

The Racial Character of Computer Graphics Research

THEODORE KIM, ALEXA SCHOR, JULIAN POSADA, and ALKA V. MENON, Yale University, USA

Computer graphics algorithms for generating photorealistic imagery are widely perceived to be universal, and capable of conjuring anything that a filmmaker or game designer can imagine. However, recent works have suggested that 3D algorithms for depicting synthetic humans are far from generic, and instead favor historically hegemonic characteristics. We present the first systematic review of human depiction in the top computer graphics conference and the journal of record (*SIGGRAPH* and *ACM Transactions on Graphics*) that confirms previous hypotheses. Algorithms that claim to be generically rendering “human skin” are in fact imagined and formulated for translucent, “high albedo” materials such as white skin. Algorithms claiming to apply generically to “human hair” are formulated for “rods,” “wires” and “threads” which are analogous to straight hair. Our analysis reveals conceptual binarization, where algorithms for white skin are treated as computational substrate for “all” skin, imposing a hierarchical assumption that all skin descends from the math and physics of white skin. Hair algorithms follow a similar historical pattern, with the first examples of computer-generated Type 4 hair only appearing after the murder of George Floyd in 2020. We offer a new conceptual label, *McDaniels Methods*, for characterizing and critiquing computer graphics algorithms that reinforce racial hierarchy under a false cover of diversity. We also offer an inverse label, *Durald Methods*, for algorithms that were closely co-designed with the people being depicted. Our analysis points the way towards several neglected avenues for future research.

ACM Reference Format:

Theodore Kim, Alexa Schor, Julian Posada, and Alka V. Menon. 2026. The Racial Character of Computer Graphics Research. In *The 2026 ACM Conference on Fairness, Accountability, and Transparency (FAccT '26)*, June 25–28, 2026, Montreal, QC, Canada. ACM, New York, NY, USA, 33 pages. <https://doi.org/10.1145/3805689.3812275>

1 Introduction

Modern algorithms for computer generated imagery (CGI) are widely considered to be universal methods that can depict anything in the physical world. Research into CGI has been ongoing for the last 50 years, with a broad constellation of algorithms developed to depict a wide variety of natural phenomena, such as plants [224], fire [249], water [315], muscles [289] and cloth [16]. Today, the algorithms regularly appear in any form of image-based media, including YouTube videos, film and games. The techniques are sufficiently mature that one VFX practitioner has claimed that “[t]here’s a general consensus within visual effects that more or less anything can now be simulated realistically” [333]. Epic Games has publicly declared [204] that the latest version of Unreal Engine overcomes the uncanny valley [239].

However, claims to universalism and neutrality can often mask systematic bias. For example, artificial intelligence systems have been described as “racialized and gendered models of the self that are falsely presented as universal” [173] that speak “to some universal yet never articulated subject when it is obvious that the authors of such systems regard themselves as the gold standard of universal subjects.” [1] Cryptographic research has seen similar criticisms for “universalist standard[s] [that] in fact implies a ‘normal’ user who conforms to normative standards of race (White), class (wealthy), gender (cisgendered male), ability (able-bodied).” *The Gender Shades*

Authors’ Contact Information: Theodore Kim, theodore.kim@yale.edu; Alexa Schor, alexa.schor@yale.edu; Julian Posada, julian.posada@yale.edu; Alka V. Menon, alka.menon@yale.edu, Yale University, New Haven, CT, USA.



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License.

FAccT '26, Montreal, QC, Canada

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2596-8/2026/06

<https://doi.org/10.1145/3805689.3812275>

study [47] punctured such universalist claims by exposing patterns of gender and racial discrimination in supposedly neutral machine learning algorithms. Gender issues have also been identified in the neutrally-coded automatic gender recognition algorithms utilized in HCI [176]. Troubling racial patterns have been found in seemingly unbiased restaurant review datasets [213]. Rather than being neutral, value systems are implicitly encoded in both datasets and research practices [24, 42, 68, 230, 251, 263]. For example, a field-level audit of [38] machine learning research found that Performance, Efficiency, and Novelty are vaunted qualities, while harms are rarely discussed. A similar audit [172] of the computer vision literature found an increasing alignment with surveillance applications, with connections simultaneously being obfuscated.

Similar problems have been found with the claim that CGI algorithms can evenhandedly depict all humans [179, 181]. Key algorithms that are generically billed as capturing the appearance of “skin” and “human skin” are often only demonstrated on white skin, and the underlying physics being targeted is pale, translucent materials. Algorithms that generically advertise their suitability for depicting the appearance and motion of “hair” are often only shown working on straight hair, and investigate the optics and mechanics of infinitely straight rods. These observations are echoed in STS scholarship that observes that, contrary to any claims about universality, CGI algorithms tend to focus on a “small subset of parts in any given world” [113].

Attempts to create CGI humans that deviate from these implicit norms have resulted in Black artists and writers criticizing the poor quality of dark skin and textured hair in video games. [72] Narcisse [248] has pointed out that this is partially a problem of representation, i.e. video game artists with insufficient backgrounds inevitably fail to make authentic-looking hair. However, the biases baked into the underlying CGI techniques also play a significant role. The algorithms were made to depict one kind of human, so creating other types that lie outside the preferred baseline becomes progressively more difficult.

We present the first field-level audit of human depiction in computer graphics research papers. We focus on the conferences of record in the field of computer graphics, *SIGGRAPH* (1974-present) and *SIGGRAPH Asia* (2008-present), and the journal of record, *ACM Transactions on Graphics (TOG)* (1982-present). The venues are inextricably linked: for much of their existence, the *SIGGRAPH* proceedings were published as a special issue of *TOG*. Through an agreement with our university library, we have obtained the full text of every article through the end of 2024. The size of this corpus is significant (6864 papers), so we perform a *thematic analysis* of this body of literature in order to build a roadmap for future research. We tag and categorize every instance of the term “skin” and “hair”, identify patterns, and analyze the images and language that are deployed around these terms. A large proportion of the early papers were written by North American and European authors (Figure 1, right), so much of what we see is WEIRD (Western, Educated, Industrialized, Rich, and Democratic) norms [144] consistent with the cultural practices of the U.S. film and video game industries.

Previous works on this topic (some written by authors of this paper) have largely been short form, either speculating [179] on the existence of racial patterns in a 1000-word op-ed, or presenting pilot analyses [181] in a 2-page extended abstract. We instead present the first detailed, literature-level study that both confirms and extends the hypotheses from these previous works. In addition to white skin being presented as false synecdoche for all skin, we find a community-level binarization where skin is conceptualized as a “white” / “non-white” dyad. Consistent with similar feminist analyses [186], the binary encodes a hierarchy that imagines the “Black body as a derivation from the white body” [219]. We find a similar dyad and hierarchy of “straight” / “non-straight” in the algorithmic conception of hair. We observe an uptick in the instances of Black skin, as well as the *first* instance of Black hair, coincident with the aftermath of the 2020 murder of George Floyd. However, these instances often appear as low-quality addenda to papers that otherwise showcase white skin and straight hair. This practice mirrors early analog cinematography [6], where Black faces were only shown when necessary to suggest that white skin lighting norms were universal.

We propose the term *McDaniels Methods* to identify algorithms that continue this practice of reinforcing racial hierarchy under a false cover of diversity. We define *McDaniels Methods* as *techniques that are exhaustively*




Fig. 1. **Left:** Percentage of total papers coming from industrial labs in *SIGGRAPH* and *TOG*. Film studios like Lucasfilm and Pixar had early influence, as well as Disney in the 2000s and 2010s. **Right:** Geographic location of authors over time. North America dominated early on, and majority currently skews WEIRD (Western, Educated, Industrialized, Rich, and Democratic). [144] The data source for these figures is described in Appendix A.1.

validated on white skin and straight hair, and relegate all other forms of skin and hair to a small set of flawed and underwhelming examples. We choose the name “McDaniels” because the one Black face in the influential 1949 cinematography manual *Painting With Light* [6] is that of Hattie McDaniel, the first Black actor to ever win an Academy Award in 1939. Reflecting the carelessness her depiction was given at the time, the book misspells her name. The caption lists her as *Hattie McDaniels*, but her last name has no trailing “s”. We call such techniques *McDaniels Methods* to underscore how the racial practices of the analog era carried forward into the digital.

We propose the contrasting term *Durald Methods*, which we define as *techniques that were co-created in close collaboration with the people being depicted*. We name the term after Autumn Durald Arkapaw, the first woman of color to ever win the Academy Award in Best Cinematography for the 2025 film *Sinners*.

2 Background

Race is a sociopolitical construct with many dimensions. From self-identification, to phenotype, to evaluations of racial group membership by others on the basis of appearance to ancestry or genetics, the word “race” is colloquially used to cover all these different dimensions [284]. In this paper, when we discuss the “racial character” of computer graphics, we draw on the dimension of race as represented through phenotype, including color and physical features like hair texture. Phenotype encompasses the visual dimension of race, which is key for representational contexts. Its mapping onto racial categories is historically and culturally contingent, and central to the construction and maintenance of racial hierarchies in the U.S. [136, 254]. This then echoes in international contexts and online [247] because of the U.S.’s economic and cultural power [229]. Attempts to represent race reveal the assumptions a society makes about bodies, values, and universality. Fields from cosmetics to medicine have sought to render the social phenomenon of race exclusively in technical and physical terms, under the guise of colorblindness [41]. The principle of colorblindness “pretends that racial recognition rather than racist rule is the problem to be solved” [69] and suppresses inquiry by “equat[ing] noticing racial inequality with *perpetuating* racial inequality” [276].

In the absence of such inquiries, a norm of whiteness becomes the default, the unmarked category [228, 297]. This stems from the enduring legacy U.S. and European domination, where white people topped the racial hierarchy [240] and came to stand in for the universal human [99]. At the same time, scholars have pointed to a preference for lightness in Africa and Asia that can overlap with and reinforce norms of whiteness, even as they arise from locally specific social hierarchies [124, 312]. Thus, representation of race, even without any explicit references, remakes and produces racial difference even as it purports to simply describe it.

This is not a new phenomenon, as whiteness as the default in imaging technologies goes back to the analog era. Roth [283] investigated its role as a baseline in analog color photography, and Yue [364] examined how it

intersected with gender when looking at “Leader Ladies” or “China Girls” as development targets in analog color motion pictures. Dyer [96] traces the provenance of these norms back through silent film era to the Dutch Old Masters. Extending these analyses, our work shows how the same racial standards and practices carried over from the analog into the digital era.

This research advances the wider critique of colorblind racism and the universalism applied to computing technology. Our analysis is notably informed by the “New Jim Code” of Benjamin [24], which critiques technological deployments portrayed as objective or progressive that actually reproduce existing racial inequalities. This resonates with the work of Noble [251] on algorithmic oppression, which challenges the perceived neutrality of computer-mediated platforms that allow structural inequality to propagate under a colorblind veneer. These authors extend the critique of colorblind racism leveled by Bonilla-Silva [41], specifically the naturalization of racial outcomes and the abstract liberalism used to justify market dynamics, which ultimately renders structural inequalities invisible. Such masking of social differentiation has been a persistent feature of computing technology for decades, as scholars continue to document how the “myth of neutrality” sustains systemic bias.

In this context, many have examined the values encoded in the research practices of different sub-disciplines in computer science. Cryptography [282], machine learning [38], computer vision [172], and human-computer interaction [94] have all seen thematic analyses of what is implicitly valued and devalued in their respective literatures. Rather than reflecting neutral, disinterested research that is taking place from “a view from nowhere” [1], many of these sub-disciplines are repeating and re-entrenching the biases and power structures of the past. For example, some machine learning datasets have been found to be overwhelmingly light-skinned [47], which produces facial analysis algorithms that perform poorly on darker-skinned females, and has led to a variety of works on bias in visual datasets [101].

While thematic analyses have been carried out most prominently on machine-learning datasets and their outputs, far less attention has been paid to computer graphics algorithms. Such analyses are urgent because CGI algorithms are increasingly being used to generate “synthetic” datasets, for example, for machine learning. Notable examples include driving simulation [93], human faces [341], and full human bodies [321]. In this sense, our analysis lies one step upstream: any biases present in the underlying algorithmic substrate will be passed on through the data.

Science and Technology Studies (STS) scholars have also begun examining the bias in CGI algorithms. Building on earlier analyses [179], Malazita [219] examines “white photorealism” in the Unreal Engine, as Gass [119] has similarly shown in the 2018 game *Detroit: Becoming Human*. Other work has shown how a normative white face is the baseline for *FaceGen*, the character creation tool used in *Fallout 3*. Phillips [267] and Ketchum [175] have shown how CGI algorithms are non-universal, and how reality is “unevenly distributed” in synthetic images [113]. STS scholars [89, 218] have collaborated with computer graphics researchers on extended abstracts at *SIGGRAPH* that bring more visibility to these issues. None of this previous work has attempted a field-level audit of representation. Thus, our work both provides empirical validation and extension to this prior scholarship.

Conference and Journal Choice. We selected *SIGGRAPH* and *TOG* for analysis because they are widely acknowledged as the leading venues in computer graphics. The annual *SIGGRAPH* conference routinely hosts over 10,000 attendees [342], including leading researchers from both industry and academia, and has historically been where major breakthroughs appear. The next largest *Eurographics* conference is over an order of magnitude smaller at roughly 500 annual attendees [100]. *TOG* is consistently the top-rated computer graphics journal, as evidenced by 2024 Google Scholar Metrics (h-index 114), Journal Citation Reports (Impact Factor 9.5), and Eigenfactor (0.031). The next highest journals trail significantly: *Computer Graphics Forum* (h-index 64, IF 2.9, EF 0.023) and *IEEE Transactions on Visualization and Computer Graphics* (h-index 94, IF 6.5, EF 0.017). The latter’s citations mostly come from data visualization, not computer graphics. While a complete analysis of all computer graphics

papers would undoubtedly yield valuable insights, any attempt to understand what is considered *top-quality, high-prestige* research in computer graphics must start by examining the papers in *SIGGRAPH* and *TOG*.

Skin Classification Systems. Numerous schemes have been developed, the most prevalent being the Fitzpatrick scale [108] that classifies skin based on how quickly it burns in the sun. Fitzpatrick [109] disputed this application of his system, because it had been developed “to classify persons *with white skin*” (emphasis his). The von Luschan [325] scale was developed for racial classification, but was “swiftly abandoned when reflectance spectrophotometry was introduced in the early 1950s” [163]. Its role in Nazi eugenics [242] probably contributed to this abandonment. The recent Monk [234] scale adopted by Google and Meta is an effort to decouple “tone” from “race.”

Classifying skin is inherently fraught. The concept of “white” varies widely by geographic region and political context because “whiteness can only be defined by state power” [67]. That is, country-specific laws, like the U.S. justice system, “adjudicate the arbitrary inclusion and exclusion of people across time,” [67] such as with Irish American immigrants in the 19th century [162]. After extensive discussion of the constraints of existing systems, the authors decided on the broad classifications of “white” and “non-white.” These labels are designed to facilitate the overarching purpose of our study: a conversation about what has remained implicit in the field.

Hair Classification. Many hair typing systems also exist with various drawbacks. For example, the popular Walker [326] system (Types 1-4) is a qualitative measure developed for haircare products, and does not correspond to physical measurements. The system has been criticized for excluding curlier hair types [71], leading to the informal addition of Types 3c and 4c. L’oreal developed a quantitative system [77] that measures the “number of waves” when a hair is pressed between two plates of glass, a process that breaks down with helical hairs.

