension to avoid appearance variations due to stretching or shearing.

**DISCUSSION** Following Sadeghi et al. [190], we plot the light scattered by yarns in the longitudinal axis  $\theta$  (Figure 70). These plots compare the reflectances from measurements against the simulated scattering of the virtual replicas. Given the inherent difficulty of capturing the transmission lobe (e.g. calibrating the collimated laser beam, avoiding the saturation of the sensor) we decide to capture only the longitudinal scattering over the upper reflected directions. This allows us to observe the most notable differences in the scattering function between different pairs of fibers-yarns. Despite the limitations of the capturing setup, and the large amount of variables involved, our measurements and their virtual replicas match in terms of relative size and shape of the scattering lobes, and remain consistent between

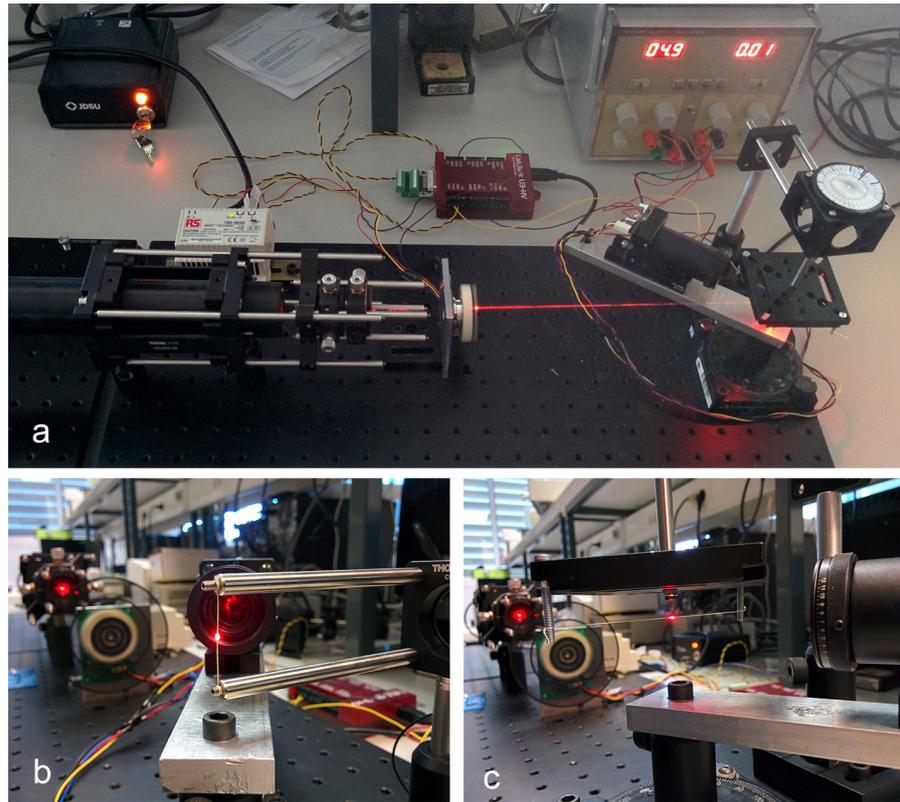


Figure 69: *Top*: Full view of the measurement setup, which includes a laser light source and a photodetector mounted on a goniometer. *Bottom*: Close ups of the two different apparatus for scanning reflectance along  $\psi$  and  $\theta$ , respectively.

yarns with similar (but different) structure and fibers, showing interesting effects such as stronger specular reflections for fibers with higher dye concentrations.

## 5.7 CONCLUSIONS

In this work we have focused on the definition of scattering functions for realistic cloth rendering, for both explicit and volumetric representations. Most previous works have been directed towards accurate modeling of the mesostructure of cloth, including the type of knitting or weaving, or in formulating models for the yarns giving form to these structures. In this regard, we have gone a step further in the level of detail, by focusing on the particular scattering functions of individual fibers.

For this purpose, we have left behind common assumptions from previous fiber scattering models, such as considering fibers as cylinders with circular cross section and, together with physically-based optical parameters directly related with the fabrication properties of the fiber, developed a bottom-up approach for defining the appearance of cloth in a predictive manner. In this regard, ours is the first

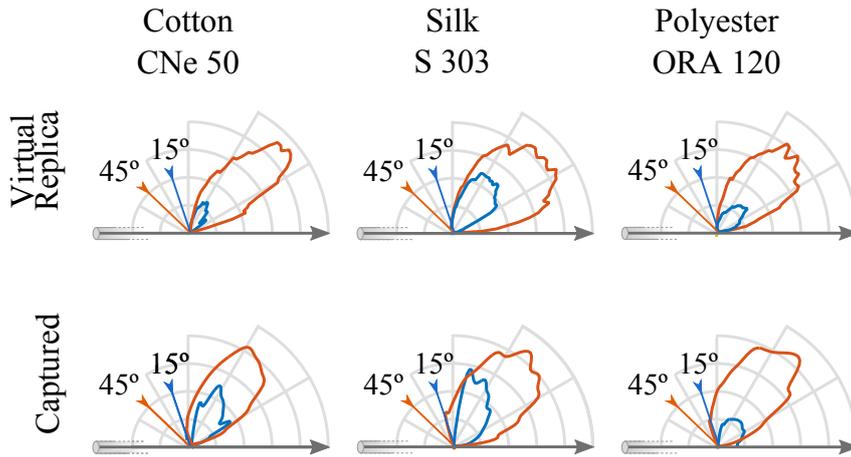


Figure 70: Longitudinal scattering plotting (following the plots in [190]) of three real yarns (bottom) compared to their replicas obtained by getting the best fit approximation (top), under incoming illumination at  $15^\circ$  and  $45^\circ$  from the normal direction of the yarn. From left to right: polyester (ORA 120), silk (S 303), and cotton (CNe50) (specifications in Table 11). Abrupt changes in the measured lobes (see the  $15^\circ$  lobe of silk) are due to the unavoidable occlusions between the sensor and the laser (see Figure 69).

appearance model for fibers accounting for such detailed input. We have shown that the scattering of individual fibers is important for defining the appearance at detailed shots, and analyzed the emerging effect of these different scattering functions on cloth appearance. Further on, we have shown the limitations of previous fiber models on describing the complex features exhibited by the BCSDf of realistic fibers.

While our work is, to our knowledge, the most detailed on defining the reflectance field of cloth fibers, several assumptions have been made to build our model. The main one is related to the geometric optics assumed on the definition of the BCSDf: while the obtained results are sound and agree with the scattering acquired from measured real-world fibers, diffraction would be expected to appear when light interacts with very thin fibers. While previous work has shown that ray-optics are a good model for closely approximating complex electromagnetic phenomena at the diffraction limit (e.g. [189]), it is unclear up to what extent diffraction would affect our model, although the fiber's surface roughness and medium absorption are likely to slightly reduce its effect. Investigating the wave-based phenomena, and whether it can be approximated using data processing as in previous work [189, 58] are interesting avenues for future work. In this context, other effects such as birefringence and fluorescence are also expected to play an important role on fibers appearance.

Despite these limitations our work has shown that more accurate physically-based BCSDfs have an important impact on the appear-

ance of cloth. We believe that, given that rendering highly realistic cloth is a matter of subtle but important details, the approach proposed in this work is a promising path for higher-quality predictive cloth rendering.

## CONCLUSIONS & FUTURE WORK

---

This thesis presents contributions in several topics relevant to the visual perception, modeling and rendering of virtual humans appearance. Starting from an artistic perspective, aiming for a better understanding of the role of stylization on the perception of the characters, we have progressively focused in photorealism and the difficulties involved in modeling the appearance of skin and cloth, two materials that are crucial in virtual humans. In the following we summarize the contributions of each part separately.

**PERCEPTION OF DIGITAL CHARACTERS** In Chapter 2 we presented a thorough study about the effect of stylization on the perception of digital faces. The study focuses in the perceptual effects of different stylizations along shape and material, two of the main aspects that define the appearance of a virtual character. To this aim, several psychophysical experiments were performed by using five manually crafted stylizations of two virtual characters, a male and a female, from very realistic (captured) to highly stylized in both their shape and materials. The results point to the shape as the main contributor to the perceived realism, while the effect of the material stylization becomes more prominent when the realism of the shape increases at the same time. The study also shows how mismatchings in the level of stylization of material and shape are much less salient on abstract characters. In the case of extreme mismatchings, very realistic materials on stylized shapes resulted in the most unappealing characters; instead, subtle stylizations over realistic faces can be preferable, since they remove impurities of the skin, making the character more appeal. Also, when shape and material stylizations are matched, the appeal quickly increases. In summary, material stylization affects strongly the appeal of the character and the shape stylization is very crucial for enhancing the expressivity of the characters.

The study was done over still images where peak expressions were performed by the characters. A very interesting future avenue of work is to validate the observed trends in videos. This would allow to also measure the effect of stylized animation, which was outside the scope of the study due to the size of the stimuli and the complexity of creating a scale of animated expressions. Other possibilities like the lighting conditions or the mood of the illumination, would also be interesting to test.

Although the obtained results are valid just for the particular set of stimuli employed, the study explores combinations of a considerable

amount of variations over shape and material, that were carefully performed following typical designs used in feature animation, with the help of professional artists. And also the observed trends are consistent between the two characters. Thus, the study shed some light on the process of creating appealing virtual characters, through practical guidelines and interesting insights that can inspire future research on the perception of digital faces.

**MODELING SKIN APPEARANCE** Chapter 3 presents a biophysically-based model of the optical properties of the human skin and how they evolve with age. It can simulate the changes in the appearance of elderly skin thanks to a multi-layered structure, together with detailed biophysical parameters coming from the extensive medical and optical tissue literature. These parameters are the ones that affect appearance the most. Namely, variations in the chemical compounds responsible for absorbing light (mainly melanin and haemoglobin); variations in the structure of collagen, which is the main scatterer of light in the skin; and structural changes like the progressive slimming of the skin layers, and other morphological changes like the flattening of the interfaces between the dermis and the epidermis. Thus, our model can be expressed not only through the optical parameters commonly used in computer graphics (absorption and scattering coefficients), but also in terms of biophysical properties or, more intuitively, through high level parameters like age, gender, skin type or lifestyle habits. When any type of skin in terms of the mentioned variables wants to be depicted, the presented model avoids the need for artistic input or costly capturing processes, being suitable for any rendering system that uses diffusion profiles, even in real-time.

An interesting extension for future work would be to consider the dependence on incoming light direction, in order to model anisotropic light transport instead of radially constant diffusion profiles. Also, analytical functions could be fitted to the obtained tabulated profiles to gain efficiency. Another possibility is to use neural networks to learn a model capable to interpolate or even extrapolate new types of skin, making use of the extensive dataset already generated in this work. Last, modeling the time-varying distribution and properties of heterogeneities, like the the increased uneven distribution of melanin with age, would really complete the model, but we are bounded by existing measurements and there is no existing bio-physical data available about it. In any case, the model in its current form represents the most detailed one for simulating changes in aging skin, and it can be potentially used in other domains than graphics, due to its predictive nature, in applications such as cosmetics, dermatology, or tissue optics.

**CLOTH APPEARANCE** Almost half of the content of this thesis is dedicated to cloth. First, from a perceptual side in Chapter 4, then from a modeling and rendering perspective in Chapter 5. The perceptual study [8] consists of a set of psychophysical experiments to explore the effect of dynamics and appearance, two of the main factors that contribute to the perception of realistic cloth. To create the stimuli, we started by a real video footage of several fabrics that cover a reasonable range of visual appearance and dynamic behaviors of the fabric types commonly used in real garments. Then, digital replicas are created to mimic both key attributes, appearance and dynamics, for a latter exploration of every combination of these two attributes. The goal is to gain knowledge about the interplay between appearance and dynamics and provide insights into how efficiency can be maximized without sacrificing plausibility. From the results one can conclude that the appearance dominates over dynamics except for very characteristic movements (e.g.: silk). Interestingly, this effect is more prominent in men, maybe because women often wear a larger variety of fabrics and flowing garments and they are more familiar to their dynamics.

There are many possibilities for future work. A similar study with animated characters wearing garments made of these fabrics could confirm our findings in a more familiar scenario. Also, the study currently focuses on high level attributes, but it would be very interesting to deeply explore the many variables involved in each factor; for instance, the material models or textile patterns / spatial frequency of the textures in the case of the appearance, or dynamics parameters in the case of motion synthesis. Last, other higher level factors could be considered, like different illumination (e.g.: environment lighting), since humans are much less used to see fabrics lit by a spot studio light than under natural lighting conditions.

On the other hand, Chapter 5 is focused on cloth appearance modeling and rendering, particularly centered in defining scattering functions for the individual fibers. This model is the most detailed on defining the reflectance field of cloth fibers, discarding common assumptions about the fibers in computer graphics, like the elliptical shape of the fibers' cross-sections. It relies on physically-based parameters directly related with the manufacturing features of real cloth, based on existing measurements and available data in the textile research community.

As a future avenue of work, the effect of diffraction is not taken into account in the current framework and it is expected to appear when light interacts with fibers as thin as several micrometers, at least the edge diffraction. In the same direction, our BCSDF model relies on geometric optics and does not account for polarization or fluorescence, both effects also expected to play some role on the appear-

ance of fibers. Nevertheless, our work shows how accurately modeling the light scattering of textile fibers in a physically-based way has an important effect on the overall appearance of cloth, both at close up views and larger distances, due to its cumulative effect. This fact was demonstrated through different common BCSDFs: silk, polyester, wool and cotton. This work opens a promising path for a predictive, bottom-up approach to define the appearance of cloth, which can help to render highly realistic cloth, but can also have potential for textile prototyping and other applications in textile industry.

**PERSONAL CONCLUSIONS** During these four years I participated in different projects, in collaboration with external universities and companies from this country and abroad, thus having the opportunity to work in different groups of people in a variety of environments, both in academia and industry. This has allowed me to experience the workflows particular to each atmosphere, helping me to have a wide vision of this field. In that sense, I always tried to retain the aspects and mechanisms I considered best from each world, giving shape to what I am today. Since the beginning, the people surrounding me at the *Graphics & Imaging Lab* and the great collaborative atmosphere and inertia of the group not only increased my passion for computer graphics, it also aroused my interest in computational imaging or computer vision. Leaving aside the invaluable technical knowledge acquired at the lab during these years of lots of work and discussions, the stay at the lab gave me the basic and most important guidelines to be a researcher, capable to think out of the box but also succeed solving concrete problems. Thanks to the multidisciplinary nature of the lab, I really developed my adaptation skills throughout the journey.

Collaborations with other institutions also help in the process of opening your mind to different ways of thinking and working, specially when you are the new incorporation to the crew, as happens with internships in companies. From my point of view, it is a perfect complement to academia during your PhD, since in industry you can experience procedures that are typical from each company but have a common distinctive aspect with respect to academia: the need of always being close to the development departments, making sure what you do is, if not directly applicable to product or the current pipeline, at least of clear help to future research and development, always with practical usefulness in mind. This aspect is very important for me: it is very rewarding when you can see that your work is applied in production, and technicians or artists are making use of it.

Along this process I went through many different emotional states, but the overall sensation of this sort of roller coaster is very positive. I feel very lucky to have met amazing people in the way, and have

lived awesome experiences that made me grow professionally, and are of great value also for my personal life.



## BIBLIOGRAPHY

---

- [1] Identification and ratings of caricatures: Implications for mental representations of faces. *Cognitive Psychology*, 19(4):473 – 497, 1987.
- [2] Seeing liquids from visual motion. *Vision Research*, 109, Part B, 2015.
- [3] Seeing liquids from static snapshots. *Vision Research*, 115, Part B, 2015.
- [4] Alonzo C Addison. Emerging trends in virtual heritage. *IEEE multimedia*, 7(2), 2000.
- [5] Erik Alerstam, William Chun Yip Lo, Tianyi David Han, Jonathan Rose, Stefan Andersson-Engels, and Lothar Lilge. Next-generation acceleration and code optimization for light transport in turbid media using gpus. *Biomedical optics express*, 1(2), 2010.
- [6] Oleg Alexander, Mike Rogers, William Lambeth, Jen-Yuan Chiang, Wan-Chun Ma, Chuan-Chang Wang, and Paul Debevec. The Digital Emily project: Achieving a photorealistic digital actor. *IEEE Computer Graphics and Applications*, 30(4), 2010.
- [7] Oleg Alexander, Graham Fyffe, Jay Busch, Xueming Yu, Ryosuke Ichikari, Andrew Jones, Paul Debevec, Jorge Jimenez, Etienne Danvoye, Bernardo Antionazzi, Mike Eheler, Zybnek Kysela, and Javier von der Pahlen. Digital ira: Creating a real-time photoreal digital actor. In *ACM SIGGRAPH 2013 Posters*, SIGGRAPH '13, 2013.
- [8] Carlos Aliaga, Carol O'Sullivan, Diego Guterrez, and Rasmus Tamstorf. Sackcloth or silk? the impact of appearance vs dynamics on the perception of animated cloth. In *Proceedings of ACM SIGGRAPH Symposium on Applied Perception*, 2015.
- [9] Carlos Aliaga, Carlos Castillo, Diego Gutierrez, Miguel A. Otaduy, Jorge Lopez-Moreno, and Adrian Jarabo. A fiber-level model for predictive cloth rendering. In *SIGGRAPH 2016 Posters*, 2016.
- [10] Carlos Aliaga, Carlos Castillo, Diego Gutierrez, Miguel A. Otaduy, Jorge Lopez-Moreno, and Adrian Jarabo. An appearance model for textile fibers. *Computer Graphics Forum (Proc. of the Eurographics Symposium on Rendering)*, 36(4), 2017.