We adopted a qualitative approach to hair types that is a reduced version of the Walker [326] system:

- *Straight and wavy hair*, containing by virtually no waves, or roughly sinusoidal waves along one dimension. This hair does *not* contain full helical twists.
- *Curly hair*, containing full helical twists, but with a sufficiently large radii that hair is straight near the scalp.
- *Textured hair*, which contains twists with millimeter-scale radii, and forms a spongy layer near the scalp.

In the Walker system, straight (Type 1) and wavy (Type 2) hair are separate categories. We combined them because it was often unclear if an undulation was inherent to the hair or originated from styling, e.g. because the hair had been lightly mussed. The co-authors with the computer graphics background performed the initial classifications of skin and hair in the articles, which were spot checked by the sociology authors.

3 Methods

Skin Papers. We initially searched for all *SIGGRAPH* and *TOG* papers containing the word “skin” that were also tagged under the ACM Computing Classification System (CCS) as **Computing Methodologies** → **Computer Graphics** → **Rendering**. Since *SIGGRAPH* began in 1974 but the current CCS system was instituted in 2012, the tags were unreliable. For example, one influential paper [167] was incorrectly tagged as “shape modeling” and “computational geometry”.

Instead, we searched the corpus with the query `skin AND (render OR rendering OR shading)` to obtain a superset of the relevant papers. The query yielded 466 papers. We manually inspected the results, and if the term “skin” referenced a photorealistic rendering in the paper, it was included. Second, if the term “skin” was used in the text to refer to the physics of skin interacting with light, the paper was included. 120 papers remained after this culling process. Table 1 gives a breakdown of the excluded papers.

The authors coded skin color as white/non-white in the images appearing in the selected papers. We chose not to outsource the coding process to crowd-based “data workers” because prior research has documented concerns with such practices, including the risk of client-induced bias in annotated outputs [230] and limited transparency

Number	Category
146	No image of a human or substantive reference to skin, e.g. “skin” only appeared in a citation.
39	Image and video editing papers. These 2D algorithms are distinct from 3D rendering methods.
37	Papers spanning 2D and 3D, e.g. using 2D images to construct a 3D avatar, remained included. Incidental image of humans, e.g. a geometry processing paper containing a human head model [307]
16	Hair papers, moved to our hair analysis.
21	Geometric skinning, elasticity, and motion capture papers on skin <i>animation</i> , not rendering.
17	Eyelid [27], and eyebrows [198], teeth, tongue, and lips papers.
14	Full-body capture papers, as issues of body normativity are out-of-scope.
13	Clothing capture and simulation papers
20	Anatomy modeling, e.g. muscles [290], diaphragm [88], and hands.
13	Illustration papers, e.g. portrait painting tools [134] or anime portrait generators [202].
5	Medical imaging papers, e.g. viewing organ shapes in MRI data.
5	Haptic rendering [210] and crowd simulation papers

Table 1. Papers excluded from our analysis from the query skin AND (render OR rendering OR shading).

surrounding the labor conditions and demographic composition of workers on major platforms [243]. Instead, two authors who are North American computer graphics researchers performed initial classifications, which were then spot checked by the sociology authors. The labels from the two researchers only disagreed on 3 papers [36, 250, 363] out of 120. The disagreements were resolved through discussion, detailed in Appendix B.1. Since the field has been long-dominated by North American scientists (Figure 1), collecting labels from researchers embedded in that population provides the most direct window into the field’s perceptions. The labels are not authoritative but subjective, and rooted in the researchers’ commonsense understanding.

Hair Papers. We performed the query hair AND (render OR rendering OR shading OR modeling OR modelling OR simulation OR animation), yielding 570 papers. We manually inspected the matches, and if the term “hair” appeared in the context of a photorealistic image, such as an attempt to achieve realistic appearance, shape, or motion, the paper was included. If no image appeared, but ‘hair’ was in reference to the optics or mechanics of hair, it was included. 160 papers remained after this culling process. Table 2 gives a breakdown of the excluded papers. Two graphics researchers coded hair as straight/wavy, curly, or textured, which were then checked by the sociology authors. Only 3 papers of 160 [143, 165, 332] saw disagreement, and were resolved through discussions described in Appendix B.2.

Limitations. Other computer graphics conferences include *Eurographics*, *Pacific Graphics*, and *Afrigraph* respectively based in Europe, Asia, and Africa. Preliminary examinations of these corpora did not significantly shift our findings. Strikingly, we did *not* find a sudden wealth of Black skin and textured hair papers within *Afrigraph*. Instead, both the African and European conferences contained the same patterns of showcasing white skin [84, 302] and straight hair [157, 261], and we were able to locate only two instances of textured hair [31, 262]. While the European (1980 - present) and Asian (1993 - present) conferences are long-running, the African conference was short-lived (2001 - 2010).

Our choice of classifications limits the patterns we can identify. Hair color was not categorized, and skin tone was not evaluated on a continuum. Elective modifications such as the straightened textured hair [148], lightened skin [313], dyed hair, or tanned skin were not tagged. Potential trends such as East Asian skin representations becoming lighter (or darker) over time, Black subjects being shown predominantly with straightened hair, or

Number	Category
130	No image of hair or substantive reference to hair, e.g. “hair” only appeared in a citation.
67	Image and video editing papers. These 2D algorithms are distinct from 3D rendering methods. Papers spanning 2D and 3D remained included.
51	Illustration and stylization papers, e.g. watercolors [65], pen and ink [291], and manga [350].
63	Face-only papers, that capture and render faces, with hair excluded [43, 128, 345].
42	Full-body and clothing capture papers where hair detail is coarse [76] or removed [205].
22	Geometry papers, referencing “hair” region on geometric model, e.g. Michaelangelo’s <i>David</i> .
24	Incidental image of humans, e.g. differentiable rendering where one of the examples is a head. [190]
7	Eyelash, eyebrow, and feather papers.
4	Crowd simulation and perception papers.

Table 2. Papers excluded from our analysis from the query hair AND (render OR rendering OR shading OR modeling OR modelling OR simulation OR animation).

blonde hair being preferred in general, cannot be detected using our labels, and are interesting avenues for future inquiry. Sub-categories of “non-white” like East Asian and South Asian were coded, but did not alter the overall findings. Additional observations from these labels are in Appendix B.

This reflexive thematic analysis [44] is necessarily a subjective and messy exercise. Outcomes may vary depending on the person performing it and the assumptions and racial commonsense they bring. Again, our point is not to establish an authoritative, replicable classification system, nor to use an existing system given the significant problems each has. Rather we seek to characterize the implicit assumptions in the computer graphics field in order to lay the groundwork for a more explicit conversation about racial representation going forward.

4 Results

4.1 Skin Papers

Our survey of the literature suggests that the term “skin” takes on at least three different meanings within the context of rendering. In order to understand their relationship, we will examine them in order of smallest to largest spatial scale:

- Skin as an *optical substrate*: the physics of light propagation through an infinitesimal physical medium. E.g. how a laser beam spreads as it impinges on a piece of marble [125].
- Skin as a *layered material*: light propagation through a mesoscale, heterogeneous medium. E.g. how light interacts with the oils and melanin on a swatch of skin, as well as the hemoglobin (blood) in deeper layers.
- Skin as a *human face*: characterizing light transport on a macro object like a head or body. This broadest category contains *facial capture*, which tries to acquire the appearance of a human subject’s face, and *reflectance field capture*, which tries to acquire the appearance of generic objects, including humans.

For all the following findings, more detailed breakdowns are available in Appendix B.

4.1.1 Skin as an optical medium. The 19 papers in this category suggest that at the basic optical level, SIGGRAPH researchers imagine skin as a pale, translucent material. Papers at this level repeatedly state equivalences to milk, marble, and snow. The early 1993 work of Hanrahan and Krueger [137] established this equivalence by grouping together “natural materials such as **human skin**, plant leaves, **snow**, sand, paint”. As a result, *subsurface scattering*, the physics of how light spreads diffusely through space as it propagates through a translucent material,




Fig. 2. 3D renderings from papers that investigate skin as an optical medium, respectively [83, 91, 111, 166, 167, 323], from *SIGGRAPH*, *SIGGRAPH Asia* and *ACM TOG*. There are no other renderings of humans in these papers.



Fig. 3. Examples from other papers presenting optical materials that are deemed analogous to “skin”. Alabaster [125], jade and wax [301], translucent plastics [140], milk [258], marble [373], and soap [82]. No humans appear in these publications.

dominates the research at this spatial scale. Subsequent to the 1993 paper [137], follow-on work goes further and draws equivalences to other extremely pale materials:

- “translucent materials, such as **marble**, cloth, paper, **skin**, **milk**, **cheese**, bread, meat, fruits, plants, fish, ocean water, **snow**, etc.” [167]
- “Translucent materials are frequently encountered in the natural world. Examples include **snow**, plants, **milk**, **cheese**, meat, **human skin**, cloth, **marble**, and jade” [166]
- “The accurate and efficient simulation of these effects is often required to achieve the color and soft appearance of media such as **skin**, hair, ocean water, **wax** and **marble**.” [83]
- “common materials are partially permeable to light such as **wax**, **skin**, teeth, **cheese**, fish, or stone” [45]

With the exception of the 1993 paper, [137] all of the subsurface scattering papers [83, 91, 111, 166, 167, 323] only show renderings of “skin” that feature pale, white skin (Figure 2). There are no images of Black (high melanin) skin, and the aesthetics and optical phenomena associated with high-melanin skin are not considered.

Many subsurface scattering papers do not contain any renderings of humans and focus on e.g. jade figurines or slabs of marble. They still consistently draw equivalences between “skin”, subsurface scattering, and pale materials. The language is consistent across papers, decades, and research groups, suggesting community-wide consensus. (Figure 3)

- “Many daily life objects (e.g., **milk**, **skin** or **marble**) are translucent” [125]
- “important for simulating the appearance of translucent materials, such as **marble**, **skin**, and **milk**.” [138]
- “materials such as **wax**, **marble**, and **skin** exhibit light scattering within the object volume.” [301]
- “subsurface scattering, which is important for translucent materials such as **wax**, **marble**, food, and **skin**.” [90]
- “this scattering is crucial to believable appearance of materials such as human **skin**, **wax**, and **marble**.” [140]
- “thick media (such as **soap**, **skin** and jade)” [158]




Fig. 4. Examples from papers presenting layered models of skin. In the left pair, swatches of different skin color are presented, but the final renders only show white skin. [92] In the right pair, a skin color lookup texture is mentioned, but the final results are only shown on white skin. [169].

The above papers were all written after 2001, after the publication of a highly influential 2001 “dipole approximation” paper [167] that saw widespread commercial adoption. The algorithm was used to render the translucent characters of Gollum in *The Two Towers* (2002) [3], Dobby in *Harry Potter and the Chamber of Secrets* (2001) [146], and won a 2004 Scientific and Technical Academy Award. It currently has 1384 citations on Google Scholar, exceeding the citation count of its predecessors (803 and 494 citations for [137] and [221]).

The equivalence of skin and snow is both implied and directly stated in commercial deployments. A 2018 special issue of *TOG* published full-length invited papers from all the major commercial and film studio rendering teams. Of these, the papers for Pixar’s *RenderMan*, Walt Disney Animation’s *Hyperion*, and Weta Digital’s *Manuka* rendering engines appeared in our keyword search. All three only label renderings of white, pale faces as instances of “skin,” and the *Hyperion* paper [48] even describes how their approach to skin evolved from their code for rendering snow.

4.1.2 Skin as a layered material. Skin is composed of multiple layers (stratum corneum, epidermis, and dermis [61]) and with the optical media abstraction of skin established in the 1990s and 2000s, subsequent papers attempted to model and render these heterogeneous layers (see Figure 4).

Throughout the 2000s and 2010s, layered material papers routinely showed only renderings of white skin as evidence of algorithmic effectiveness [61, 91, 245]. In one instance, the algorithm was based on data collected from “one Caucasian female, 33 years old; three Caucasian males, 26, 33, and 35 years old” [169]. In multiple instances, the existence of darker skin is acknowledged in the form of a “African, skin type V” swatch [92], a “skin color lookup texture” [169] and a model for melanin variation [61]. *None* of these papers showed final renders of anything except white hands and faces.

Weyrich et al. [337] from 2006 is worth closer examination, as one of the few papers that renders a Black face. The paper collects data from 149 subjects across a wide range of race, gender, and age. One of the main contributions is a handheld instrument for measuring translucency values, but the measurements show that the effect is *significantly* less pronounced in darker skin (Figures 11 and 12 in [337]). Black skin is not a block of wax, but as subsurface scattering becomes less pronounced, the possibility that a different visual feature might come to the fore is not explored.

However, it is well-known that the *gloss* and *shine* on darker skin, i.e. sharp and bright glints, are the more important visual cues. Many Black cinematographers have emphasized [139, 191, 214] that these features better convey shape and performance [338] with dark skin. Nevertheless, Weyrich et al. [337] does not posit the need for a better models for shine on skin, and instead re-uses established models for metals [317] and plastics [39]. One explanation for this is that existing models were sufficiently mature that new works were not necessary, i.e. the problem of glossiness was already solved. Later papers contradict this hypothesis. Starting in 2014, several papers performed detailed investigations into the optics of layered, glossy materials, and consistently acknowledged its importance in skin appearance:

- “layered biological structures like leaves, flower petals, or **skin**” [164]
- “Previous works on layered materials ... such as **human skin**” [366]
- “Hanrahan and Krueger [1993] proposed a general model for layered materials ... Stam [2001] generalized Hanrahan and Krueger’s work taking into account rough boundaries, in the context of **skin rendering**” [130]
- “many approximate models targeting specific layered materials such as ... **human skin**” [131]

None of these papers render a glossy-skinned human, Black or otherwise. The rendered objects are instead consistently household items like vases [131, 164], espresso machines [366], and pearlescent soap dispensers [130]. We discuss this reluctance in §5, relative to Thompson’s notion of the “slave sublime” [313]. Finally, we found one paper [324] on the effects of “darker skin tones” on blood flow monitoring and one [269] on tattooing across multiple skin tones. Both appeared post-George Floyd, which we will also discuss in §5.1.4.

4.1.3 Skin as a Human Face. As the largest category of “skin,” we sorted these papers into two broad sub-categories: *facial capture* and *reflectance field capture*. Since both dealt with the rendering humans on the scale of a face or body, the categories contained overlaps. All these papers used the *optical substrate* and *layered material* definitions of “skin” as both computational and epistemic baselines. We examined their racial distribution of examples to understand what was deemed sufficient evidence that an algorithm worked on “humans.”

Facial Capture. We examined 67 papers on facial capture, animation and rendering. For inclusion in this subcategory, a paper’s goal must be the generation of a photorealistic human face from 3D data. This included taking 2D photos, generating a textured 3D model, and then rendering with novel viewpoints, lighting, and/or facial expressions.

Within this set, 27% of the papers (18 of 67) contained only white examples. The papers are not localized in time, and span the 1980s [334], 1990s [128], and 2010s [310], with the most recent appearing in 2023 [7]. The remaining papers all contained non-white examples, but their distribution is highly uneven. Overall, 28% (10 of 67) of the papers contained Asian, but no Black, examples, while 4% (3 of 67) contained a only single Black example: Barack Obama.

Among the 39% of papers (26 of 61) that did contain other Black examples, almost all (21 of the 26) were published after 2020, again the post-George Floyd pattern we will discuss further in §5.1.4. Finally, 4% of the papers (3 of 67) feature *only* light-skinned East Asians as examples, and all appeared after 2020. We discuss this instance of racial elasticity further in §5.1.3. Under the current *SIGGRAPH* norms, it is still acceptable to publish a paper containing solely pale-skinned (white [7] or East Asian [63]) examples as evidence that an algorithm can render “humans.”

Facial Reflectance and Relighting. We examined 19 papers on the topic of facial reflectance field capture, whose initial application was the problem of “relighting” in a Hollywood films, i.e. changing the lighting on an actor’s face in post-production. Early versions of this technology as used in, e.g. *Spiderman 2* (2004) [280].

Within this set, we found 3 papers (16%) contained white-only final renders, with the most recent published in 2024 [199]. The rest contained non-white examples, but their distribution was again highly uneven. Asians were the only non-white examples in 3 papers (16%), one contained a single Latina example [132], and of the remaining 12 papers (63%), the vast majority (9 of the 12) were published after 2020.

In the papers showing a range of tones, the Black examples usually appear at the end, after the algorithm has been extensively presented and verified on white examples. They are also consistently flat and underlit (Figure 5), suggesting that the input photographs were taken under the same lighting setups as the white examples, and the outputs validated under the same conditions. The low contrast and flat colors recall the 2020 controversy [294, 314] surrounding a *Vogue* [2] article where white photographer Annie Leibovitz shot underlit and unflattering images of Black gymnast Simone Biles. The relighting papers similarly display Black skin under white lighting.



Fig. 5. Examples of Black skin in relighting algorithms, respectively [35, 293, 360, 362]. The skin is consistently underlit (too dark to see details), and sharp, well-defined highlights (“gloss”) characteristic of dark skin are often muted or missing.




Fig. 6. Inconsistent labels for “curly” hair. From left to right, “curly ponytail” [182], “curly hair” [335], “curly” [329], “curls” [145] and “curly” [279]. Many only contain light waves, and do not contain a full helical revolution (i.e. curl).

4.2 Hair Results

Whereas the computer graphics conception of “skin” is most commonly associated with its appearance, research into “hair” can reference its appearance, shape or motion. We decomposed hair papers into three categories:

- *Geometry* methods for constructing, editing, and capturing the *shape* of hair.
- *Simulation* methods for computing the *motion* of hair.
- *Rendering* methods for computing the *appearance* of hair, particularly how light interacting with hair.

These broad topic categories already exist in the *SIGGRAPH* literature, so we are largely following the categorizations that the papers self-report.

Hair Geometry. We identified 61 hair geometry papers. Of these, 27 papers (44%) showed results exclusively on straight or wavy hair, the most recent from 2024 [46]. Two additional papers (3%) showed only straight and wavy examples, plus Barack Obama. Another two (3%) dealt with “braids,” which meant straight hair braided into ponytails or pigtails. Braided textured hair styles such as box braids or cornrows were not considered.

We found 26 papers (43%) that showed “curly” examples in addition to straight and wavy examples. However, the definition of “curly” was *highly inconsistent* across the literature, and sometimes corresponded to hair that our computer graphics author labeled as “wavy.” They did not contain anything approaching a full helical revolution (Figure 6). Textured hair appeared in 4 papers (7%), two of which were as failure cases, where the paper’s algorithm was shown to fail on box braids [288] and “frizzy” hair [56]. Only one 2024 paper [346] explicitly attempted to model textured hair.

Hair Simulation. We identified 28 simulation papers that contained images of hair. Of these, 19 papers (68%) exclusively showed straight or wavy hair, while 5 (18%) papers also included curly hair. Two papers [70, 151] included textured hair, and were respectively from 2023 and 2024.

Another 7 simulation papers did not show images of hair, but made reference to “hair” that illustrated how the concept is situated in the scientific imagination:

- “physical objects such as **pieces of rope and strands of hair**” [232]
- “**hair, chains**, and other articulated models” [277]

- “wide class of 2d objects ranging from **wires and ropes to hair**” [86]
- “**ropes, chains, belts, cables, tendons, hair**, and other thin, curve-like physical objects” [303]
- “one-dimensional elastic **rods and hair**” [19]
- “**chains** and chainmail, **threads, knots, hair, and necklaces**” [272]
- “beams, **cables, yarn, and hair**” [135]

“Hair” is imagined to be an entirely straight solid, such as ropes, chains, wires, belts, cables, and rods. The main model for hair simulation at both Walt Disney Animation [316] and Weta Digital [197] is the *Discrete Elastics Rods* [26] model, so this imaginary of “straight as a rod” carries over into films such as *Moana* (2016).

One apparently missing paper is the curly hair system [160] Pixar developed for Princess Merida in the film *Brave* (2012). The journal-length version of that paper was rejected by both *SIGGRAPH* and *ACM Transactions on Graphics*, and instead appeared in the secondary conference the *Symposium on Computer Animation*. [159]

Hair Rendering. We found 64 papers on rendering, 35 of which dealt directly with hair rendering algorithms. Of these, 21 papers (60% of 35) showed exclusively straight and wavy hair, and 6 papers (17% of 35) also included curly hair. There were 8 papers (23% of 35) showing the full range of straight and wavy, curly, and textured hair, and all but one [208] were published after 2021. Within these rendering papers, we found that hair was initially abstracted as a translucent cylinder [171], and later refined to a cylinder covered in cuticles [220]. This conception of hair as a straight cylinder was reflected in 3 other papers on cloth rendering that directly compared hair to threads:

- “Like hair, textile fibers are long cylindrical structures made of dielectric material” [178]
- “unlike hair, threads have no consistent cuticle that displaces their specular reflection” [286]
- “This is similar to previous work on hair rendering” [235]

The special issue of *TOG* on production renderers appeared again, this time including all 5 papers from the issue. The Arnold renderer [120], as well as Sony Imageworks’ fork of Arnold [189], appeared in addition to the Pixar, Weta, and Disney renderers from §4.1.1. All of the papers showed exclusively straight or wavy hair [48, 64, 103, 120, 189]. A 2004 paper on the DreamWorks rendering system [306] for *Shrek 2* (2004) also appeared, and the only examples of hair were Prince Charming’s blonde bob, and the fur on Puss In Boots.

5 Discussion

5.1 Binarization and Hierarchization

5.1.1 Binarization. Our results show binarization in the algorithmic conceptualization of skin. The physical substrate of skin is formulated as a high albedo, highly scattering material analogous to snow, milk, wax and marble: white skin. This conceptualization appears to gain momentum circa 2011 with the commercial success of Jensen et al. [167]. Prior to that paper, the imaginary surrounding skin rendering included darker (South Asian) [221], and “dark complexion” skin [137]. After 2001, the possibility of humans with darker skin are acknowledged, but the algorithmic results put forward suggested that it was unnecessary to demonstrate that a technique can render equally convincing white and Black humans. The consistent, albeit unspoken, evidentiary benchmark is that results on white skin are acceptable synecdoche. Other shades of skin appear in large tables of disparate skin tones with no consistent standard for inclusion, aside from being darker than the white skin examples. We observe binarization into *white* and *non-white* skin.

We see similar binarization for hair. While for classification purposes we combined the categories for straight and wavy, the literature draws unambiguous analogies to ropes, wires, strands, and yarn: straight hair. Other hair types appear as derivations from this canonical type, often as an inconsistently labeled table that conflates the categories of wavy and curly. The only unifying characteristic among the “curly” examples is that they are less straight than the straight hair examples. We observe binarization into *straight* and *non-straight* hair.

5.1.2 Hierarchization. Klein and D'Ignazio [186] observe that “binaries are often hiding hierarchies.” With skin, the non-white form is conceived to descend from the white form. With one exception [339], every skin paper included an example of white skin. Many *only* showed white skin, while claiming to apply generically to “skin.”

The hierarchy is visible in Figure 4, where non-white skin receives a nod in the form of swatches and charts, but the final empirical validations only occur on white skin. These papers were accepted despite this gap in empiricism, suggesting that the hierarchy is a community-wide norm, not one researcher’s idiosyncrasy. The same hierarchy implicitly appears in literature gaps. When high-albedo phenomena is shown to be less salient in high-melanin skin [337], alternative visual phenomena associated with dark skin (gloss) goes unexplored.

Similarly, almost every paper published at *SIGGRAPH* includes straight hair. Non-straight hair is sufficiently subordinate that showing it is not necessary to claim that an algorithm applies to generic “hair.” The fungibility within the “non-straight” category is apparent in the images that get labeled as “curly” (Figure 6). In at least two known instances [160, 299], submitted papers that centered curly hair were rejected from *SIGGRAPH*. In one instance, the failure to acknowledge the centrality of straight hair was a reviewer rationale for rejection [180].

5.1.3 Racial Elasticity. Racial categories are not fixed, but shrink and grow in order to preserve existing power structures [67]. Privileging whiteness does not preclude representation of any other racial group [40, 318]. Our survey showed that in the 1980s to 2000s, white-only facial capture papers were most common, while starting in 2008 and accelerating in the 2010s, white and East Asian-only papers became acceptable. In the 2020s, East Asian-only facial capture papers became consistently acceptable at *SIGGRAPH*. Over the same time, no South Asian-only or Black-only papers appeared. This is consistent with scholarship on how the racial category “white” expands, and how its dominance can be buffered by other racial categories [40, 147]. The ongoing maintenance of a racial hierarchy privileging whiteness remains, to the exclusion of other (most notably Black or dark) racial groups. Under a theory of elasticity, “white” has expanded to include a light East Asian baseline, following a period of economic growth in East Asia. Pursuing a growing global market, multinational companies have circulated aspirational portrayals of Asian “lightness” [124, 194, 292].

This inclusion has occurred within the existing assumptions of the computer graphics field, rather than shifting the status quo. This fits the pattern observed by Lawrence [192] that these sorts of expansions are “not being driven by inclusionary ideals, it is about accessing an emerging profit center.” The profit incentives are also visible in the author affiliations of the *Facial Capture* papers, where of the 26 with inclusive skin tones, 5 came from Disney Research, 4 from Meta, 3 from Shanghai Tech startup Deemos Technology, 2 from USC startup Pinsceen, and 1 paper each came from Tencent, Technicolor, Snap, Google, and Kuaishou Technology. However, the high albedo models for skin and straight rod assumptions for hair were not revisited. Darker skin and curlier hair were deemed worthy of biometric and economic *capture*, but not of scientific *study*.

5.1.4 Post-George Floyd Tokenism. Our review also found a surge in tokenism after the 2020 murder of George Floyd that coincided with the ensuing worldwide protests for racial justice. Many skin tone tables in facial capture papers finally expanded to include examples of Black skin the following year (18 of 25). The first instances of highly coiled hair (Walker Type 4 [326]) appeared in 2023 [151, 378], following significant internal [181, 309] and external [179] organizing that directly cited Floyd’s murder. The only instance [346] of a *SIGGRAPH* paper treating highly coiled hair as the main research question came directly from this advocacy, i.e. some of the organizers authored the paper (as well as this one).

Despite this apparent progress, in at least two post-2020 instances [34, 74], examples of highly coiled hair were shown that do not correspond to any form of human hair. Highly coiled hair is composed of tight helices, but in the given examples, they were modeled as undifferentiated noise layered atop straight hair. (Figure 7) These examples passed peer review as evidence that the techniques apply generically to “hair,” underscoring how ignorance of even the basic shape of highly coiled hair remains an entrenched community-wide norm.




Fig. 7. Two tokenism examples. In the left pair [74], a simulation technique is extensively validated on a straight hair example before presenting a perplexing noisy-haired example at the end. In the right pair [34], a rendering technique is shown scaling up to an army of straight-haired models, while a perplexing “curly” example appears at the end.

5.2 McDaniels and Durald Methods

Binarization, hierarchization, and tokenism have antecedents in analog film and photography. Dyer [96] has described how “[i]nnovation in the photographic media has generally taken ... the white face as the norm.” This pattern repeated in §5.1 with the standardization of high albedo materials as the optical substrate of skin. The influence of film and photography has been considerable in the history of *SIGGRAPH* (Figure 1). Two of the most prolific labs from 1980-1990 were Lucasfilm and Pixar, with Disney becoming influential in the 2000s and 2010s. Microsoft’s rise coincides with its release of the Xbox, and the largest current industrial player is Adobe, the company behind Photoshop.

Dyer [96] suggests that the underlying ideology of a technology is visible in its standard operating manuals. One prominent lighting manual that successfully spanned both the analog and digital eras is *Painting With Light* by Alton [6] in 1949. The phrase goes back further [233], but Alton’s work is the one referenced in the graphics literature [80, 211, 265]. Alton displays the same hierarchization and tokenism in Chapter 5, “The Hollywood Close-Up,” when he exhaustively documents how to light a “face,” and every face is white. The only exception appears when Alton briefly notes that “[t]here is a widespread belief that close-ups of colored people have to be overlit. Nothing can be farther from the truth. By lighting them normally, we get a bronze-like skin texture.” The techniques used to light white actors can be transferred, without modification, to Black actors as well. Later scholarship has shown this to be false; cinematographers experienced in lighting Black skin consistently note that larger, diffuse light sources are preferred in order to achieve better shaping on darker skin [139, 338]. The token example in Alton [6] serves the rhetorical role of waving away concerns about the generalizability of a white skin technique.

The same flow appears in *SIGGRAPH* papers when algorithms get exhaustively validated on white skin and straight hair, with a token example of dark skin or curly hair inserted at the end. In both cases (Figures 5 and 7), the results can be visually disappointing and missing key features. While *an attempt was made* on a non-white example, whether it was successful is considered immaterial. While *SIGGRAPH* papers are not directly imitating Alton [6], they are transferring widespread norms of whiteness, as documented in Alton [6] and elsewhere, from the analog era into the digital.

As described in §1, we label techniques that deploy this rhetorical maneuver **McDaniels Methods**: *A technique that is exhaustively validated on white skin and straight hair, and relegates all other types of skin and hair to a small set of flawed and underwhelming examples.* To reiterate, we chose this name because the actress Hattie McDaniel is the only Black subject in Alton [6], but in a reflection of the carelessness given to her depiction, her name is misspelled. Attention to Black representation has historically been an afterthought in cinematography; the evidence presented here suggests it has remained so in computer graphics research.

We assign the label **Durald Methods** to the *techniques that were co-created in close collaboration with the people being depicted*. We named this after Autumn Durald Arkapaw, the winner of the 2025 Academy Award for Best Cinematography for the film *Sinners*. The lighting techniques for the film were designed in close collaboration with director Ryan Coogler to tell a story of racism and vampires in 1930s Mississippi [285].

5.3 Opportunities for Future Research

We examined the 6864 papers from *SIGGRAPH* and *TOG* because they are the most influential conference and journal in computer graphics. Expanding our study would be valuable future work. We estimate from Crossref.org that the journal corpus is ~26602 articles across *Computer Graphics Forum* (7104 articles), *IEEE Transactions on Visualization and Computer Graphics* (7512), *The Visual Computer* (5458), *Computers & Graphics* (5691) and *Graphical Models* (837).