- [11] Carlos Aliaga, Carol O'Sullivan, Diego Guterrez, and Rasmus Tamstorf. Evaluating the impact of appearance vs dynamics on cloth perception. *ACM Transactions on Applied Perception*, 2017.
- [12] Thomas R. Alley and Michael R. Cunningham. Averaged faces are attractive, but very attractive faces are not average. *Psychological Science*, 2(2), 1991.
- [13] R Rox Anderson and John A Parrish. The optics of human skin. *Journal of Investigative Dermatology*, 77(1), 1981.
- [14] E. Angelopoulo, R. Molana, and K. Daniilidis. Multispectral skin color modeling. In *Computer Vision and Pattern Recognition, 2001. CVPR 2001. Proceedings of the 2001 IEEE Computer Society Conference on*, volume 2, 2001.
- [15] M Aslan, J Yamada, MP Mengüç, and JA Thomasson. Characterization of individual cotton fibers via light-scattering experiments. *Journal of thermophysics and heat transfer*, 17(4), 2003.
- [16] Mustafa Aslan, Jun Yamada, Pinar Menguc, and Alex Thomasson. Radiative properties of individual cotton fibers: Experiments and predictions. In *8th AIAA/ASME Joint Thermophysics and Heat Transfer Conference*, page 3325, 2002.
- [17] David Baraff and Andrew Witkin. Large steps in cloth simulation. In *Proc. SIGGRAPH'98*, 1998.
- [18] Gladimir VG Baranoski and Aravind Krishnaswamy. *Light and Skin Interactions: Simulations for Computer Graphics Applications*. Morgan Kaufmann, 2010.
- [19] Alexey N Bashkatov, Elina A Genina, and Valery V Tuchin. Optical properties of skin, subcutaneous, and muscle tissues: a review. *Journal of Innovative Optical Health Sciences*, 4(01), 2011.
- [20] AN Bashkatov, EA Genina, VI Kochubey, and VV Tuchin. Optical properties of human skin, subcutaneous and mucous tissues in the wavelength range from 400 to 2000 nm. *Journal of Physics D: Applied Physics*, 38(15), 2005.
- [21] Thabo Beeler, Bernd Bickel, Paul Beardsley, Bob Sumner, and Markus Gross. High-quality single-shot capture of facial geometry. *ACM Transactions on Graphics*, 29(4), 2010.
- [22] Thabo Beeler, Bernd Bickel, Paul Beardsley, Bob Sumner, and Markus Gross. High-quality single-shot capture of facial geometry. *ACM Trans. Graph.*, 29(4), July 2010.

- [23] Thabo Beeler, Bernd Bickel, Gioacchino Noris, Paul Beardsley, Steve Marschner, Robert W Sumner, and Markus Gross. Coupled 3d reconstruction of sparse facial hair and skin. *ACM Transactions on Graphics (ToG)*, 31(4), 2012.
- [24] Philip Benson and David Perrett. Face to face with the perfect image. *New Scientist*, 133(1809), 1992.
- [25] Pascal Bérard, Derek Bradley, Maurizio Nitti, Thabo Beeler, and Markus H Gross. High-quality capture of eyes. *ACM Trans. Graph.*, 33(6), 2014.
- [26] Wenyan Bi and Bei Xiao. Perceptual constancy of mechanical properties of cloth under variation of external forces. In *Proceedings of the ACM Symposium on Applied Perception, SAP '16*, 2016.
- [27] Bernd Bickel, Peter Kaufmann, Mélina Skouras, Bernhard Thomaszewski, Derek Bradley, Thabo Beeler, Phil Jackson, Steve Marschner, Wojciech Matusik, and Markus Gross. Physical face cloning. *ACM Transactions on Graphics (TOG)*, 31(4), 2012.
- [28] Laurence Boissieux, Gergo Kiss, Nadia Magnenat Thalmann, and Prem Kalra. Simulation of skin aging and wrinkles with cosmetics insight. In *Computer Animation and Simulation 2000*. 2000.
- [29] Stephan A Bolliger, Michael J Thali, Steffen Ross, Ursula Buck, Silvio Naether, and Peter Vock. Virtual autopsy using imaging: bridging radiologic and forensic sciences. a review of the virtopsy and similar projects. *European radiology*, 18(2), 2008.
- [30] RA D Booth, BA Goddard, and A Paton. Measurement of fat thickness in man: a comparison of ultrasound, harpenden calipers and electrical conductivity. *British Journal of Nutrition*, 20(04), 1966.
- [31] Katherine L Bouman, Bei Xiao, Peter Battaglia, and William T Freeman. Estimating the material properties of fabric from video. In *IEEE Conf. on Comp. Vision (ICCV)*, pages 1984–1991, 2013.
- [32] MC Branchet, S Boisnic, C Frances, C Lesty, and L Robert. Morphometric analysis of dermal collagen fibers in normal human skin as a function of age. *Archives of gerontology and geriatrics*, 13(1), 1991.
- [33] Robert Bridson, Ronald Fedkiw, and John Anderson. Robust treatment of collisions, contact and friction for cloth animation. *ACM Trans. Graph.*, 21(3), 2002.

- [34] Richard Burdett. Amplitude modulated signals: The lock-in amplifier. In *Handbook of Measuring System Design*. John Wiley & Sons, Ltd, 2005. ISBN 9780471497394.
- [35] Tyler J. Burleigh, Jordan R. Schoenherr, and Guy L. Lacroix. Does the Uncanny Valley exist? An empirical test of the relationship between eeriness and the human likeness of digitally created faces. *Computers in Human Behavior*, 29(3), 2013.
- [36] Andrew J Calder. Facial emotion recognition after bilateral amygdala damage: Differentially severe impairment of fear. *Cognitive Neuropsychology*, 13(5), 1996.
- [37] H. Chen, R. Russell, and K. Nakayama. Crossing the “Uncanny Valley”: Adaptation to cartoon faces can influence perception of human faces. *Perception*, (39), 2010.
- [38] Matt Jen-Yuan Chiang, Benedikt Bitterli, Chuck Tappan, and Brent Burley. A practical and controllable hair and fur model for production path tracing. In *ACM SIGGRAPH 2015 Talks*, SIGGRAPH '15, 2015.
- [39] Matt Jen-Yuan Chiang, Benedikt Bitterli, Chuck Tappan, and Brent Burley. A practical and controllable hair and fur model for production path tracing. *Computer Graphics Forum*, 35(2), 2016.
- [40] Jin Ho Chung, Kiichiro Yano, Mi Kyung Lee, Choon Shik Youn, Jin Young Seo, Kyu Han Kim, Kwang Hyun Cho, Hee Chul Eun, and Michael Detmar. Differential effects of photoaging vs intrinsic aging on the vascularization of human skin. *Archives of dermatology*, 138(11), 2002.
- [41] Gabriel Cirio, Jorge Lopez-Moreno, David Miraut, and Miguel A. Otaduy. Yarn-level simulation of woven cloth. *ACM Trans. Graph.*, 33(6), 2014.
- [42] Gabriel Cirio, Jorge Lopez-Moreno, David Miraut, and Miguel A. Otaduy. Yarn-level simulation of woven cloth. *ACM Trans. Graph.*, 33(6), 2014.
- [43] Gabriel Cirio, Jorge Lopez-Moreno, and Miguel A. Otaduy. Efficient simulation of knitted cloth using persistent contacts. In *Eurographics Symposium on Computer Animation*, pages 55–61, 2015.
- [44] Robert L Cook and Kenneth E. Torrance. A reflectance model for computer graphics. *ACM Trans. Graph.*, 1(1), 1982.
- [45] Douglas Cunningham and Christian Wallraven. *Experimental Design: From User Studies to Psychophysics*. AK Peters/CRC Press, 2011.