Stepping away from a norm of whiteness opens up a variety of graphics research directions. A relative absence of research into the complexities of melanin has already been observed in physics [270], and while some graphics work also investigated its role in appearance, they are consistently McDaniels Methods [5, 61]. There has been no *SIGGRAPH* paper on the optical phenomena *characteristic of high melanin skin*, such as the well-known “blue undertones” [270]. The topic is ripe for investigation. In addition, there has been no investigation into the appearance of gloss in high melanin skin. The representational conflict between a diffuse, ethereal “glow” and a hard, well-defined “shine” was examined by Dyer [96] in film and photography, where the ideal has always been “white women ... bathed in and permeated by light ... in short, they glow.” Tracing these visual norms back to painting, Thompson [313] described how shine was used in “the slave sublime,” where “slave traders actually greased the bodies of enslaved Africans, using sweet oil or greasy water ‘to make them shine,’ as freedman Moses Roper described it ... to blind buyers, if you will, to the slave’s humanity.” Glowing skin as the preferred mode of human depiction predates computer graphics, while investigating shine violates a hegemonic visual norm in European art. However, for scientists willing to step outside of this norm, *developing a theory of high melanin shine* would advance the science of depiction.

Textured hair also offers multiple opportunities for investigation. The mechanics of *why hair curls* is still poorly understood, as the hypothesis that it was due to the hair’s elliptical cross-section was recently shown to be correlative, not causative [343]. A better understanding of where textured hair comes from would even have implications outside of computer graphics, in the same way that botanists have identified multiple mechanisms for tendril coiling in plants [102]. The spontaneous self-organization exhibited in textured hair also remains poorly understood. However, understanding *how thin structures self-order* is of interest in soft-matter physics [129] and has applications such as DNA packing into cell nuclei [183]. Detailed understanding of this phenomenon could potentially have wide impacts.

More broadly, the way forward is to reject the false promise of colorblind research and instead work to close heretofore neglected algorithmic gaps. Doing so will require sustained engagement with practitioners skilled in the art of depicting Black skin and hair, and a commitment to Durald Methods over McDaniels Methods. We know of at least one such instance, involving authors from this paper, where graphics researchers entered into a years-long collaboration with an artist who specializes in the digital depiction of Black hair. A sustained conversation on the aesthetics of Black hair, its characteristic shapes and motions, and what features are considered culturally desirable, yielded the first paper on textured hair to ever appear at *SIGGRAPH*. [346] In a reflection of the public appetite for such Durald Methods, the work went on to garner widespread press coverage. [13, 114, 252, 255, 275] We hope that by identifying these positive and negative patterns, this current work can spur similar advances in the science of depiction.

Generative AI Usage Statement

No generative AI was used in the preparation of this submission.

Positionality Statement

Our research team comprises scholars from multiple disciplines, including computer science, information science, sociology, and American studies. Across these fields, our work is informed by critical science and technology studies, with particular attention to bias, race, and the social implications of computational systems. Our prior research has examined labor relations in AI development, racial boundary making in the U.S. and Asia, and includes dozens of technical publications in computer graphics, including the *SIGGRAPH* and *Transactions on Graphics* venues examined here.

We conducted this study from our positions as researchers at Yale University, a well-resourced elite institution in the United States. The team has different backgrounds; one team member is an immigrant from Colombia working in the United States; another identifies as South Asian American, and the others as Korean American and White American. These institutional, disciplinary, and personal positions shaped how we approached the research questions, selected and interpreted methods, and understood the broader significance of the subject under study. We share these positions not to suggest that particular experiences are representative of the groups mentioned, but 1) to showcase that we do not all come from the position and 2) because our experiences in these positions contribute to our team's collective racial commonsense understandings, which we implicitly drew upon when making skin and racial classifications in the study.

There is no objective way of measuring difference, and we characterize our approach as subjective. Every "objective" classification system for skin is rooted in particular ideas about racial hierarchies, including in some cases eugenics and Nazism, as we note in §2 of the paper. Accordingly, as researchers, we eschewed these classification systems and instead applied blunt, high-level binary classifications: of skin as white/non-white and hair as straight/non-straight. Work in science and technology studies suggest that standard and classification efforts are never neutral, so we recognize that our classification efforts are necessarily subjective and not neutral either. We chose a classification strategy that allows us to open a broader conversation about race, while recognizing that the binary categories are too limited to do more than start that conversation.

Acknowledgments

This work was supported by NSF IIS-2132280, the Bungie Foundation, the Teng and Han Family Fund, and the Social Science Research Council's Just Tech Program, with funds provided by the Ford Foundation, MacArthur Foundation, and Surdna Foundation. The views, findings, conclusions, or recommendations expressed in this paper are those of the authors and do not necessarily represent those of the International Development Research Centre, the Social Science Research Council, Ford Foundation, MacArthur Foundation, and the Surdna Foundation or their governing boards. The authors also would like to thank Lauren F. Klein for early discussions and pointers to other field-level audits.

References

- [1] Alison Adam. 2006. *Artificial knowing: Gender and the thinking machine*. Routledge.
- [2] Abby Aguirre. 2020. Simone Biles on Overcoming Abuse, the Postponed Olympics, and Training During a Pandemic. *Vogue* (2020).
- [3] Matt Aitken, Greg Butler, Dan Lemmon, Eric Saindon, Dana Peters, and Guy Williams. 2004. The Lord of the Rings: the visual effects that brought middle earth to the screen. In *ACM SIGGRAPH Courses*. 11–es.
- [4] O Alexander, M Rogers, W Lambeth, J Chiang, W Ma, C Wang, and P Debevec. 2009. The digital emily project: Achieving a photoreal digital actor. *IEEE Computer Graphics and Applications* 30 (2009).
- [5] Carlos Aliaga, Mengqi Xia, Hao Xie, Adrian Jarabo, Gustav Braun, and Christophe Hery. 2023. A Hyperspectral Space of Skin Tones for Inverse Rendering of Biophysical Skin Properties. *Computer Graphics Forum* (2023). doi:10.1111/cgf.14887
- [6] John Alton. 1949. *Painting with light*. Macmillan.

- [7] Shivangi Aneja, Justus Thies, Angela Dai, and Matthias Nießner. 2023. Clipface: Text-guided editing of textured 3d morphable models. In *ACM SIGGRAPH 2023 Conference Proceedings*. 1–11.
- [8] Ken-ichi Anjyo, Yoshiaki Usami, and Tsuneya Kurihara. 1992. A simple method for extracting the natural beauty of hair. In *Proceedings of SIGGRAPH*. 111–120.
- [9] Jonghee Back, Binh-Son Hua, Toshiya Hachisuka, and Bochang Moon. 2022. Self-supervised post-correction for Monte Carlo denoising. In *ACM SIGGRAPH 2022 Conference Proceedings*. 1–8.
- [10] Kelian Baert, Shrisha Bharadwaj, Fabien Castan, Benoit Maujean, Marc Christie, Victoria Abrevaya, and Adnane Boukhayma. 2024. SPARK: Self-supervised Personalized Real-time Monocular Face Capture. In *SIGGRAPH Asia 2024 Conference Papers (Tokyo, Japan) (SA '24)*. Association for Computing Machinery, New York, NY, USA, Article 113, 12 pages. doi:10.1145/3680528.3687704
- [11] Shaojie Bai, Te-Li Wang, Chenghui Li, Akshay Venkatesh, Tomas Simon, Chen Cao, Gabriel Schwartz, Jason Saragih, Yaser Sheikh, and Shih-En Wei. 2024. Universal Facial Encoding of Codec Avatars from VR Headsets. *ACM Trans. Graph. (TOG)* 43, 4 (2024), 1–22.
- [12] Arthur E Balbao and Marcelo Walter. 2022. A Biologically Inspired Hair Aging Model. *ACM Trans. Graph. (TOG)* 41, 6 (2022), 1–9.
- [13] Brent Bambury. 2025. Realistic animation of Black hair has lagged far behind other hair types for years. This artist has a plan for how to fix that. *CBC, Day 6* (2025).
- [14] David C Banks. 1994. Illumination in diverse codimensions. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*. 327–334.
- [15] Linchao Bao, Xiangkai Lin, Yajing Chen, Haoxian Zhang, Sheng Wang, Xuefei Zhe, Di Kang, Haozhi Huang, Xinwei Jiang, Jue Wang, et al. 2021. High-fidelity 3d digital human head creation from rgb-d selfies. *ACM Trans. Graph. (TOG)* 41, 1 (2021), 1–21.
- [16] David Baraff and Andrew Witkin. 1998. Large steps in cloth simulation. In *Proceedings of SIGGRAPH*. 43–54.
- [17] Rasmus Barringer and Tomas Akenine-Möller. 2013. A4: asynchronous adaptive anti-aliasing using shared memory. *ACM Trans. Graph. (TOG)* 32, 4 (2013), 1–10.
- [18] Rasmus Barringer, Carl Johan Gribel, and Tomas Akenine-Möller. 2012. High-quality curve rendering using line sampled visibility. *ACM Trans. Graph. (TOG)* 31, 6 (2012), 1–10.
- [19] Christopher Batty, Andres Uribe, Basile Audoly, and Eitan Grinspun. 2012. Discrete viscous sheets. *ACM Trans. Graph. (TOG)* 31, 4 (2012), 1–7.
- [20] Thabo Beeler, Bernd Bickel, Paul Beardsley, Bob Sumner, and Markus Gross. 2010. High-quality single-shot capture of facial geometry. *ACM Trans. Graph.* 29, 4 (2010), 1–9.
- [21] Thabo Beeler, Bernd Bickel, Gioacchino Noris, Paul Beardsley, Steve Marschner, Robert W Sumner, and Markus Gross. 2012. Coupled 3D reconstruction of sparse facial hair and skin. *ACM Trans. Graph. (TOG)* 31, 4 (2012), 1–10.
- [22] Thabo Beeler and Derek Bradley. 2014. Rigid stabilization of facial expressions. *ACM Trans. Graph. (TOG)* 33, 4 (2014), 1–9.
- [23] Thabo Beeler, Fabian Hahn, Derek Bradley, Bernd Bickel, Paul A Beardsley, Craig Gotsman, Robert W Sumner, and Markus H Gross. 2011. High-quality passive facial performance capture using anchor frames. *ACM Trans. Graph.* 30, 4 (2011), 75.
- [24] Ruha Benjamin. 2019. *Race After Technology: Abolitionist Tools for the New Jim Code*. John Wiley & Sons.
- [25] Miklós Bergou, Basile Audoly, Etienne Vouga, Max Wardetzky, and Eitan Grinspun. 2010. Discrete viscous threads. *ACM Trans. Graph. (TOG)* 29, 4 (2010), 1–10.
- [26] Miklós Bergou, Max Wardetzky, Stephen Robinson, Basile Audoly, and Eitan Grinspun. 2008. Discrete Elastic Rods. *ACM Trans. Graph.* 27, 3 (aug 2008), 63:1–63:12.
- [27] Amit Bermano, Thabo Beeler, Yeara Kozlov, Derek Bradley, Bernd Bickel, and Markus Gross. 2015. Detailed spatio-temporal reconstruction of eyelids. *ACM Trans. Graph. (TOG)* 34, 4 (2015), 1–11.
- [28] Amit Bermano, Philipp Brüscheiler, Anselm Grundhöfer, Daisuke Iwai, Bernd Bickel, and Markus Gross. 2013. Augmenting physical avatars using projector-based illumination. *ACM Trans. Graph. (TOG)* 32, 6 (2013), 1–10.
- [29] Amit H Bermano, Derek Bradley, Thabo Beeler, Fabio Zund, Derek Nowrouzezahrai, Ilya Baran, Olga Sorkine-Hornung, Hanspeter Pfister, Robert W Sumner, Bernd Bickel, et al. 2014. Facial performance enhancement using dynamic shape space analysis. *ACM Trans. Graph. (TOG)* 33, 2 (2014), 1–12.
- [30] Florence Bertails, Basile Audoly, Marie-Paule Cani, Bernard Querleux, Frédéric Leroy, and Jean-Luc Lévêque. 2006. Super-helices for predicting the dynamics of natural hair. *ACM Trans. Graph.* 25, 3 (jul 2006), 1180–1187.
- [31] Florence Bertails, Basile Audoly, Bernard Querleux, Frédéric Leroy, Jean-Luc Lévêque, and Marie-Paule Cani. 2005. Predicting natural hair shapes by solving the statics of flexible rods. In *Eurographics short papers*. Eurographics.
- [32] Florence Bertails-Descoubes, Florent Cadoux, Gilles Daviet, and Vincent Acary. 2011. A nonsmooth Newton solver for capturing exact Coulomb friction in fiber assemblies. *ACM Trans. Graph. (TOG)* 30, 1 (2011), 1–14.
- [33] Shrisha Bharadwaj, Yufeng Zheng, Otmaz Hilliges, Michael J Black, and Victoria Fernández Abrevaya. 2023. FLARE: Fast Learning of Animatable and Relightable Mesh Avatars. *ACM Trans. Graph.* 42, 6 (2023), 204.
- [34] Gaurav Bhokare, Eisen Montalvo, Elie Diaz, and Cem Yuksel. 2024. Real-time hair rendering with hair meshes. In *Proceedings of SIGGRAPH*. 1–10.