- [46] Kristin J Dana, Bram Van Ginneken, Shree K Nayar, and Jan J Koenderink. Reflectance and texture of real-world surfaces. *ACM Trans. Graph.*, 18(1), 1999.
- [47] Kristin J Dana, Bram Van Ginneken, Shree K Nayar, and Jan J Koenderink. Reflectance and texture of real-world surfaces. *ACM Trans. Graph.*, 18(1), 1999.
- [48] Paul Debevec. The light stages and their applications to photo-real digital actors. *SIGGRAPH Asia*, 2(4), 2012.
- [49] Paul Debevec, Tim Hawkins, Chris Tchou, Haarm-Pieter Duiker, Westley Sarokin, and Mark Sagar. Acquiring the reflectance field of a human face. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '00*, 2000.
- [50] Eugene d'Eon and Geoffrey Irving. A quantized-diffusion model for rendering translucent materials. *ACM Transactions on Graphics*, 30(4), 2011.
- [51] Eugene d'Eon and David Luebke. Advanced techniques for realistic real-time skin rendering. In *GPU Gems 3*, chapter 14. 2007.
- [52] Eugene d'Eon, David Luebke, and Eric Enderton. Efficient rendering of human skin. *EGSR '07*, 2007.
- [53] Eugene d'Eon, Guillaume Francois, Martin Hill, Joe Letteri, and Jean-Marie Aubry. An energy-conserving hair reflectance model. *Computer Graphics Forum*, 30(4), 2011.
- [54] Eugene d'Eon, Guillaume Francois, Martin Hill, Joe Letteri, and Jean-Marie Aubry. An energy-conserving hair reflectance model. *Computer Graphics Forum*, 30(4), 2011.
- [55] Angélica Díaz, Ramaz Katsarava, and Jordi Puiggalí. Synthesis, properties and applications of biodegradable polymers derived from diols and dicarboxylic acids: From polyesters to poly (ester amide) s. *International Journal of Molecular Sciences*, 15(5), 2014.
- [56] Vanderson Dill, Laura Mattos Flach, Rafael Hocevar, Christian Lykawka, Soraia Raupp Musse, and Márcio Sarroglia Pinho. Evaluation of the Uncanny Valley in CG characters. In *Intelligent Virtual Agents*, 2012.
- [57] John Dingliana and Carol O'Sullivan. Graceful degradation of collision handling in physically based animation. *Comp. Graph. Forum*, 19(3), 2000.

- [58] Zhao Dong, Bruce Walter, Steve Marschner, and Donald P. Greenberg. Predicting appearance from measured microgeometry of metal surfaces. *ACM Trans. Graph.*, 35(1), 2015.
- [59] Craig Donner and Henrik Wann Jensen. Light diffusion in multi-layered translucent materials. *ACM Transactions on Graphics*, 24(3), 2005.
- [60] Craig Donner and Henrik Wann Jensen. A spectral bssrdf for shading human skin. In *EGSR '06*, 2006.
- [61] Craig Donner, Tim Weyrich, Eugene d'Eon, Ravi Ramamoorthi, and Szymon Rusinkiewicz. A layered, heterogeneous reflectance model for acquiring and rendering human skin. *ACM Transactions on Graphics*, 27(5), 2008.
- [62] Sylvain Duchêne, Carlos Aliaga, Tania Pouli, and Patrick Perez. Mixed illumination analysis in single image for interactive color grading. In *Expressive 2017 The Joint Symposium on Computational Aesthetics and Sketch-Based Interfaces and Modeling and Non-Photorealistic Animation and Rendering*. The Eurographics Association, 2017.
- [63] Jose I Echevarria, Derek Bradley, Diego Gutierrez, and Thabo Beeler. Capturing and stylizing hair for 3d fabrication. *ACM Transactions on Graphics (ToG)*, 33(4), 2014.
- [64] C Edwards, R Heggie, and R Marks. A study of differences in surface roughness between sun-exposed and unexposed skin with age. *Photodermatology, photoimmunology & photomedicine*, 19(4), 2003.
- [65] Mariko Egawa, Tetsuji Hirao, and Motoji Takahashi. In vivo estimation of stratum corneum thickness from water concentration profiles obtained with raman spectroscopy. *Acta dermatovenereologica*, 87(1), 2007.
- [66] Christian Eisenacher, Gregory Nichols, Andrew Selle, and Brent Burley. Sorted deferred shading for production path tracing. *Comp. Graph. Forum*, 32(4), 2013.
- [67] Paul Ekman. Universal and cultural differences in facial expression of emotion. *Proceedings of Nebraska Symposium on Motivation*, 19, 1972.
- [68] Shaojing Fan, Tian-Tsong Ng, Jonathan S. Herberg, Bryan L. Koenig, and Shiqing Xin. Real or fake? Human judgments about photographs and computer-generated images of faces. In *SIGGRAPH Asia 2012 Technical Briefs*, 2012.

- [69] Miranda A Farage, Kenneth W Miller, and Howard I Maibach. Degenerative changes in aging skin. In *Textbook of aging skin*. Springer, 2010.
- [70] Miranda A Farage, Kenneth W Miller, Howard I Maibach, Miranda A Farage, Kenneth W Miller, and Howard I Maibach. *Textbook of aging skin*. Springer, 2010.
- [71] Neil A Fenske and Clifford W Lober. Structural and functional changes of normal aging skin. *Journal of the American Academy of Dermatology*, 15(4), 1986.
- [72] Anne E Ferrey, Tyler J Burleigh, and Mark J Fenske. Stimulus-category competition, inhibition, and affective devaluation: a novel account of the uncanny valley. *Frontiers in psychology*, 6, 2015.
- [73] Ylva Ferstl, Elena Kokkinara, and Rachel McDonnell. Do i trust you, abstract creature?: a study on personality perception of abstract virtual faces. In *Proceedings of the ACM Symposium on Applied Perception*, 2016.
- [74] James A Ferwerda, Fabio Pellacini, and Donald P Greenberg. Psychophysically based model of surface gloss perception. In *Proc. Photonics West—Electronic Imaging*, 2001.
- [75] Jirí Filip, Mike J Chantler, Patrick R Green, and Michal Haindl. A psychophysically validated metric for bidirectional texture data reduction. *ACM Trans. Graph.*, 27(5), 2008.
- [76] B. Fink, L. Bunse, P. J. Matts, and D. D’Emiliano. Visible skin colouration predicts perception of male facial age, health and attractiveness. *International Journal of Cosmetic Science*, 34(4), 2012.
- [77] Bernhard Fink and Paul J. Matts. The effects of skin colour distribution and topography cues on the perception of female facial age and health. *Journal of the European Academy of Dermatology and Venereology*, 22(4), 2008.
- [78] Bernhard Fink, Karl Grammer, and Paul J. Matts. Visible skin color distribution plays a role in the perception of age, attractiveness, and health in female faces. *Evolution and Human Behavior*, 27(6), 2006.
- [79] Thomas B Fitzpatrick. The validity and practicality of sun-reactive skin types i through vi. *Archives of dermatology*, 124(6), 1988.
- [80] Reuben Fleming, Betty J Mohler, Javier Romero, Michael J Black, Martin Breidt, IV Piryankova, JK Stefanucci, J Romero, S de la Rosa, MJ Black, et al. Appealing female avatars from 3d body

- scans: Perceptual effects of stylization. *ACM Transactions on Applied Perception*, 11, 2017.
- [81] Roland W Fleming and Heinrich H Bülthoff. Low-level image cues in the perception of translucent materials. *ACM Trans. Applied Perception*, 2(3), 2005.
- [82] Roland W Fleming, Ron O Dror, and Edward H Adelson. How do humans determine reflectance properties under unknown illumination? In *Proc. IEEE Workshop on Identifying Objects across Variations in Lighting: Psychophysics and Computation*, 2001.
- [83] Roland W Fleming, Ron O Dror, and Edward H Adelson. Real-world illumination and the perception of surface reflectance properties. *Journal of Vision*, 3(5), 2003.
- [84] Marcos Garcia, John Dingliana, and Carol O'Sullivan. Perceptual evaluation of cartoon physics: accuracy, attention, appeal. In *Proc. Symp. on Applied Perception in Graph. and Visualization*, 2008.
- [85] Tom Geller. Overcoming the Uncanny Valley. *IEEE Computer Graphics and Applications*, 28(4), 2008.
- [86] Abhijeet Ghosh, Tim Hawkins, Pieter Peers, Sune Frederiksen, and Paul Debevec. Practical modeling and acquisition of layered facial reflectance. *ACM Transactions on Graphics*, 27(5), 2008.
- [87] Abhijeet Ghosh, Graham Fyffe, Borom Tunwattanapong, Jay Busch, Xueming Yu, and Paul Debevec. Multiview face capture using polarized spherical gradient illumination. *ACM Transactions on Graphics*, 30(6), 2011.
- [88] Adam Giangreco, Stephen J Goldie, Virgilio Failla, Gaëlle Saintigny, and Fiona M Watt. Human skin aging is associated with reduced expression of the stem cell markers  $\beta 1$  integrin and mcsp. *Journal of Investigative Dermatology*, 130(2), 2010.
- [89] Barbara A Gilchrest and George Szabo. Effects of aging and chronic sun exposure on melanocytes in human skin. *Journal of Investigative Dermatology*, 73(2), 1979.
- [90] Gunilla Gillberg and David Kemp. Surface characterization of polyester fibers. *Journal of Applied Polymer Science*, 26(6), 1981.
- [91] Ioannis Gkioulekas, Bei Xiao, Shuang Zhao, Edward H Adelson, Todd Zickler, and Kavita Bala. Understanding the role of phase function in translucent appearance. *ACM Trans. Graph.*, 32(5), 2013.

- [92] Aleksey Golovinskiy, Wojciech Matusik, Hanspeter Pfister, Szymon Rusinkiewicz, and Thomas Funkhouser. A statistical model for synthesis of detailed facial geometry. *ACM Transactions on Graphics*, 25(3), 2006.
- [93] Bruce Gooch, Erik Reinhard, and Amy Gooch. Human facial illustrations: Creation and psychophysical evaluation. *ACM Transactions on Graphics*, 23(1), 2004.
- [94] Paul Graham, Borom Tunwattanapong, Jay Busch, Xueming Yu, Andrew Jones, Paul Debevec, and Abhijeet Ghosh. Measurement-based synthesis of facial microgeometry. *Computer Graphics Forum*, 32(2pt3), 2013.
- [95] Maureen M Grasso, Bruce D Hunn, and Ardis M Rewerts. Effect of textile properties in evaluating a directional shading fabric. *Textile research journal*, 67(4):233–247, 1997.
- [96] Robert D. Green, Karl F. MacDorman, Chin-Chang Ho, and Sandosh Vasudevan. Sensitivity to the proportions of faces that vary in human likeness. *Computers in Human Behavior*, 24(5), 2008.
- [97] Diego Gutierrez, Francisco Seron, Adolfo Muñoz, and Oscar Anson. Visualizing underwater ocean optics. *Computer Graphics Forum*, 27(2), 2008.
- [98] Ralf Habel, Per H. Christensen, and Wojciech Jarosz. Photon beam diffusion: A hybrid monte carlo method for subsurface scattering. *Computer Graphics Forum*, 32(4), 2013.
- [99] AA Hamza, TZN Sokkar, MA El-Morsy, and MAE Nawareg. Automatic determination of refractive index profile of fibers having regular and/or irregular transverse sections considering the refraction of light rays by the fiber. *Optics Communications*, 282(1), 2009.
- [100] Donghui Han, Shu-wei Hsu, Ann McNamara, and John Keyser. Believability in simplifications of large scale physically based simulation. In *Proc. ACM Symp. on Applied Perception*, 2013.
- [101] David Hanson. Expanding the Aesthetics Possibilities for Humanlike Robots. In *Proceedings of IEEE Humanoid Robotics Conference, special session on the Uncanny Valley*, 2005.
- [102] Abu Naser Md Ahsanul Haque, MA Hannan, and Md Masud Rana. Compatibility analysis of reactive dyes by exhaustion-fixation and adsorption isotherm on knitted cotton fabric. *Fashion and Textiles*, 2(1), 2015.

- [103] David Harmon, Etienne Vouga, Rasmus Tamstorf, and Eitan Grinspun. Robust treatment of simultaneous collisions. *ACM Trans. Graph.*, 27(3), 2008.
- [104] John WS Hearle and William Ernest Morton. *Physical properties of textile fibres*. 2008.
- [105] Eric Heitz, Jonathan Dupuy, Cyril Crassin, and Carsten Dachsbacher. The SGGX microflake distribution. *ACM Trans. Graph.*, 34(4), 2015.
- [106] Chin-Chang Ho and Karl F. MacDorman. Revisiting the uncanny valley theory: Developing and validating an alternative to the godspeed indices. *Computers in Human Behavior*, 26(6), 2010.
- [107] Helene Hoffman and Dzung Vu. Virtual reality: teaching tool of the twenty-first century?. *Academic Medicine*, 72(12), 1997.
- [108] Ludovic Hoyet, Rachel McDonnell, and Carol O'Sullivan. Push it real: perceiving causality in virtual interactions. *ACM Trans. Graph.*, 31(4), 2012.
- [109] Matthias B. Hullin, Johannes Hanika, Boris Ajdin, Hans-Peter Seidel, Jan Kautz, and Hendrik P. A. Lensch. Acquisition and analysis of bispectral bidirectional reflectance and reradiation distribution functions. *ACM Trans. Graph.*, 29(4), 2010.
- [110] Takanori Igarashi, Ko Nishino, and Shree K Nayar. The appearance of human skin: A survey. *Foundations and Trends in Computer Graphics and Vision*, 3(1), 2007.
- [111] Jose A. Iglesias-Guitian, Carlos Aliaga, Adrian Jarabo, and Diego Gutierrez. A biophysically-based model of the optical properties of skin aging. *Computer Graphics Forum (EUROGRAPHICS 2015)*, 34(2), 2015.
- [112] Piti Irawan and Steve Marschner. Specular reflection from woven cloth. *ACM Trans. Graph.*, 31(1), 2012.
- [113] Piti Irawan and Steve Marschner. Specular reflection from woven cloth. *ACM Trans. Graph.*, 31(1), 2012.
- [114] Dr Rachael Jack, Mr Oliver Garrod, and Philippe Schyns. Dynamic facial expressions of emotion transmit an evolving hierarchy of signals over time. *Current Biology*, 24(2), 2014.
- [115] Steven L Jacques. Optical properties of biological tissues: a review. *Physics in medicine and biology*, 58(11), 2013.
- [116] Wenzel Jakob. Mitsuba renderer, 2010. <http://www.mitsuba-renderer.org>.

- [117] Wenzel Jakob, Adam Arbree, Jonathan T Moon, Kavita Bala, and Steve Marschner. A radiative transfer framework for rendering materials with anisotropic structure. *ACM Trans. Graph.*, 29(4), 2010.
- [118] Adrian Jarabo, Tom Van Eyck, Veronica Sundstedt, Kavita Bala, Diego Gutierrez, and Carol O’Sullivan. Crowd light: Evaluating the perceived fidelity of illuminated dynamic scenes. *Comp. Graph. Forum*, 31(2pt3), 2012.
- [119] Adrian Jarabo, Hongzhi Wu, Julie Dorsey, Holly Rushmeier, and Diego Gutierrez. Effects of approximate filtering on the appearance of bidirectional texture functions. *IEEE Trans. Visualization and Comp. Graph.*, 20(6), 2014.
- [120] Henrik Wann Jensen and Juan Buhler. A rapid hierarchical rendering technique for translucent materials. *ACM Transactions on Graphics (TOG)*, 21(3), 2002.
- [121] Henrik Wann Jensen, Stephen R. Marschner, Marc Levoy, and Pat Hanrahan. A practical model for subsurface light transport. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH ’01*, 2001.
- [122] Jorge Jimenez, Veronica Sundstedt, and Diego Gutierrez. Screen-space perceptual rendering of human skin. *ACM Transactions on Applied Perception*, 6(4), 2009.
- [123] Jorge Jimenez, Timothy Scully, Nuno Barbosa, Craig Donner, Xenxo Alvarez, Teresa Vieira, Paul Matts, Verónica Orvalho, Diego Gutierrez, and Tim Weyrich. A practical appearance model for dynamic facial color. *ACM Transactions on Graphics*, 29(6), 2010.
- [124] Jorge Jimenez, David Whelan, Veronica Sundstedt, and Diego Gutierrez. Real-time realistic skin translucency. *IEEE Computer Graphics and Applications*, 30(4), 2010.
- [125] Jorge Jimenez, Karoly Zsolnai, Adrian Jarabo, Christian Freude, Thomas Auzinger, Xian-Chun Wu, Javier von der Pahlen, Michael Wimmer, and Diego Gutierrez. Separable subsurface scattering. *Computer Graphics Forum*, 2015.
- [126] Jorge Jimenez, Károly Zsolnai, Adrian Jarabo, Christian Freude, Thomas Auzinger, Xian-Chun Wu, Javier von der Pahlen, Michael Wimmer, and Diego Gutierrez. Separable subsurface scattering. *Computer Graphics Forum*, 2015.
- [127] James T Kajiya and Timothy L Kay. Rendering fur with three dimensional textures. In *Proceedings of SIGGRAPH*, 1989.