- [35] Sai Bi, Stephen Lombardi, Shunsuke Saito, Tomas Simon, Shih-En Wei, Kevyn Mcphail, Ravi Ramamoorthi, Yaser Sheikh, and Jason Saragih. 2021. Deep relightable appearance models for animatable faces. *ACM Trans. Graph. (ToG)* 40, 4 (2021), 1–15.
- [36] Bernd Bickel, Mario Botsch, Roland Angst, Wojciech Matusik, Miguel Otaduy, Hanspeter Pfister, and Markus Gross. 2007. Multi-scale capture of facial geometry and motion. *ACM Trans. Graph. (TOG)* 26, 3 (2007), 33–es.
- [37] Bernd Bickel, Peter Kaufmann, Mélina Skouras, Bernhard Thomaszewski, Derek Bradley, Thabo Beeler, Phil Jackson, Steve Marschner, Wojciech Matusik, and Markus Gross. 2012. Physical face cloning. *ACM Trans. Graph. (TOG)* 31, 4 (2012), 1–10.
- [38] Abeba Birhane, Pratyusha Kalluri, Dallas Card, William Agnew, Ravit Dotan, and Michelle Bao. 2022. The values encoded in machine learning research. In *Proceedings of the 2022 ACM conference on fairness, accountability, and transparency*. 173–184.
- [39] James F Blinn. 1982. Light reflection functions for simulation of clouds and dusty surfaces. *Proceedings of SIGGRAPH* 16, 3 (1982), 21–29.
- [40] Eduardo Bonilla-Silva. 2004. From bi-racial to tri-racial: Towards a new system of racial stratification in the USA. *Ethnic and racial studies* 27, 6 (2004), 931–950.
- [41] Eduardo Bonilla-Silva. 2017. *Racism without Racists: Color-Blind Racism and the Persistence of Racial Inequality in America* (6th ed.). Rowman & Littlefield, Lanham, MD.
- [42] Geoffrey C Bowker and Susan Leigh Star. 2000. *Sorting things out: Classification and its consequences*. MIT press.
- [43] Derek Bradley, Wolfgang Heidrich, Tiberiu Popa, and Alla Sheffer. 2010. High resolution passive facial performance capture. *ACM Trans. Graph.* 29, 4 (2010), 41.
- [44] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative research in psychology* 3, 2 (2006), 77–101.
- [45] Alan Brunton, Can Ates Arikian, Tejas Madan Tanksale, and Philipp Urban. 2018. 3D printing spatially varying color and translucency. *ACM Trans. Graph. (TOG)* 37, 4 (2018), 1–13.
- [46] Marcel C Buehler, Gengyan Li, Erroll Wood, Leonhard Helming, Xu Chen, Tanmay Shah, Daoye Wang, Stephan Garbin, Sergio Orts-Escolano, Otmar Hilliges, et al. 2024. Cafca: High-quality novel view synthesis of expressive faces from casual few-shot captures. In *SIGGRAPH Asia 2024 Conference Papers*. 1–12.
- [47] Joy Buolamwini and Timnit Gebru. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*. PMLR, 77–91.
- [48] Brent Burley, David Adler, Matt Jen-Yuan Chiang, Hank Driskill, Ralf Habel, Patrick Kelly, Peter Kutz, Yining Karl Li, and Daniel Teece. 2018. The design and evolution of disney’s hyperion renderer. *ACM Trans. Graph. (TOG)* 37, 3 (2018), 1–22.
- [49] Chen Cao, Vasu Agrawal, Fernando De La Torre, Lele Chen, Jason Saragih, Tomas Simon, and Yaser Sheikh. 2021. Real-time 3D neural facial animation from binocular video. *ACM Trans. Graph. (TOG)* 40, 4 (2021), 1–17.
- [50] Chen Cao, Derek Bradley, Kun Zhou, and Thabo Beeler. 2015. Real-time high-fidelity facial performance capture. *ACM Trans. Graph. (ToG)* 34, 4 (2015), 1–9.
- [51] Chen Cao, Menglei Chai, Oliver Woodford, and Linjie Luo. 2018. Stabilized real-time face tracking via a learned dynamic rigidity prior. *ACM Trans. Graph. (TOG)* 37, 6 (2018), 1–11.
- [52] Chen Cao, Tomas Simon, Jin Kyu Kim, Gabe Schwartz, Michael Zollhoefer, Shunsuke Saito, Stephen Lombardi, Shih-En Wei, Danielle Belko, Shou-I Yu, et al. 2022. Authentic volumetric avatars from a phone scan. *ACM Trans. Graph. (TOG)* 41, 4 (2022), 1–19.
- [53] Chen Cao, Hongzhi Wu, Yanlin Weng, Tianjia Shao, and Kun Zhou. 2016. Real-time facial animation with image-based dynamic avatars. *ACM Trans. Graph.* 35, 4 (2016).
- [54] Edoardo Carra and Fabio Pellacini. 2019. SceneGit: a practical system for diffing and merging 3D environments. *ACM Trans. Graph. (TOG)* 38, 6 (2019), 1–15.
- [55] Romain Casati and Florence Bertails-Descoubes. 2013. Super space clothoids. *ACM Trans. Graph. (TOG)* 32, 4 (2013), 1–12.
- [56] Menglei Chai, Linjie Luo, Kalyan Sunkavalli, Nathan Carr, Sunil Hadap, and Kun Zhou. 2015. High-quality hair modeling from a single portrait photo. *ACM Trans. Graph.* 34, 6 (2015), 204–1.
- [57] Menglei Chai, Tianjia Shao, Hongzhi Wu, Yanlin Weng, and Kun Zhou. 2016. AutoHair: fully automatic hair modeling from a single image. *ACM Trans. Graph.*, Article 116 (July 2016), 12 pages.
- [58] Menglei Chai, Lvdi Wang, Yanlin Weng, Yizhou Yu, Baining Guo, and Kun Zhou. 2012. Single-view hair modeling for portrait manipulation. *ACM Trans. Graph. (TOG)* 31, 4 (2012), 1–8.
- [59] Menglei Chai, Changxi Zheng, and Kun Zhou. 2014. A reduced model for interactive hairs. *ACM Trans. Graph. (TOG)* 33, 4 (2014), 1–11.
- [60] Prashanth Chandran, Sebastian Winberg, Gaspard Zoss, Jérémy Riviere, Markus Gross, Paulo Gotardo, and Derek Bradley. 2021. Rendering with style: combining traditional and neural approaches for high-quality face rendering. *ACM Trans. Graph. (ToG)* 40, 6 (2021), 1–14.
- [61] Tenn F Chen, Gladimir VG Baranoski, Bradley W Kimmel, and Erik Miranda. 2015. Hyperspectral modeling of skin appearance. *ACM Trans. Graph. (TOG)* 34, 3 (2015), 1–14.
- [62] Xiao-Song Chen, Chen-Feng Li, Geng-Chen Cao, Yun-Tao Jiang, and Shi-Min Hu. 2020. A moving least square reproducing kernel particle method for unified multiphase continuum simulation. *ACM Trans. Graph. (TOG)* 39, 6 (2020), 1–15.

- [63] Yufan Chen, Lizhen Wang, Qijing Li, Hongjiang Xiao, Shengping Zhang, Hongxun Yao, and Yebin Liu. 2024. Monogaussianavatar: Monocular gaussian point-based head avatar. In *ACM SIGGRAPH 2024 Conference Papers*. 1–9.
- [64] Per Christensen, Julian Fong, Jonathan Shade, Wayne Wooten, Brenden Schubert, Andrew Kensler, Stephen Friedman, Charlie Kilpatrick, Cliff Ramshaw, Marc Bannister, et al. 2018. Renderman: An advanced path-tracing architecture for movie rendering. *ACM Trans. Graph. (TOG)* 37, 3 (2018), 1–21.
- [65] Nelson S-H Chu and Chiew-Lan Tai. 2005. Moxi: real-time ink dispersion in absorbent paper. *ACM Trans. Graph. (TOG)* 24, 3 (2005), 504–511.
- [66] Robert L Cook, John Halstead, Maxwell Planck, and David Ryu. 2007. Stochastic simplification of aggregate detail. *ACM Trans. Graph. (TOG)* 26, 3 (2007), 79–es.
- [67] Tressie McMillan Cottom. 2019. *Thick: And other essays*. The New Press.
- [68] Kate Crawford and Trevor Paglen. 2021. Excavating AI: The politics of images in machine learning training sets. *Ai & Society* 36, 4 (2021), 1105–1116.
- [69] Kimberlé Williams Crenshaw, Luke Charles Harris, Daniel Martinez HoSang, and George Lipsitz. 2019. *Seeing race again: Countering colorblindness across the disciplines*. Univ of California Press.
- [70] Octave Crespel, Emile Hohnadel, Thibaut Métivet, and Florence Bertails-Descoubes. 2024. Contact detection between curved fibres: high order makes a difference. In *Proceedings of SIGGRAPH*.
- [71] CurlCentric. 2014. *Hair Type Chart: How to Find Your Curl Pattern with Pictures*. <https://www.curlcentric.com/hair-typing-system/>
- [72] A.M. Darke, Isaac Olander, and Theodore Kim. 2024. More Than Killmonger Locs: A Style Guide for Black Hair (in Computer Graphics). In *ACM SIGGRAPH Courses (Denver, CO, USA) (SIGGRAPH Courses)*. Article 18, 251 pages.
- [73] Gilles Daviet. 2020. Simple and scalable frictional contacts for thin nodal objects. *ACM Trans. Graph.* 39, 4 (2020), 61–1.
- [74] Gilles Daviet. 2023. Interactive Hair Simulation on the GPU using ADMM. In *Proceedings of SIGGRAPH*. 1–11.
- [75] Gilles Daviet, Florence Bertails-Descoubes, and Laurence Boissieux. 2011. A hybrid iterative solver for robustly capturing coulomb friction in hair dynamics. *ACM Trans. Graph.* 30, 6 (2011), 1–12.
- [76] Edilson de Aguiar, Carsten Stoll, Christian Theobalt, Naveed Ahmed, Hans-Peter Seidel, and Sebastian Thrun. 2008. Performance capture from sparse multi-view video. *ACM Trans. Graph.* 27, 3 (2008), 1–10.
- [77] R. De La Mettrie, D. Saint-Léger, G. Loussouarn, A. Garcel, C. Porter, and A. Langaney. 2007. Shape variability and classification of human hair: a worldwide approach. *Human biology* 79, 3 (2007), 265–281.
- [78] Paul Debevec, Tim Hawkins, Chris Tchou, Haarm-Pieter Duiker, Westley Sarokin, and Mark Sagar. 2000. Acquiring the reflectance field of a human face. In *Proceedings of SIGGRAPH*. 145–156.
- [79] Paul Debevec, Andreas Wenger, Chris Tchou, Andrew Gardner, Jamie Waese, and Tim Hawkins. 2002. A lighting reproduction approach to live-action compositing. *ACM Trans. Graph. (TOG)* 21, 3 (2002), 547–556.
- [80] Christopher DeCoro, Forrester Cole, Adam Finkelstein, and Szymon Rusinkiewicz. 2007. Stylized shadows. In *Proceedings of the 5th international symposium on Non-photorealistic animation and rendering*. 77–83.
- [81] Ryane Defalco. 2025. *What 'Republican Hair' Really Means (& The Worst Offenders)*. <https://www.glam.com/1917335/republican-hair-trend-explained/>
- [82] Xi Deng, Fujun Luan, Bruce Walter, Kavita Bala, and Steve Marschner. 2022. Reconstructing translucent objects using differentiable rendering. In *Proceedings of SIGGRAPH*. 1–10.
- [83] Eugene d'Eon and Geoffrey Irving. 2011. A quantized-diffusion model for rendering translucent materials. *ACM Trans. Graph. (TOG)* 30, 4 (2011), 1–14.
- [84] Eugene d'Eon, David Luebke, and Eric Enderton. 2007. Efficient rendering of human skin. In *Proceedings of the 18th Eurographics conference on Rendering Techniques*. 147–157.
- [85] Alexandre Derouet-Jourdan, Florence Bertails-Descoubes, Gilles Daviet, and Joëlle Thollot. 2013. Inverse dynamic hair modeling with frictional contact. *ACM Trans. Graph. (TOG)* 32, 6 (2013), 1–10.
- [86] Alexandre Derouet-Jourdan, Florence Bertails-Descoubes, and Joëlle Thollot. 2010. Stable inverse dynamic curves. *ACM Trans. Graph. (TOG)* 29, 6 (2010), 1–10.
- [87] Piotr Didyk, Elmar Eisemann, Tobias Ritschel, Karol Myszkowski, and Hans-Peter Seidel. 2010. Apparent display resolution enhancement for moving images. *ACM Trans. Graph. (TOG)* 29, 4 (2010), 1–8.
- [88] Paul C DiLorenzo, Victor B Zordan, and Benjamin L Sanders. 2008. Laughing out loud: Control for modeling anatomically inspired laughter using audio. In *Proceedings of SIGGRAPH Asia*. 1–8.
- [89] Ana Dodik, Silvia Sellán, Theodore Kim, and Amanda Phillips. 2022. Sex and gender in the computer graphics research literature. In *ACM SIGGRAPH 2022 Talks*. 1–2.
- [90] Yue Dong, Jiaping Wang, Fabio Pellacini, Xin Tong, and Baining Guo. 2010. Fabricating spatially-varying subsurface scattering. *ACM Trans. Graph.* 29, 4 (2010), 1–10.
- [91] Craig Donner and Henrik Wann Jensen. 2005. Light diffusion in multi-layered translucent materials. *ACM Trans. Graph. (ToG)* 24, 3 (2005), 1032–1039.

- [92] Craig Donner, Tim Weyrich, Eugene d'Eon, Ravi Ramamoorthi, and Szymon Rusinkiewicz. 2008. A layered, heterogeneous reflectance model for acquiring and rendering human skin. *ACM Trans. Graph. (TOG)* 27, 5 (2008), 1–12.
- [93] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. 2017. CARLA: An open urban driving simulator. In *Conference on robot learning*. PMLR, 1–16.
- [94] Paul Dourish. 2006. Implications for design. In *Proceedings of the SIGCHI conference on Human Factors in computing systems*. 541–550.
- [95] Hao-Bin Duan, Miao Wang, Jin-Chuan Shi, Xu-Chuan Chen, and Yan-Pei Cao. 2023. Bakedavatar: Baking neural fields for real-time head avatar synthesis. *ACM Trans. Graph. (ToG)* 42, 6 (2023), 1–17.
- [96] Richard Dyer. 1997. *White*. Routledge.
- [97] Jose I Echevarria, Derek Bradley, Diego Gutierrez, and Thabo Beeler. 2014. Capturing and stylizing hair for 3D fabrication. *ACM Trans. Graph. (ToG)* 33, 4 (2014), 1–11.
- [98] Bernhard Egger, William AP Smith, Ayush Tewari, Stefanie Wuhrer, Michael Zollhoefer, Thabo Beeler, Florian Bernard, Timo Bolkart, Adam Kortylewski, Sami Romdhani, et al. 2020. 3d morphable face models—past, present, and future. *ACM Trans. Graph. (ToG)* 39, 5 (2020), 1–38.
- [99] Steven Epstein. 2024. *Inclusion: The politics of difference in medical research*. University of Chicago Press.
- [100] Eurographics. 2014. *Organization of Eurographics 2014*. <https://eg2014.unistra.fr/html/organization/organization.html>
- [101] Simone Fabbrizzi, Symeon Papadopoulos, Eirini Ntoutsis, and Ioannis Kompatsiaris. 2022. A survey on bias in visual datasets. *Computer Vision and Image Understanding* 223 (2022), 103552.
- [102] Muhammad Farhan, Frederike Klimm, Marc Thielen, Andraž Rešetič, Anil Bastola, Marc Behl, Thomas Speck, and Andreas Lendlein. 2023. Artificial tendrils mimicking plant movements by mismatching modulus and strain in core and shell. *Advanced Materials* 35, 22 (2023), 2211902.
- [103] Luca Fascione, Johannes Hanika, Mark Leone, Marc Droske, Jorge Schwarzhaupt, Tomáš Davidovič, Andrea Weidlich, and Johannes Meng. 2018. Manuka: A batch-shading architecture for spectral path tracing in movie production. *ACM Trans. Graph. (TOG)* 37, 3 (2018), 1–18.
- [104] Yun Fei, Christopher Batty, Eitan Grinspun, and Changxi Zheng. 2019. A multi-scale model for coupling strands with shear-dependent liquid. *ACM Trans. Graph. (TOG)* 38, 6 (2019), 1–20.
- [105] Yun Fei, Qi Guo, Rundong Wu, Li Huang, and Ming Gao. 2021. Revisiting integration in the material point method: a scheme for easier separation and less dissipation. *ACM Trans. Graph. (TOG)* 40, 4 (2021), 1–16.
- [106] Yun Fei, Henrique Teles Maia, Christopher Batty, Changxi Zheng, and Eitan Grinspun. 2017. A multi-scale model for simulating liquid-hair interactions. *ACM Trans. Graph. (TOG)* 36, 4 (2017), 1–17.
- [107] Yao Feng, Haiwen Feng, Michael J Black, and Timo Bolkart. 2021. Learning an animatable detailed 3D face model from in-the-wild images. *ACM Trans. Graph. (ToG)* 40, 4 (2021), 1–13.
- [108] Thomas B Fitzpatrick. 1975. Soleil et peau. *J. Med. Esthet.* 2 (1975), 33–34.
- [109] Thomas B Fitzpatrick. 1988. The validity and practicality of sun-reactive skin types I through VI. *Archives of dermatology* 124, 6 (1988), 869–871.
- [110] Kurt W Fleischer, David H Laidlaw, Bena L Currin, and Alan H Barr. 1995. Cellular texture generation. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*. 239–248.
- [111] Roald Frederickx and Philip Dutré. 2017. A forward scattering dipole model from a functional integral approximation. *ACM Trans. Graph. (TOG)* 36, 4 (2017), 1–13.
- [112] Graham Fyffe, Andrew Jones, Oleg Alexander, Ryosuke Ichikari, and Paul Debevec. 2014. Driving High-Resolution Facial Scans with Video Performance Capture. *ACM Trans. Graph.* 34, 1 (2014), 1–14.
- [113] Jacob Gaboury. 2021. *Image objects: An archaeology of computer graphics*. MIT Press.
- [114] Essence Gant. 2024. We're Much Closer to A Disney Princess With Type 4 Hair. *Allure* (2024).
- [115] Lin Gao, Feng-Lin Liu, Shu-Yu Chen, Kaiwen Jiang, Chungpeng Li, Yukun Lai, and Hongbo Fu. 2023. SketchFaceNeRF: Sketch-based facial generation and editing in neural radiance fields. *ACM Trans. Graph.* 42, 4 (2023).
- [116] Xuan Gao, Chenglai Zhong, Jun Xiang, Yang Hong, Yudong Guo, and Juyong Zhang. 2022. Reconstructing personalized semantic facial nerf models from monocular video. *ACM Trans. Graph. (TOG)* 41, 6 (2022), 1–12.
- [117] Pablo Garrido, Levi Valgaerts, Chenglei Wu, and Christian Theobalt. 2013. Reconstructing detailed dynamic face geometry from monocular video. *ACM Trans. Graph.* 32, 6 (2013), 158–1.
- [118] Pablo Garrido, Michael Zollhöfer, Dan Casas, Levi Valgaerts, Kiran Varanasi, Patrick Pérez, and Christian Theobalt. 2016. Reconstruction of personalized 3D face rigs from monocular video. *ACM Trans. Graph. (TOG)* 35, 3 (2016), 1–15.
- [119] Arianna B Gass. 2022. *The Body in Play: Embodiment in and Through Videogames*. Ph.D. Dissertation. The University of Chicago.
- [120] Iliyan Georgiev, Thiago Ize, Mike Farnsworth, Ramón Montoya-Vozmediano, Alan King, Brecht Van Lommel, Angel Jimenez, Oscar Anson, Shinji Ogaki, Eric Johnston, et al. 2018. Arnold: A brute-force production path tracer. *ACM Trans. Graph. (TOG)* 37, 3 (2018), 1–12.