- [128] CW Kan. Surface morphological study of low temperature plasma treated wool—a time dependence study. 2007.
- [129] Takahiro Kawabe and Shin'ya Nishida. Seeing jelly: Judging elasticity of a transparent object. In *Proceedings of the ACM Symposium on Applied Perception, SAP '16*, 2016.
- [130] Pramook Khungurn and Steve Marschner. Azimuthal scattering from elliptical hair fibers. *ACM Trans. Graph.*, 2016.
- [131] Pramook Khungurn, Daniel Schroeder, Shuang Zhao, Kavita Bala, and Steve Marschner. Matching real fabrics with micro-appearance models. *ACM Trans. Graph.*, 35(1), 2015.
- [132] Pramook Khungurn, Daniel Schroeder, Shuang Zhao, Kavita Bala, and Steve Marschner. Matching real fabrics with micro-appearance models. *ACM Trans. Graph.*, 35(1), 2015.
- [133] Oliver Klehm, Fabrice Rousselle, Marios Papas, Derek Bradley, Christophe Hery, Bernd Bickel, Wojciech Jarosz, and Thabo Beeler. Recent advances in facial appearance capture. *Computer Graphics Forum (Proceedings of Eurographics)*, 34(2), 2015.
- [134] Oliver Klehm, Fabrice Rousselle, Marios Papas, Derek Bradley, Christophe Hery, Bernd Bickel, Wojciech Jarosz, and Thabo Beeler. Recent advances in facial appearance capture. In *Computer Graphics Forum*, volume 34, 2015.
- [135] Aravind Krishnaswamy and Gladimir VG Baranoski. A biophysically-based spectral model of light interaction with human skin. *Computer Graphics Forum*, 23(3), 2004.
- [136] Jaroslav Křivánek, James A Ferwerda, and Kavita Bala. Effects of global illumination approximations on material appearance. *ACM Trans. Graph.*, 29(4), 2010.
- [137] Zhengxiong Li, Yanjun Xing, and Jinjin Dai. Superhydrophobic surfaces prepared from water glass and non-fluorinated alkylsilane on cotton substrates. *Applied Surface Science*, 254(7), 2008.
- [138] Anthony C Little and Peter JB Hancock. The role of masculinity and distinctiveness in judgments of human male facial attractiveness. *British Journal of Psychology*, 93(4), 2002.
- [139] Xiaosong Liu and Fumei Wang. Visible light shielding performance of fabrics with non-circular cross section fiber. *Journal of Engineered Fabrics & Fibers (JEFF)*, 7(3), 2012.
- [140] Patric Ljung, Calle Winskog, Anders Persson, Claes Lundstrom, and Anders Ynnerman. Full body virtual autopsies using a state-of-the-art volume rendering pipeline. *IEEE transactions on visualization and computer graphics*, 12(5), 2006.

- [141] J Lock-Andersen, P Therkildsen, F Olivarius, M Gniadecka, K Dahlstrom, T Poulsen, and H-C Wulf. Epidermal thickness, skin pigmentation and constitutive photosensitivity. *Photodermatology, photoimmunology & photomedicine*, 13(4), 1997.
- [142] Jorge Lopez-Moreno, David Miraut, Gabriel Cirio, and Miguel A. Otaduy. Sparse GPU voxelization of yarn-level cloth. *Computer Graphics Forum*, 2015.
- [143] Wan-Chun Ma, Tim Hawkins, Pieter Peers, Charles-Felix Chabert, Malte Weiss, and Paul Debevec. Rapid acquisition of specular and diffuse normal maps from polarized spherical gradient illumination. In *EGSR '07*, 2007.
- [144] Karl F. MacDorman. Subjective ratings of robot video clips for human likeness, familiarity, and eeriness: An exploration of the Uncanny Valley. In *Proceedings of ICCS/CogSci Long Symposium: Toward Social Mechanisms of Android Science*, 2006.
- [145] Karl F. MacDorman, Robert D. Green, Chin-Chang Ho, and Clinton T. Koch. Too real for comfort? Uncanny responses to computer generated faces. *Computers in Human Behavior*, 25(3), 2009.
- [146] Toshiro Makino and Jun-ichi Horiba. Scattering of radiation by a fiber with a rough surface. *Heat Transfer - Asian Research*, 28(4), 1999.
- [147] Stephen R Marschner, Stephen H Westin, Eric PF Lafortune, Kenneth E Torrance, and Donald P Greenberg. Image-based BRDF measurement including human skin. In *Rendering Techniques 1999*. Springer, 1999.
- [148] Stephen R Marschner, Henrik Wann Jensen, Mike Cammarano, Steve Worley, and Pat Hanrahan. Light scattering from human hair fibers. *ACM Trans. Graph.*, 22(3), 2003.
- [149] Stephen R. Marschner, Henrik Wann Jensen, Mike Cammarano, Steve Worley, and Pat Hanrahan. Light scattering from human hair fibers. In *ACM SIGGRAPH 2003 Papers*, SIGGRAPH '03, 2003.
- [150] Belen Masia, Gordon Wetzstein, Carlos Aliaga, Ramesh Raskar, and Diego Gutierrez. Display adaptive 3d content remapping. *Computers and Graphics*, 37(8):983 – 996, 2013.
- [151] SJ Matcher, M Cope, and DT Delpy. In vivo measurements of the wavelength dependence of tissue-scattering coefficients between 760 and 900 nm measured with time-resolved spectroscopy. *Applied Optics*, 36(1), 1997.

- [152] Scott McCloud. *Understanding Comics*. 1993.
- [153] Rachel McDonnell, Simon Dobbyn, Steven Collins, and Carol O'Sullivan. Perceptual evaluation of lod clothing for virtual humans. In *Proc. ACM SIGGRAPH/Eurographics Symp. Computer Animation*, 2006.
- [154] Rachel McDonnell, Martin Breidt, and Heinrich H. Bühlhoff. Render me real? Investigating the effect of render style on the perception of animated virtual humans. *ACM Transactions on Graphics*, 31(4), 2012.
- [155] Igor V Meglinski and Stephen J Matcher. Quantitative assessment of skin layers absorption and skin reflectance spectra simulation in the visible and near-infrared spectral regions. *Physiological measurement*, 23(4), 2002.
- [156] Johannes Meng, Marios Papas, Ralf Habel, Carsten Dachsbacher, Steve Marschner, Markus Gross, and Wojciech Jarosz. Multi-scale modeling and rendering of granular materials. *ACM Trans. Graph.*, 34(4), 2015.
- [157] G.A. Milliken and D.E. Johnson. *Analysis of Messy Data*. Number v. 1. 1992.
- [158] Masahiro Mori, Karl F. MacDorman, and Norri Kageki. The Uncanny Valley [From the field]. *IEEE Robotics and Automation Magazine*, 19(2), 2012.
- [159] Judith R Mourant, Tamika Fuselier, James Boyer, Tamara M Johnson, and Irving J Bigio. Predictions and measurements of scattering and absorption over broad wavelength ranges in tissue phantoms. *Applied Optics*, 36(4), 1997.
- [160] Thomas Müller, Marios Papas, Markus Gross, Wojciech Jarosz, and Jan Novák. Efficient rendering of heterogeneous polydisperse granular media. *ACM Trans. Graph.*, 35(6), 2016.
- [161] Maryam Mustafa and Marcus Magnor. Eeg based analysis of the perception of computer-generated faces. In *Proceedings of the 13th European Conference on Visual Media Production (CVMP 2016)*, page 4, 2016.
- [162] Karol Myszkowski. Perception-based global illumination, rendering, and animation techniques. In *Proc. Spring Conf. on Comp. Graph.*, 2002.
- [163] Koki Nagano, Graham Fyffe, Oleg Alexander, Jernej Barbič, Hao Li, Abhijeet Ghosh, and Paul Debevec. Skin microstructure deformation with displacement map convolution. *ACM Trans. Graph.*, 34(4), 2015.

- [164] Giljoo Nam, Joo Ho Lee, Hongzhi Wu, Diego Gutierrez, and Min H. Kim. Simultaneous acquisition of microscale reflectance and normals. *ACM Trans. Graph.*, 35(6), 2016.
- [165] Geoffrey RS Naylor, Margaret Pate, and Graham J Higgerson. Determination of cotton fiber maturity and linear density (fineness) by examination of fiber cross-sections. part 2: A comparison optical and scanning electron microscopy. *Textile Research Journal*, 84(18), 2014.
- [166] Howard L Needles. *Textile fibers, dyes, finishes, and processes*. 2001.
- [167] Birgit Nitzsche and Karsten Rose. *Bodyshop: The Photoshop Retouching Guide for the Face and Body*. 2011.
- [168] James J Nordlund, Raymond E Boissy, Vincent J Hearing, Richard A King, and JP Ortonne. *The pigmentary system*. Oxford University Press New York NY, 1998.
- [169] Shinji Ogaki, Yusuke Tokuyoshi, and Sebastian Schoellhammer. An empirical fur shader. In *ACM SIGGRAPH ASIA 2010 Sketches*, 2010.
- [170] Michael Oren and Shree K Nayar. Generalization of Lambert's reflectance model. In *Proceedings of SIGGRAPH*, 1994.
- [171] Carol O'Sullivan and John Dingliana. Collisions and perception. *ACM Trans. Graph.*, 20(3), 2001.
- [172] Carol O'Sullivan, Ralph Radach, and Steven Collins. A model of collision perception for real-time animation. In *Proc. Eurographics Workshop on Comp. Anim. and Sim.*, 1999.
- [173] Carol O'Sullivan, John Dingliana, Thanh Giang, and Mary K Kaiser. Evaluating the visual fidelity of physically based animations. *ACM Trans. Graph.*, 22(3), 2003.
- [174] JO Otutu, EM Efurhievwe, SU Ameru, EK Ossai, et al. Synthesis and application of hetaryl disazo disperse dyes derived from 2-amino-5-mercapto-1, 3, 4-thiadiazole and 3-chloroaniline on synthetic polymer-fibres. *Chemistry and Materials Research*, 6(4), 2014.
- [175] Thomas D Parsons and Albert A Rizzo. Affective outcomes of virtual reality exposure therapy for anxiety and specific phobias: A meta-analysis. *Journal of behavior therapy and experimental psychiatry*, 39(3), 2008.
- [176] Fabio Pellacini, James A Ferwerda, and Donald P Greenberg. Toward a psychophysically-based light reflection model for image synthesis. In *Proc. SIGGRAPH'00*, 2000.

- [177] Bo Peng, Tianhuai Ding, and Peng Wang. Propagation of polarized light through textile material. *Applied optics*, 51(26), 2012.
- [178] Rudolf Penndorf. Angular Mie scattering. *JOSA*, 52(4), 1962.
- [179] Matt Pharr and Greg Humphreys. *Physically Based Rendering, Second Edition: From Theory To Implementation*. 2nd edition, 2010.
- [180] Peter E Pochi, John S Strauss, and Donald T Downing. Age-related changes in sebaceous gland activity. *Journal of Investigative Dermatology*, 73(1), 1979.
- [181] Robin M Pope and Edward S Fry. Absorption spectrum (380–700 nm) of pure water. ii. integrating cavity measurements. *Applied optics*, 36(33), 1997.
- [182] Ganesh Ramanarayanan, James Ferwerda, Bruce Walter, and Kavita Bala. Visual equivalence: towards a new standard for image fidelity. *ACM Trans. Graph.*, 26(3), 2007.
- [183] Ganesh Ramanarayanan, Kavita Bala, and James A Ferwerda. Perception of complex aggregates. *ACM Trans. Graph.*, 27(3), 2008.
- [184] Freddy Alberto Monroy Ramírez and Jorge Garcia-Sucerquia. Optical fibre characterization through digital holographic microscopy. In *Biomedical Optics*, page JMA18. Optical Society of America, 2010.
- [185] Paul SA Reitsma and Nancy S Pollard. Perceptual metrics for character animation: sensitivity to errors in ballistic motion. *ACM Trans. Graph.*, 22(3), 2003.
- [186] Roland Ruiters and Reinhard Klein. Btf compression via sparse tensor decomposition. *Computer Graphics Forum*, 28(4), 2009.
- [187] Iman Sadeghi, Heather Pritchett, Henrik Wann Jensen, and Rasmus Tamstorf. An artist friendly hair shading system. *ACM Trans. Graph.*, 29(4), 2010.
- [188] Iman Sadeghi, Heather Pritchett, Henrik Wann Jensen, and Rasmus Tamstorf. An artist friendly hair shading system. *ACM Trans. Graph.*, 29(4), 2010.
- [189] Iman Sadeghi, Adolfo Munoz, Philip Laven, Wojciech Jarosz, Francisco Seron, Diego Gutierrez, and Henrik Wann Jensen. Physically-based simulation of rainbows. *ACM Trans. Graph.*, 31(1), 2012.
- [190] Iman Sadeghi, Oleg Bisker, Joachim de Deken, and Henrik Wann Jensen. A practical microcylinder appearance model for cloth rendering. *ACM Trans. Graph.*, 32(2), 2013.

- [191] Iman Sadeghi, Oleg Bisker, Joachim de Deken, and Henrik Wann Jensen. A practical microcylinder appearance model for cloth rendering. *ACM Trans. Graph.*, 32(2), 2013.
- [192] Iyad Salam Saidi. *Transcutaneous optical measurement of hyperbilirubinemia in neonates*. PhD thesis, Rice, University, 1992.
- [193] Elena Salomatina, Brian Jiang, John Novak, and Anna N Yaroslavsky. Optical properties of normal and cancerous human skin in the visible and near-infrared spectral range. *Journal of Biomedical Optics*, 11(6), 2006.
- [194] Mirko Sattler, Ralf Sarlette, and Reinhard Klein. Efficient and realistic visualization of cloth. In *Proceedings of EGSR*, 2003.
- [195] A.P. Saygin, T. Chaminade, H. Ishiguro, J. Driver, and C. Frith. The thing that should not be: Predictive coding and the Uncanny Valley in perceiving human and humanoid robot actions. *Social Cognitive Affective Neuroscience*, 7(4), 2012.
- [196] Kristina Scherbaum, Tobias Ritschel, Matthias Hullin, Thorsten Thormählen, Volker Blanz, and Hans-Peter Seidel. Computer-suggested facial makeup. In *Computer Graphics Forum*, volume 30, 2011.
- [197] Kristina Scherbaum, James Petterson, Rogerio S Feris, Volker Blanz, and Hans-Peter Seidel. Fast face detector training using tailored views. In *Proceedings of the IEEE International Conference on Computer Vision*, 2013.
- [198] Edward Schneider, Yifan Wang, and Shanshan Yang. Exploring the Uncanny Valley with japanese video game characters. In *Situated Play: Proceedings of the Digital Games Research Association Conference*, 2007.
- [199] Kai Schröder, Reinhard Klein, and Arno Zinke. A volumetric approach to predictive rendering of fabrics. *Computer Graphics Forum*, 30(4), 2011.
- [200] Kai Schröder, Reinhard Klein, and Arno Zinke. A volumetric approach to predictive rendering of fabrics. *Computer Graphics Forum (Proceedings of EGSR 2011)*, 30(4):1277–1286, 2011.
- [201] Kai Schröder, Shuang Zhao, and Arno Zinke. Recent advances in physically-based appearance modeling of cloth. In *SIGGRAPH Asia 2012 Courses*, 2012.
- [202] Kai Schröder, Arno Zinke, and Reinhard Klein. Image-based reverse engineering and visual prototyping of woven cloth. *IEEE Transactions on Visualization and Computer Graphics*, PP(99), 2014.

- [203] Kai Schröder, Arno Zinke, and Reinhard Klein. Image-based reverse engineering and visual prototyping of woven cloth. *IEEE Trans. Vis. Comp. Graph.*, 21(2), 2015.
- [204] Ana Serrano, Diego Gutierrez, Karol Myszkowski, Hans-Peter Seidel, and Belen Masia. An intuitive control space for material appearance. *ACM Transactions on Graphics (SIGGRAPH ASIA 2016)*, 35(6), 2016.
- [205] Jun'ichiro Seyama and Ruth S. Nagayama. The Uncanny Valley: Effect of realism on the impression of artificial human faces. *Presence: Teleoperators and Virtual Environments*, 16(4), 2007.
- [206] Neal E Seymour, Anthony G Gallagher, Sanziana A Roman, Michael K O'brien, Vipin K Bansal, Dana K Andersen, and Richard M Satava. Virtual reality training improves operating room performance: results of a randomized, double-blinded study. *Annals of surgery*, 236(4), 2002.
- [207] Darshil U. Shah, David Porter, and Fritz Vollrath. Opportunities for silk textiles in reinforced biocomposites: Studying through-thickness compaction behaviour. *Composites Part A: Applied Science and Manufacturing*, 62, 2014.
- [208] Leonid Sigal, Moshe Mahler, Spencer Diaz, Kyna McIntosh, Elizabeth Carter, Timothy Richards, and Jessica Hodgins. A perceptual control space for garment simulation. *ACM Transactions on Graphics (TOG)*, 34(4), 2015.
- [209] Marc Spicker, Diana Arellano, Ulrich Schaller, Reinhold Rauh, Volker Helzle, and Oliver Deussen. Emotion recognition in autism spectrum disorder: does stylization help? In *Proceedings of the ACM Symposium on Applied Perception*, 2016.
- [210] Jos Stam. An illumination model for a skin layer bounded by rough surfaces. In *Rendering Techniques 2001*. 2001.
- [211] William A Stokes, James A Ferwerda, Bruce Walter, and Donald P Greenberg. Perceptual illumination components: a new approach to efficient, high quality global illumination rendering. *ACM Trans. Graph.*, 23(3), 2004.
- [212] JA Thomasson, S Manickavasagam, and MP Mengüç. Cotton fiber quality characterization with light scattering and fourier transform infrared techniques. *Applied spectroscopy*, 63(3), 2009.
- [213] Lihui Tian and Shuangjiu Xiao. Facial feature exaggeration according to social psychology of face perception. In *Computer Graphics Forum*, volume 35, 2016.