- [121] Abhijeet Ghosh, Graham Fyffe, Borom Tunwattanapong, Jay Busch, Xueming Yu, and Paul E Debevec. 2011. Multiview face capture using polarized spherical gradient illumination. *ACM Trans. Graph.* (2011).
- [122] Abhijeet Ghosh, Tim Hawkins, Pieter Peers, Sune Frederiksen, and Paul Debevec. 2008. Practical modeling and acquisition of layered facial reflectance. *ACM Trans. Graph.* 27, 5 (2008), 1–10.
- [123] Simon Giebenhain, Tobias Kirschstein, Martin Rünz, Lourdes Agapito, and Matthias Nießner. 2024. Npga: Neural parametric gaussian avatars. In *SIGGRAPH Asia 2024 Conference Papers*. 1–11.
- [124] Evelyn Nakano Glenn. 2008. Yearning for lightness: Transnational circuits in the marketing and consumption of skin lighteners. *Gender & society* 22, 3 (2008), 281–302.
- [125] Michael Goesele, Hendrik PA Lensch, Jochen Lang, Christian Fuchs, and Hans-Peter Seidel. 2004. DISCO: acquisition of translucent objects. *ACM Trans. Graph.* 23, 3 (2004), 835–844.
- [126] Dan B Goldman. 1997. Fake fur rendering. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*. 127–134.
- [127] Paulo Gotardo, Jérémy Riviere, Derek Bradley, Abhijeet Ghosh, and Thabo Beeler. 2018. Practical dynamic facial appearance modeling and acquisition. *ACM Trans. Graph. (TOG)* 37, 6 (2018), 1–13.
- [128] Brian Guenter, Cindy Grimm, Daniel Wood, Henrique Malvar, and Fredric Pighin. 1998. Making faces. In *Proceedings of SIGGRAPH*. 55–66.
- [129] Arman Guerra, Anja C Slim, Douglas P Holmes, and Ousmane Kodio. 2023. Self-ordering of buckling, bending, and bumping beams. *Physical Review Letters* 130, 14 (2023), 148201.
- [130] Ibón Guillén, Julio Marco, Diego Gutierrez, Wenzel Jakob, and Adrian Jarabo. 2020. A general framework for pearlescent materials. *ACM Trans. Graph.* 39, 6 (2020), 1–15.
- [131] Jie Guo, Zeru Li, Xueyan He, Beibei Wang, Wenbin Li, Yanwen Guo, and Ling-Qi Yan. 2023. MetaLayer: a meta-learned BSDF model for layered materials. *ACM Trans. Graph. (TOG)* 42, 6 (2023), 1–15.
- [132] Kaiwen Guo, Peter Lincoln, Philip Davidson, Jay Busch, Xueming Yu, Matt Whalen, Geoff Harvey, Sergio Orts-Escolano, Rohit Pandey, Jason Dourgarian, et al. 2019. The relightables: Volumetric performance capture of humans with realistic relighting. *ACM Trans. Graph. (ToG)* 38, 6 (2019), 1–19.
- [133] Kunal Gupta, Milos Hasan, Zexiang Xu, Fujun Luan, Kalyan Sunkavalli, Xin Sun, Manmohan Chandraker, and Sai Bi. 2023. MCNeRF: Monte Carlo rendering and denoising for real-time NeRFs. In *SIGGRAPH Asia 2023 Conference Papers*. 1–11.
- [134] Paul Haerberli. 1990. Paint by numbers: Abstract image representations. In *Proceedings of SIGGRAPH*. 207–214.
- [135] Christian Hafner and Bernd Bickel. 2023. The design space of Kirchhoff rods. *ACM Trans. Graph.* 42, 5 (2023), 1–20.
- [136] Evelyn M Hammonds. 2013. New technologies of race. In *The Gendered Cyborg*. Routledge, 305–318.
- [137] Pat Hanrahan and Wolfgang Krueger. 1993. Reflection from layered surfaces due to subsurface scattering. In *Proceedings of SIGGRAPH*. 165–174.
- [138] Xuejun Hao and Amitabh Varshney. 2004. Real-time rendering of translucent meshes. *ACM Trans. Graph. (TOG)* 23, 2 (2004), 120–142.
- [139] Xavier Harding. 2017. Keeping “Insecure” lit: HBO cinematographer Ava Berkofsky on properly lighting black faces. *Mic.* 6 (2017).
- [140] Miloš Hašan, Martin Fuchs, Wojciech Matusik, Hanspeter Pfister, and Szymon Rusinkiewicz. 2010. Physical reproduction of materials with specified subsurface scattering. *ACM Trans. Graph.* 29, 4 (2010), 1–10.
- [141] Miloš Hašan, Fabio Pellacini, and Kavita Bala. 2006. Direct-to-indirect transfer for cinematic relighting. *ACM Trans. Graph. (TOG)* 25, 3 (2006), 1089–1097.
- [142] Mingming He, Pascal Clausen, Ahmet Levent Taşel, Li Ma, Oliver Pilarski, Wenqi Xian, Laszlo Rikker, Xueming Yu, Ryan Burgert, Ning Yu, and Paul Debevec. 2024. Diffrelight: Diffusion-based facial performance relighting. In *Proceedings of SIGGRAPH Asia*. 1–12.
- [143] Eric Heitz, Jonathan Dupuy, Cyril Crassin, and Carsten Dachsbacher. 2015. The SGGX microflake distribution. *ACM Trans. Graph. (TOG)* 34, 4 (2015), 1–11.
- [144] Joseph Henrich, Steven J Heine, and Ara Norenzayan. 2010. The weirdest people in the world? *Behavioral and brain sciences* 33, 2-3 (2010), 61–83.
- [145] Tomas Lay Herrera, Arno Zinke, and Andreas Weber. 2012. Lighting hair from the inside: A thermal approach to hair reconstruction. *ACM Trans. Graph. (TOG)* 31, 6 (2012), 1–9.
- [146] Christophe Hery. 2004. Rendering evolution at industrial light & magic. In *Proceedings of the Fifteenth Eurographics conference on Rendering Techniques*. 19–21.
- [147] Cathy Park Hong. 2021. *Minor feelings: an Asian American reckoning*. One World.
- [148] bell hooks. 1989. *Talking back: Thinking feminist, thinking black*. South End Press.
- [149] Qiming Hou, Hao Qin, Wenyao Li, Baining Guo, and Kun Zhou. 2010. Micropolygon ray tracing with defocus and motion blur. *ACM Trans. Graph. (TOG)* 29, 4 (2010), 1–10.
- [150] Jerry Hsu, Nghia Truong, Cem Yuksel, and Kui Wu. 2022. A general two-stage initialization for sag-free deformable simulations. *ACM Trans. Graph. (TOG)* 41, 4 (2022), 1–13.

- [151] Jerry Hsu, Tongtong Wang, Zherong Pan, Xifeng Gao, Cem Yuksel, and Kui Wu. 2023. Sag-Free Initialization for Strand-Based Hybrid Hair Simulation. *ACM Trans. Graph.* 42, 4 (2023), 1–14.
- [152] Jerry Hsu, Tongtong Wang, Zherong Pan, Xifeng Gao, Cem Yuksel, and Kui Wu. 2024. Real-time Physically Guided Hair Interpolation. *ACM Trans. Graph.* 43, 4, Article 95 (July 2024), 11 pages. doi:10.1145/3658176
- [153] Liwen Hu, Chongyang Ma, Linjie Luo, and Hao Li. 2014. Robust hair capture using simulated examples. *ACM Trans. Graph. (TOG)* 33, 4 (2014), 1–10.
- [154] Liwen Hu, Chongyang Ma, Linjie Luo, and Hao Li. 2015. Single-view hair modeling using a hairstyle database. *ACM Trans. Graph. (ToG)* 34, 4 (2015), 1–9.
- [155] Liwen Hu, Chongyang Ma, Linjie Luo, Li-Yi Wei, and Hao Li. 2014. Capturing braided hairstyles. *ACM Trans. Graph. (TOG)* 33, 6 (2014), 1–9.
- [156] Liwen Hu, Shunsuke Saito, Lingyu Wei, Koki Nagano, Jaewoo Seo, Jens Fursund, Iman Sadeghi, Carrie Sun, Yen-Chun Chen, and Hao Li. 2017. Avatar digitization from a single image for real-time rendering. *ACM Trans. Graph. (ToG)* 36, 6 (2017), 1–14.
- [157] W. Huang, M. B. Hullin, and J. Hanika. 2022. A Microfacet-based Hair Scattering Model. *Computer Graphics Forum* 41, 4 (2022), 79–91. arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.14588> doi:10.1111/cgf.14588
- [158] Yuchi Huo, Rui Wang, Tianlei Hu, Wei Hua, and Hujun Bao. 2016. Adaptive matrix column sampling and completion for rendering participating media. *ACM Trans. Graph. (TOG)* 35, 6 (2016), 1–11.
- [159] Hayley Iben. 2013. Personal Communication.
- [160] Hayley Iben, Mark Meyer, Lena Petrovic, Olivier Soares, John Anderson, and Andrew Witkin. 2013. Artistic simulation of curly hair. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*. 63–71.
- [161] Alexandru Eugen Ichim, Sofien Bouaziz, and Mark Pauly. 2015. Dynamic 3D avatar creation from hand-held video input. *ACM Trans. Graph. (ToG)* 34, 4 (2015), 1–14.
- [162] Noel Ignatiev. 2012. *How the Irish became white*. Routledge.
- [163] Nina G Jablonski. 2004. The evolution of human skin and skin color. *Annu. Rev. Anthropol.* 33, 1 (2004), 585–623.
- [164] Wenzel Jakob, Eugene d'Eon, Otto Jakob, and Steve Marschner. 2014. A comprehensive framework for rendering layered materials. *ACM Trans. Graph. (ToG)* 33, 4 (2014), 1–14.
- [165] Wenzel Jakob, Jonathan T Moon, and Steve Marschner. 2009. Capturing hair assemblies fiber by fiber. *ACM Trans. Graph. (TOG)* 28, 5 (2009), 1–9.
- [166] Henrik Wann Jensen and Juan Buhler. 2002. A rapid hierarchical rendering technique for translucent materials. In *Proceedings of SIGGRAPH*. 576–581.
- [167] Henrik Wann Jensen, Stephen R Marschner, Marc Levoy, and Pat Hanrahan. 2001. A practical model for subsurface light transport. In *Proceedings of SIGGRAPH*. 511–518.
- [168] Chenfanfu Jiang, Theodore Gast, and Joseph Teran. 2017. Anisotropic elastoplasticity for cloth, knit and hair frictional contact. *ACM Trans. Graph. (TOG)* 36, 4 (2017), 1–14.
- [169] Jorge Jimenez, Timothy Scully, Nuno Barbosa, Craig Donner, Xenxo Alvarez, Teresa Vieira, Paul Matts, Verónica Orvalho, Diego Gutierrez, and Tim Weyrich. 2010. A practical appearance model for dynamic facial color. *ACM Trans. Graph. (TOG)* 29, 6 (2010), 1–10.
- [170] Mattie Kahn. 2025. *The Rise of Big Washington Hair*. <https://www.townandcountrymag.com/style/beauty-products/a64220319/washington-dc-hair-trends-explained/>
- [171] James T Kajiya and Timothy L Kay. 1989. Rendering fur with three dimensional textures. *Proceedings of SIGGRAPH* 23, 3 (1989), 271–280.
- [172] Pratyusha Ria Kalluri, William Agnew, Myra Cheng, Kentrell Owens, Luca Soldaini, and Abeba Birhane. 2025. Computer-vision research powers surveillance technology. *Nature* (2025), 1–7.
- [173] Yarden Katz. 2020. *Artificial whiteness: Politics and ideology in artificial intelligence*. Columbia University Press.
- [174] Danny M. Kaufman, Rasmus Tamstorf, Breannan Smith, Jean-Marie Aubry, and Eitan Grinspun. 2014. Adaptive Nonlinearity for Collisions in Complex Rod Assemblies. *ACM Trans. Graph.* (2014).
- [175] Karyl E Ketchum. 2009. Facegen and the technovisual politics of embodied surfaces. *WSQ: Women's Studies Quarterly* 37, 1 (2009), 183–199.
- [176] Os Keyes. 2018. The misgendering machines: Trans/HCI implications of automatic gender recognition. *Proceedings of the ACM on human-computer interaction* 2, CSCW (2018), 1–22.
- [177] Pramook Khungurn and Steve Marschner. 2017. Azimuthal scattering from elliptical hair fibers. *ACM Trans. Graph. (TOG)* 36, 2 (2017), 1–23.
- [178] Pramook Khungurn, Daniel Schroeder, Shuang Zhao, Kavita Bala, and Steve Marschner. 2015. Matching real fabrics with micro-appearance models. *ACM Trans. Graph.* 35, 1 (2015), 1–1.
- [179] Theodore Kim. 2020. The Racist Legacy of Computer-Generated Humans. *Scientific American* (2020).
- [180] Theodore Kim. 2023. The Character of Computer Animation Research. Keynote, Symposium on Geometry Processing.