- [214] A. Tinwell and M. Grimshaw. Bridging the uncanny: An impossible traverse? In *Proceedings of the 13th International MindTrek Conference: Everyday Life in the Ubiquitous Era*, 2009.
- [215] Angela Tinwell. *The Uncanny Valley in Games and Animation*. 2014.
- [216] Angela Tinwell, Mark Grimshaw, Debbie Abdel Nabi, and Andrew Williams. Facial expression of emotion and perception of the Uncanny Valley in virtual characters. *Computers in Human Behavior*, 27(2), 2011.
- [217] Angela Tinwell, Mark Grimshaw, and Andrew Williams. The Uncanny Wall. *International Journal of Arts and Technology*, 4(3), 2011.
- [218] Kenneth E Torrance and Ephraim M Sparrow. Theory for off-specular reflection from roughened surfaces. *JOSA*, 57(9), 1967.
- [219] Edward R Trotman. *Dyeing and chemical technology of textile fibres*. 1984.
- [220] Norimichi Tsumura, Hideaki Haneishi, and Yoichi Miyake. Independent-component analysis of skin color image. *JOSA A*, 16(9), 1999.
- [221] Norimichi Tsumura, Nobutoshi Ojima, Kayoko Sato, Mitsuhiro Shiraishi, Hideto Shimizu, Hirohide Nabeshima, Syuuichi Akazaki, Kimihiko Hori, and Yoichi Miyake. Image-based skin color and texture analysis/synthesis by extracting hemoglobin and melanin information in the skin. *ACM Transactions on Graphics*, 22(3), 2003.
- [222] Valeriĭ Viktorovich Tuchin. *Tissue optics: light scattering methods and instruments for medical diagnosis*, volume 13. SPIE press Bellingham, 2007.
- [223] Peter Vangorp, Jurgen Laurijssen, and Philip Dutré. The influence of shape on the perception of material reflectance. *ACM Trans. Graph.*, 26(3):77, 2007.
- [224] Michele Vicovaro, Ludovic Hoyet, Luigi Burigana, and Carol O'Sullivan. Evaluating the plausibility of edited throwing animations. In *Proc. ACM SIGGRAPH/Eurographics Symp. on Comp. Anim.*, 2012.
- [225] Pascal Volino, Frederic Cordier, and Nadia Magnenat-Thalmann. From early virtual garment simulation to interactive fashion design. *Computer-aided design*, 37(6), 2005.

- [226] Jeanette M Waller and Howard I Maibach. Age and skin structure and function, a quantitative approach (i): blood flow, ph, thickness, and ultrasound echogenicity. *Skin Research and Technology*, 11(4), 2005.
- [227] Christian Wallraven, Heinrich H. Bühlhoff, Douglas W. Cunningham, Jan Fischer, and Dirk Bartz. Evaluation of real-world and computer-generated stylized facial expressions. *ACM Transactions on Applied Perception*, 4(3), 2007.
- [228] Christian Wallraven, Martin Breidt, Douglas W. Cunningham, and Heinrich H. Bühlhoff. Evaluating the perceptual realism of animated facial expressions. *ACM Transactions on Applied Perception*, 4(4), 2008.
- [229] Lihong Wang, Steven L Jacques, and Liqiong Zheng. Monte carlo modeling of light transport in multi-layered tissues. *Computer methods and programs in biomedicine*, 47(2), 1995.
- [230] Andrea Weidlich and Alexander Wilkie. Realistic rendering of birefringency in uniaxial crystals. *ACM Trans. Graph.*, 27(1), 2008.
- [231] Thibaut Weise, Sofien Bouaziz, Hao Li, and Mark Pauly. Real-time performance-based facial animation. *ACM Transactions on Graphics*, 30(4), 2011.
- [232] Harold B Westlund and Gary W Meyer. Applying appearance standards to light reflection models. In *Proc. SIGGRAPH'01*, 2001.
- [233] Tim Weyrich, Wojciech Matusik, Hanspeter Pfister, Bernd Bickel, Craig Donner, Chien Tu, Janet McAndless, Jinho Lee, Addy Ngan, Henrik Wann Jensen, et al. Analysis of human faces using a measurement-based skin reflectance model. *ACM Transactions on Graphics*, 25(3), 2006.
- [234] Judi T Whitton and JD Everall. The thickness of the epidermis. *British Journal of dermatology*, 89(5), 1973.
- [235] Brenda K Wiederhold and Mark D Wiederhold. *Virtual reality therapy for anxiety disorders: Advances in evaluation and treatment*. 2005.
- [236] Josh Wills, Sameer Agarwal, David Kriegman, and Serge Belongie. Toward a perceptual space for gloss. *ACM Trans. Graph.*, 28(4), 2009.
- [237] Kui Wu and Cem Yuksel. Real-time fiber-level cloth rendering. In *Proceedings of I3D*, 2017.

- [238] Ying-Qing Xu, Yanyun Chen, Stephen Lin, Hua Zhong, Enhua Wu, Baining Guo, and Heung-Yeung Shum. Photorealistic rendering of knitwear using the lumislice. In *Proceedings of SIGGRAPH*, 2001.
- [239] Jun Yamada. Radiative properties of fibers with non-circular cross sectional shapes. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 73(2):261–272, 2002.
- [240] Jun Yamada and Yasuo Kurosaki. Radiative characteristics of fibers with a large size parameter. *International journal of heat and mass transfer*, 43(6):981–991, 2000.
- [241] Jun Yamada, Yasuo Kurosaki, and Masaaki Take-Uchi. Radiation transfer in a fibrous medium with fiber orientation. In *28th National Heat Transfer Conference and Exhibition*, 1992.
- [242] Ling-Qi Yan, Chi-Wei Tseng, Henrik Wann Jensen, and Ravi Ramamoorthi. Physically-accurate fur reflectance: modeling, measurement and rendering. *ACM Trans. Graph.*, 34(6), 2015.
- [243] Thomas Y Yeh, Glenn Reinman, Sanjay J Patel, and Petros Faloutsos. Fool me twice: Exploring and exploiting error tolerance in physics-based animation. *ACM Trans. Graph.*, 29(1), 2009.
- [244] Kenichiro Yoshida, Masahiro Miyaki, Nobutoshi Ojima, and Kayoko Iwata. Relationship between microstructure of the skin surface and surface reflection based on geometric optics. *Journal of Dermatological Science*, 66(3), 2012.
- [245] Eduard Zell and Mario Botsch. ElastiFace: Matching and blending textured faces. In *Proceedings of the Symposium on Non-Photorealistic Animation and Rendering*, 2013.
- [246] Eduard Zell, Carlos Aliaga, Adrian Jarabo, Katja Zibrek, Diego Gutierrez, Rachel McDonnell, and Mario Botsch. To stylize or not to stylize?: The effect of shape and material stylization on the perception of computer-generated faces. *ACM Trans. Graph.*, 34(6).
- [247] Shuang Zhao, Wenzel Jakob, Steve Marschner, and Kavita Bala. Building volumetric appearance models of fabric using micro ct imaging. *ACM Trans. Graph.*, 30(4), 2011.
- [248] Shuang Zhao, Wenzel Jakob, Steve Marschner, and Kavita Bala. Building volumetric appearance models of fabric using micro ct imaging. *ACM Trans. Graph.*, 30(4), 2011.
- [249] Shuang Zhao, Wenzel Jakob, Steve Marschner, and Kavita Bala. Structure-aware synthesis for predictive woven fabric appearance. *ACM Trans. Graph.*, 31(4), 2012.

- [250] Shuang Zhao, Milos Hasan, Ravi Ramamoorthi, and Kavita Bala. Modular flux transfer: efficient rendering of high-resolution volumes with repeated structures. *ACM Trans. Graph.*, 32(4), 2013.
- [251] Shuang Zhao, Fujun Luan, and Kavita Bala. Fitting procedural yarn models for realistic cloth rendering. *ACM Trans. Graph.*, 35(4), 2016.
- [252] Shuang Zhao, Fujun Luan, and Kavita Bala. Fitting procedural yarn models for realistic cloth rendering. *ACM Trans. Graph.*, 35(4), 2016.
- [253] Jun-Yan Zhu, Aseem Agarwala, Alexei A Efros, Eli Shechtman, and Jue Wang. Mirror mirror: Crowdsourcing better portraits. *ACM Transactions on Graphics*, 33(6), 2014.
- [254] Katja Zibrek and Rachel McDonnell. Does render style affect perception of personality in virtual humans? In *Proceedings of the ACM Symposium on Applied Perception*, 2014.
- [255] Arno Zinke and Andreas Weber. Light scattering from filaments. *IEEE Trans. Vis. Comp. Graph.*, 13(2), 2007.
- [256] Arno Zinke and Andreas Weber. Light scattering from filaments. *Visualization and Computer Graphics, IEEE Transactions on*, 13(2):342–356, 2007.
- [257] George Zonios and Aikaterini Dimou. Light scattering spectroscopy of human skin in vivo. *Optics Express*, 17(3), 2009.