- [181] Theodore Kim, Holly Rushmeier, Julie Dorsey, Derek Nowrouzezahrai, Raqi Syed, Wojciech Jarosz, and AM Darke. 2022. Countering racial bias in computer graphics research. In *ACM SIGGRAPH Talks*. 1–2.
- [182] Tae-Yong Kim and Ulrich Neumann. 2002. Interactive multiresolution hair modeling and editing. *ACM Trans. Graph.* 21, 3 (July 2002), 620–629. doi:10.1145/566654.566627
- [183] James Kindt, Shelly Tzllil, Avinoam Ben-Shaul, and William M Gelbart. 2001. DNA packaging and ejection forces in bacteriophage. *Proceedings of the National Academy of Sciences* 98, 24 (2001), 13671–13674.
- [184] Tobias Kirschstein, Simon Giebenhain, Jiapeng Tang, Markos Georgopoulos, and Matthias Nießner. 2024. Gghead: Fast and generalizable 3d gaussian heads. In *SIGGRAPH Asia 2024 Conference Papers*. 1–11.
- [185] Tobias Kirschstein, Shenhan Qian, Simon Giebenhain, Tim Walter, and Matthias Nießner. 2023. Nersemble: Multi-view radiance field reconstruction of human heads. *ACM Trans. Graph. (TOG)* 42, 4 (2023), 1–14.
- [186] Lauren Klein and Catherine D'Ignazio. 2024. Data Feminism for AI. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency* (Rio de Janeiro, Brazil) (FAccT '24). Association for Computing Machinery, New York, NY, USA, 100–112. doi:10.1145/3630106.3658543
- [187] Devon Knightner, Allie J. Hoffman, Julia Stone, and Mimi Bascom. 2018. *The ULTIMATE Guide to Pageant Hair*. <https://www.pageantplanet.com/article/the-ultimate-guide-to-pageant-hair>
- [188] Zhiyi Kuang, Yiyang Chen, Hongbo Fu, Kun Zhou, and Youyi Zheng. 2022. DeepMVSHair: Deep Hair Modeling from Sparse Views. In *Proceedings of SIGGRAPH* (Daegu, Republic of Korea) (SA '22). Association for Computing Machinery, New York, NY, USA, Article 10, 8 pages. doi:10.1145/3550469.3555385
- [189] Christopher Kulla, Alejandro Conty, Clifford Stein, and Larry Gritz. 2018. Sony pictures imageworks Arnold. *ACM Trans. Graph. (TOG)* 37, 3 (2018), 1–18.
- [190] Samuli Laine, Janne Hellsten, Tero Karras, Yeongho Seol, Jaakko Lehtinen, and Timo Aila. 2020. Modular primitives for high-performance differentiable rendering. *ACM Trans. Graph. (TOG)* 39, 6 (2020), 1–14.
- [191] Nadia Latif. 2017. It's lit! How film finally learned to light black skin. *The Guardian* 21 (2017).
- [192] Halcyon Lawrence. 2021. Siri Disciplines. In *Your Computer Is On Fire*. MIT Press.
- [193] Jason Lawrence, Danb Goldman, Supreeth Achar, Gregory Major Blascovich, Joseph G Desloge, Tommy Fortes, Eric M Gomez, Sascha Häberling, Hugues Hoppe, Andy Huibers, et al. 2021. Project starline: a high-fidelity telepresence system. *ACM Trans. Graph. (TOG)* 40, 6 (2021), 1–16.
- [194] S Hejjin Lee, Christina H Moon, and Thuy Linh Nguyen Tu. 2019. *Fashion and Beauty in the Time of Asia*. Vol. 6. NYU Press.
- [195] Aaron E Lefohn, Shubhabrata Sengupta, and John D Owens. 2007. Resolution-matched shadow maps. *ACM Trans. Graph. (TOG)* 26, 4 (2007), 20–es.
- [196] Chloe LeGendre, Xueming Yu, Dai Liu, Jay Busch, Andrew Jones, Sumanta Pattanaik, and Paul Debevec. 2016. Practical multispectral lighting reproduction. *ACM Trans. Graph. (TOG)* 35, 4 (2016), 1–11.
- [197] Steve Lesser, Alexey Stomakhin, Gilles Daviet, Joel Wretborn, John Edholm, Noh-Hoon Lee, Eston Schweickart, Xiao Zhai, Sean Flynn, and Andrew Moffat. 2022. Loki: a unified multiphysics simulation framework for production. *ACM Trans. Graph. (TOG)* 41, 4 (2022), 1–20.
- [198] Chenghong Li, Leyang Jin, Yujian Zheng, Yizhou Yu, and Xiaoguang Han. 2023. EMS: 3D Eyebrow Modeling from Single-view Images. *ACM Trans. Graph. (TOG)* 42, 6 (2023), 1–19.
- [199] Junxuan Li, Chen Cao, Gabriel Schwartz, Rawal Khirodkar, Christian Richardt, Tomas Simon, Yaser Sheikh, and Shunsuke Saito. 2024. Uravatar: Universal relightable gaussian codec avatars. In *SIGGRAPH Asia 2024 Conference Papers*. 1–11.
- [200] Jiaman Li, Zhengfei Kuang, Yajie Zhao, Mingming He, Karl Bladin, and Hao Li. 2020. Dynamic facial asset and rig generation from a single scan. *ACM Trans. Graph.* 39, 6 (2020), 215–1.
- [201] Minchen Li, Danny M Kaufman, and Chenfanfu Jiang. 2021. Codimensional incremental potential contact. *ACM Trans. Graph.* 40, 4 (2021), 1–24.
- [202] Zhansheng Li, Yangyang Xu, Nanxuan Zhao, Yang Zhou, Yongtuo Liu, Dahua Lin, and Shengfeng He. 2023. Parsing-conditioned anime translation: A new dataset and method. *ACM Trans. Graph.* 42, 3 (2023), 1–14.
- [203] Shu Liang, Xiufeng Huang, Xianyu Meng, Kunyao Chen, Linda G. Shapiro, and Ira Kemelmacher-Shlizerman. 2018. Video to fully automatic 3D hair model. *ACM Trans. Graph.* 37, 6, Article 206 (Dec. 2018), 14 pages.
- [204] Ming Lin, Sophie Jörg, Jihun Yu, and Tuur Stuyck. 2023. From Traditional Animation to modern generative AI: How Technology is Changing the Face of Animation. In *Industry Panel, ACM SIGGRAPH / Eurographics Symposium on Computer Animation*.
- [205] Lijuan Liu, Xiangyu Xu, Zhijie Lin, Jiabin Liang, and Shuicheng Yan. 2023. Towards garment sewing pattern reconstruction from a single image. *ACM Trans. Graph. (TOG)* 42, 6 (2023), 1–15.
- [206] Shichen Liu, Yunxuan Cai, Haiwei Chen, Yichao Zhou, and Yajie Zhao. 2022. Rapid face asset acquisition with recurrent feature alignment. *ACM Trans. Graph. (TOG)* 41, 6 (2022), 1–17.
- [207] T Locovic and Eric Veach. 2000. Deep Shadow Maps. In *Proceedings of ACM SIGGRAPH*.

- [208] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. 2019. Neural volumes: learning dynamic renderable volumes from images. *ACM Trans. Graph. (TOG)* 38, 4 (2019), 1–14.
- [209] Stephen Lombardi, Tomas Simon, Gabriel Schwartz, Michael Zollhoefer, Yaser Sheikh, and Jason Saragih. 2021. Mixture of volumetric primitives for efficient neural rendering. *ACM Trans. Graph. (ToG)* 40, 4 (2021), 1–13.
- [210] Benjamin Long, Sue Ann Seah, Tom Carter, and Sriram Subramanian. 2014. Rendering volumetric haptic shapes in mid-air using ultrasound. *ACM Trans. Graph. (TOG)* 33, 6 (2014), 1–10.
- [211] Jorge Lopez-Moreno, Jorge Jimenez, Sunil Hadap, Erik Reinhard, Ken Anjyo, and Diego Gutierrez. 2010. Stylized depiction of images based on depth perception. In *Proceedings of the 8th International Symposium on Non-Photorealistic Animation and Rendering*. 109–118.
- [212] Linjie Luo, Hao Li, and Szymon Rusinkiewicz. 2013. Structure-aware hair capture. *ACM Trans. Graph. (TOG)* 32, 4 (2013), 1–12.
- [213] Yiwei Luo, Kristina Gligorić, and Dan Jurafsky. 2024. Othering and low status framing of immigrant cuisines in US restaurant reviews and large language models. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 18. 985–998.
- [214] Linda Lynton. 1988. School Daze: Black College is Background. *American Cinematographer* 59, 2 (1988), 67–70.
- [215] Chongyang Ma, Li-Yi Wei, Sylvain Lefebvre, and Xin Tong. 2013. Dynamic element textures. *ACM Trans. Graph. (TOG)* 32, 4 (2013), 1–10.
- [216] Li Ma, Xiaoyu Li, Jing Liao, Xuan Wang, Qi Zhang, Jue Wang, and Pedro V Sander. 2022. Neural parameterization for dynamic human head editing. *ACM Trans. Graph. (TOG)* 41, 6 (2022), 1–15.
- [217] Wan-Chun Ma, Andrew Jones, Jen-Yuan Chiang, Tim Hawkins, Sune Frederiksen, Pieter Peers, Marko Vukovic, Ming Ouhyoung, and Paul Debevec. 2008. Facial performance synthesis using deformation-driven polynomial displacement maps. *ACM Trans. Graph. (TOG)* 27, 5 (2008), 1–10.
- [218] James Malazita. 2022. Using STS to Bridge Long Histories of Blackness, Specularity, and Rendering. In *ACM SIGGRAPH 2022 Talks*. 1–2.
- [219] James Malazita. 2024. *Enacting Platforms: Feminist Technoscience and the Unreal Engine*. MIT Press.
- [220] Stephen R Marschner, Henrik Wann Jensen, Mike Cammarano, Steve Worley, and Pat Hanrahan. 2003. Light scattering from human hair fibers. *ACM Trans. Graph. (TOG)* 22, 3 (2003), 780–791.
- [221] Stephen R Marschner, Stephen H Westin, Eric PF LaFortune, Kenneth E Torrance, and Donald P Greenberg. 1999. Image-based BRDF measurement including human skin. In *Eurographics Workshop on Rendering Techniques*. Springer, 131–144.
- [222] Wojciech Matusik, Hanspeter Pfister, Addy Ngan, Paul Beardsley, Remo Ziegler, and Leonard McMillan. 2002. Image-based 3D photography using opacity hulls. *ACM Trans. Graph. (TOG)* 21, 3 (2002), 427–437.
- [223] A. McAdams, A. Selle, K. Ward, E. Sifakis, and J. Teran. 2009. Detail preserving continuum simulation of straight hair. *ACM Trans. Graph.* 28, 3 (2009), 1–6.
- [224] Radomír Měch and Przemyslaw Prusinkiewicz. 1996. Visual models of plants interacting with their environment. In *Proceedings of SIGGRAPH*. 397–410.
- [225] Abhimitra Meka, Christian Haene, Rohit Pandey, Michael Zollhöfer, Sean Fanello, Graham Fyffe, Adarsh Kowdle, Xueming Yu, Jay Busch, Jason Dourgarian, et al. 2019. Deep reflectance fields: high-quality facial reflectance field inference from color gradient illumination. *ACM Trans. Graph. (TOG)* 38, 4 (2019), 1–12.
- [226] Abhimitra Meka, Rohit Pandey, Christian Haene, Sergio Orts-Escolano, Peter Barnum, Philip David-Son, Daniel Erickson, Yinda Zhang, Jonathan Taylor, Sofien Bouaziz, et al. 2020. Deep relightable textures: volumetric performance capture with neural rendering. *ACM Trans. Graph. (TOG)* 39, 6 (2020), 1–21.
- [227] Mohit Mendiratta, Xingang Pan, Mohamed Elgharib, Kartik Teotia, Mallikarjun B R, Ayush Tewari, Vladislav Golyanik, Adam Kortylewski, and Christian Theobalt. 2023. Avatarstudio: Text-driven editing of 3d dynamic human head avatars. *ACM Trans. Graph. (ToG)* 42, 6 (2023), 1–18.
- [228] Alka Menon. 2017. Reconstructing race and gender in American cosmetic surgery. *Ethnic and Racial Studies* 40, 4 (2017), 597–616.
- [229] Alka Vaid Menon. 2023. *Refashioning Race: How Global Cosmetic Surgery Crafts New Beauty Standards*. Univ of California Press.
- [230] Milagros Miceli and Julian Posada. 2022. The data-production dispositif. *Proceedings of the ACM on human-computer interaction* 6, CSCW2 (2022), 1–37.
- [231] Dominik L Michels, Vu Thai Luan, and Mayya Tokman. 2017. A stiffly accurate integrator for elastodynamic problems. *ACM Trans. Graph. (TOG)* 36, 4 (2017), 1–14.
- [232] Gavin SP Miller. 1988. The motion dynamics of snakes and worms. In *Proceedings of SIGGRAPH*. 169–173.
- [233] Victor Milner. 1930. Painting with Light. *Cinematographic Annual* 1 (1930), 91–108.
- [234] Ellis Monk. 2023. The monk skin tone scale. (2023).
- [235] Zahra Montazeri, Søren B Gammelmark, Shuang Zhao, and Henrik Wann Jensen. 2020. A practical ply-based appearance model of woven fabrics. *ACM Trans. Graph. (TOG)* 39, 6 (2020), 1–13.
- [236] Bochang Moon, Yongyoung Byun, Tae-Joon Kim, Pio Claudio, Hye-Sun Kim, Yun-Ji Ban, Seung Woo Nam, and Sung-Eui Yoon. 2010. Cache-oblivious ray reordering. *ACM Trans. Graph. (TOG)* 29, 3 (2010), 1–10.
- [237] Jonathan T Moon and Stephen R Marschner. 2006. Simulating multiple scattering in hair using a photon mapping approach. *ACM Trans. Graph. (TOG)* 25, 3 (2006), 1067–1074.

- [238] Jonathan T Moon, Bruce Walter, and Steve Marschner. 2008. Efficient multiple scattering in hair using spherical harmonics. *ACM Trans. Graph.* 27, 3 (2008), 1–7.
- [239] Masahiro Mori, Karl F MacDorman, and Norri Kageki. 2012. The uncanny valley [from the field]. *IEEE Robotics & automation magazine* 19, 2 (2012), 98–100.
- [240] Ann Morning. 2011. *The nature of race: How scientists think and teach about human difference*. Univ of California Press.
- [241] Lucio Moser, Chinyu Chien, Mark Williams, Jose Serra, Darren Hendler, and Doug Roble. 2021. Semi-supervised video-driven facial animation transfer for production. *ACM Trans. Graph. (TOG)* 40, 6 (2021), 1–18.
- [242] United States Holocaust Memorial Museum. 2017. *Tool used to classify skin color in racial studies conducted in Nazi Germany*. <https://collections.ushmm.org/search/catalog/irn564926>
- [243] James Muldoon, Callum Cant, Mark Graham, and Funda Ustek Spilda. 2025. The poverty of ethical AI: impact sourcing and AI supply chains. *AI & society* 40, 2 (2025), 529–543.
- [244] Krishna Mullia, Fujun Luan, Xin Sun, and Miloš Hašan. 2024. Rna: Relightable neural assets. *ACM Trans. Graph.* 44, 1 (2024), 1–19.
- [245] Koki Nagano, Graham Fyffe, Oleg Alexander, Jernej Barbic, Hao Li, Abhijeet Ghosh, and Paul E Debevec. 2015. Skin microstructure deformation with displacement map convolution. *ACM Trans. Graph.* 34, 4 (2015), 109–1.
- [246] Koki Nagano, Jaewoo Seo, Jun Xing, Lingyu Wei, Zimo Li, Shunsuke Saito, Aviral Agarwal, Jens Fursund, Hao Li, Richard Roberts, et al. 2018. paGAN: real-time avatars using dynamic textures. *ACM Trans. Graph.* 37, 6 (2018), 258.
- [247] Lisa Nakamura. 2007. *Digitizing race: Visual cultures of the Internet*. Vol. 23. U of Minnesota Press.
- [248] Evan Narcisse. 2017. The natural: The trouble portraying blackness in video games. *Kotaku* (2017).
- [249] Michael B Nielsen, Morten Bojsen-Hansen, Konstantinos Stamatelos, and R. Bridson. 2022. Physics-based combustion simulation. *ACM Trans. Graph. (TOG)* 41, 5 (2022), 1–21.
- [250] Ko Nishino and Shree K Nayar. 2004. Eyes for relighting. *ACM Trans. Graph. (TOG)* 23, 3 (2004), 704–711.
- [251] Safiya Umoja Noble. 2018. *Algorithms of Oppression: How Search Engines Reinforce Racism*. NYU Press.
- [252] Gloria Oladipo. 2025. ‘This is Black hair’: technological advances are making waves in animation. *The Guardian* (2025).
- [253] Kyle Olszewski, Joseph J Lim, Shunsuke Saito, and Hao Li. 2016. High-fidelity facial and speech animation for VR HMDs. *ACM Trans. Graph. (TOG)* 35, 6 (2016), 1–14.
- [254] Michael Omi and Howard Winant. 2014. *Racial formation in the United States*. Routledge.
- [255] Lizette Ortega. 2024. Animations of coiled hair for Black film characters improve with new algorithms. *The Washington Post* (2024).
- [256] Matthew O’Toole and Kiriakos N Kutulakos. 2010. Optical computing for fast light transport analysis. *ACM Trans. Graph. (ToG)* 29, 6 (2010), 1–12.
- [257] Rohit Pandey, Sergio Orts-Escolano, Chloe Legendre, Christian Haene, Sofien Bouaziz, Christoph Rhemann, Paul E Debevec, and Sean Ryan Fanello. 2021. Total relighting: learning to relight portraits for background replacement. *ACM Trans. Graph.* 40, 4 (2021), 43–1.
- [258] Marios Papas, Christian Regg, Wojciech Jarosz, Bernd Bickel, Philip Jackson, Wojciech Matusik, Steve Marschner, and Markus Gross. 2013. Fabricating translucent materials using continuous pigment mixtures. *ACM Trans. Graph. (TOG)* 32, 4 (2013), 1–12.
- [259] Sylvain Paris, Hector M Briceno, and François X Sillion. 2004. Capture of hair geometry from multiple images. *ACM Trans. Graph. (TOG)* 23, 3 (2004), 712–719.
- [260] Sylvain Paris, Will Chang, Oleg I Kozhushnyan, Wojciech Jarosz, Wojciech Matusik, Matthias Zwicker, and Frédo Durand. 2008. Hair photobooth: geometric and photometric acquisition of real hairstyles. *ACM Trans. Graph.* 27, 3 (2008), 30.
- [261] Deborah Patrick and Shaun Bangay. 2003. A lightweight 3D plug-in for modeling long hair on virtual humans. In *Proceedings of the 2nd International Conference on Computer Graphics, Virtual Reality, Visualisation and Interaction in Africa* (Cape Town, South Africa) (AFRIGRAPH ’03). Association for Computing Machinery, New York, NY, USA, 161–187. doi:10.1145/602330.602360
- [262] Deborah Patrick, Shaun Bangay, and Adele Lobb. 2004. Modelling and rendering techniques for african hairstyles. In *Proceedings of Afrigraph*. 115–124.
- [263] Amandalynne Paullada, Inioluwa Deborah Raji, Emily M Bender, Emily Denton, and Alex Hanna. 2021. Data and its (dis) contents: A survey of dataset development and use in machine learning research. *Patterns* 2, 11 (2021).
- [264] Pieter Peers, Naoki Tamura, Wojciech Matusik, and Paul Debevec. 2007. Post-production facial performance relighting using reflectance transfer. *ACM Trans. Graph. (TOG)* 26, 3 (2007), 52–es.
- [265] Fabio Pellacini, Frank Battaglia, R Keith Morley, and Adam Finkelstein. 2007. Lighting with paint. *ACM Trans. Graph. (TOG)* 26, 2 (2007), 9–es.
- [266] Ken Perlin and Eric M Hoffert. 1989. Hypertexture. In *Proceedings of SIGGRAPH*. 253–262.
- [267] Amanda Phillips. 2020. Gamer trouble: Feminist confrontations in digital culture. In *Gamer Trouble*. New York University Press.
- [268] Frédéric Pighin, Jamie Hecker, Dani Lischinski, Richard Szeliski, and David H Salesin. 1998. Synthesizing realistic facial expressions from photographs. In *Proceedings of SIGGRAPH*.
- [269] Michal Piovarci, Alexandre Chapiro, and Bernd Bickel. 2023. Skin-screen: A computational fabrication framework for color tattoos. *ACM Trans. Graph. (TOG)* 42, 4 (2023), 1–13.

- [270] Chanda Prescod-Weinstein. 2021. *The Disordered Cosmos: A Journey into Dark Matter, Spacetime, and Dreams Deferred*. Bold Type Books.
- [271] Zesong Qiu, Yuwei Li, Dongming He, Qixuan Zhang, Longwen Zhang, Yinghao Zhang, Jingya Wang, Lan Xu, Xudong Wang, Yuyao Zhang, et al. 2022. SCULPTOR: Skeleton-consistent face creation using a learned parametric generator. *ACM Trans. Graph. (TOG)* 41, 6 (2022), 1–17.
- [272] Ante Qu and Doug L James. 2021. Fast linking numbers for topology verification of loopy structures. *ACM Trans. Graph.* 40, 4 (2021), 106–1.
- [273] Jonathan Ragan-Kelley, Charlie Kilpatrick, Brian W Smith, Doug Epps, Paul Green, Christophe Hery, and Frédo Durand. 2007. The lightspeed automatic interactive lighting preview system. In *ACM SIGGRAPH 2007 papers*. 25–es.
- [274] Pramod Rao, Gereon Fox, Abhimitra Meka, Mallikarjun BR, Fangneng Zhan, Tim Weyrich, Bernd Bickel, Hanspeter Pfister, Wojciech Matusik, Mohamed Elgharib, et al. 2024. Lite2Relight: 3D-aware single image portrait relighting. In *Proceedings of SIGGRAPH*. 1–12.
- [275] Ayesha Rascoe. 2025. An animation breakthrough makes it possible to more accurately illustrate Black hair. *NPR* (2025).
- [276] Victor Ray. 2023. *On critical race theory: Why it matters & why you should care*. Random House Trade Paperbacks.
- [277] Stephane Redon, Nico Galoppo, and Ming C Lin. 2005. Adaptive Dynamics of Articulated Bodies. *ACM Trans. Graph.* 24, 3 (2005), 936–945.
- [278] Zhong Ren, Kun Zhou, Tengfei Li, Wei Hua, and Baining Guo. 2010. Interactive hair rendering under environment lighting. *ACM Trans. Graph. (TOG)* 29, 4 (2010), 1–8.
- [279] Alexander Reshetov and David Hart. 2024. Modeling Hair Strands with Roving Capsules. In *Proceedings of SIGGRAPH (Denver, CO, USA) (SIGGRAPH '24)*. Association for Computing Machinery, New York, NY, USA, Article 62, 9 pages. doi:10.1145/3641519.3657450
- [280] Daniel Restuccio. 2004. Visual effects - Sony Pictures Imageworks Takes Spider-man 2 to New Heights. *Post Magazine* (2004).
- [281] Jérémy Riviere, Paulo FU Gotardo, Derek Bradley, Abhijeet Ghosh, and Thabo Beeler. 2020. Single-shot high-quality facial geometry and skin appearance capture. *ACM Trans. Graph.* 39, 4 (2020), 81.
- [282] Phillip Rogaway. 2015. The moral character of cryptographic work. *Cryptology ePrint Archive* (2015).
- [283] Lorna Roth. 2009. Looking at Shirley, the ultimate norm: Colour balance, image technologies, and cognitive equity. *Canadian Journal of Communication* 34, 1 (2009), 111–136.
- [284] Wendy D Roth. 2016. The multiple dimensions of race. *Ethnic and Racial Studies* 39, 8 (2016), 1310–1338.
- [285] Melena Ryzik. 2026. The ‘Sinners’ Cinematographer Says They ‘All Had a Lot on the Line’. *New York Times* (2026).
- [286] Iman Sadeghi, Oleg Bisker, Joachim De Deken, and Henrik Wann Jensen. 2013. A practical microcylinder appearance model for cloth rendering. *ACM Trans. Graph. (TOG)* 32, 2 (2013), 1–12.
- [287] Iman Sadeghi, Heather Pritchett, Henrik Wann Jensen, and Rasmus Tamstorf. 2010. An artist friendly hair shading system. *ACM Trans. Graph. (TOG)* 29, 4 (2010), 1–10.
- [288] Shunsuke Saito, Liwen Hu, Chongyang Ma, Hikaru Ibayashi, Linjie Luo, and Hao Li. 2018. 3D hair synthesis using volumetric variational autoencoders. *ACM Trans. Graph. (TOG)* 37, 6 (2018), 1–12.
- [289] Shunsuke Saito, Zi-Ye Zhou, and Ladislav Kavan. 2015. Computational bodybuilding: anatomically-based modeling of human bodies. *ACM Trans. Graph.* 34, 4, Article 41 (July 2015), 12 pages. doi:10.1145/2766957
- [290] Shunsuke Saito, Zi-Ye Zhou, and Ladislav Kavan. 2015. Computational bodybuilding: Anatomically-based modeling of human bodies. *ACM Trans. Graph. (TOG)* 34, 4 (2015), 1–12.
- [291] Michael P Salisbury, Michael T Wong, John F Hughes, and David H Salesin. 1997. Orientable textures for image-based pen-and-ink illustration. In *Proceedings of SIGGRAPH*. 401–406.
- [292] L Ayu Saraswati. 2020. Cosmopolitan whiteness: The effects and affects of skin-whitening advertisements in a transnational women’s magazine in Indonesia. *Meridians* 19, S1 (2020), 363–388.
- [293] Kripasindhu Sarkar, Marcel C Bühler, Gengyan Li, Daoye Wang, Delio Vicini, Jérémy Riviere, Yinda Zhang, Sergio Orts-Escolano, Paulo Gotardo, Thabo Beeler, et al. 2023. Litnerf: Intrinsic radiance decomposition for high-quality view synthesis and relighting of faces. In *SIGGRAPH Asia 2023 Conference Papers*. 1–11.
- [294] Darcy Schild. 2020. Fans are applauding Simone Biles’ Vogue cover story, but many argue her photos are proof that magazines need to hire Black photographers. *Business Insider* (2020).
- [295] Gabriel Schwartz, Shih-En Wei, Te-Li Wang, Stephen Lombardi, Tomas Simon, Jason Saragih, and Yaser Sheikh. 2020. The eyes have it: An integrated eye and face model for photorealistic facial animation. *ACM Trans. Graph. (TOG)* 39, 4 (2020), 91–1.
- [296] Andrew Selle, Michael Lentine, and Ronald Fedkiw. 2008. A mass spring model for hair simulation. *ACM Trans. Graph.* 27, 3 (aug 2008), 1–11.
- [297] Jared Sexton. 2008. *Amalgamation schemes: Antiracism and the critique of multiracialism*. U of Minnesota Press.
- [298] Yuefan Shen, Shunsuke Saito, Ziyang Wang, Olivier Mauray, Chenglei Wu, Jessica Hodgins, Youyi Zheng, and Giljoo Nam. 2023. CT2Hair: High-Fidelity 3D Hair Modeling using Computed Tomography. *ACM Trans. Graph.* 42, 4, Article 75 (July 2023), 13 pages. doi:10.1145/3592106

- [299] Alvin Shi, Haomiao Wu, Jarred Parr, A. M. Darke, and Theodore Kim. 2023. Lifted Curls: A Model for Tightly Coiled Hair Simulation. *Proc. ACM Comput. Graph. Interact. Tech.* 6, 3, Article 42 (aug 2023), 19 pages. doi:10.1145/3606920
- [300] Fuhao Shi, Hsiang-Tao Wu, Xin Tong, and Jinxiang Chai. 2014. Automatic acquisition of high-fidelity facial performances using monocular videos. *ACM Trans. Graph. (TOG)* 33, 6 (2014), 1–13.
- [301] Ying Song, Xin Tong, Fabio Pellacini, and Pieter Peers. 2009. Subedit: a representation for editing measured heterogeneous subsurface scattering. *ACM Trans. Graph. (TOG)* 28, 3 (2009), 1–10.
- [302] Florian Struck, Christian-A. Bohn, Sebastian Schmidt, and Volker Helzle. 2004. Realistic shading of human skin in real time. In *Proceedings of the 3rd International Conference on Computer Graphics, Virtual Reality, Visualisation and Interaction in Africa* (Stellenbosch, South Africa) (*AFRIGRAPH '04*). Association for Computing Machinery, New York, NY, USA, 93–97. doi:10.1145/1029949.1029967
- [303] Shinjiro Sueda, Garrett L. Jones, David IW Levin, and Dinesh K. Pai. 2011. Large-scale dynamic simulation of highly constrained strands. *ACM Trans. Graph. (TOG)* 30, 4 (2011), 1–10.
- [304] Jingxiang Sun, Xuan Wang, Yichun Shi, Lizhen Wang, Jue Wang, and Yebin Liu. 2022. Ide-3d: Interactive disentangled editing for high-resolution 3d-aware portrait synthesis. *ACM Trans. Graph. (TOG)* 41, 6 (2022), 1–10.
- [305] Tiancheng Sun, Zexiang Xu, Xiuming Zhang, Sean Fanello, Christoph Rhemann, Paul Debevec, Yun-Ta Tsai, Jonathan T. Barron, and Ravi Ramamoorthi. 2020. Light stage super-resolution: continuous high-frequency relighting. *ACM Trans. Graph. (TOG)* 39, 6 (2020), 1–12.
- [306] Eric Tabellion and Arnauld Lamorlette. 2004. An approximate global illumination system for computer generated films. *ACM Trans. Graph. (TOG)* 23, 3 (2004), 469–476.
- [307] Tolga Tasdizen, Ross Whitaker, Paul Burchard, and Stanley Osher. 2003. Geometric surface processing via normal maps. *ACM Trans. Graph. (TOG)* 22, 4 (2003), 1012–1033.
- [308] Kartik Teotia, Mallikarjun B. R., Xingang Pan, Hyeonwoo Kim, Pablo Garrido, Mohamed Elgharib, and Christian Theobalt. 2024. Hq3davatar: High-quality implicit 3d head avatar. *ACM Trans. Graph.* 43, 3 (2024), 1–24.
- [309] Raqi Syed, Derek Nowrouzezahrai, Theodore Kim, Holly Rushmeier. 2022. Countering Racial Bias in Computer Graphics Research (the BOF): One Year Later. In *ACM SIGGRAPH Birds of a Feather*. ACM, 1–2.
- [310] Justus Thies, Michael Zollhöfer, Matthias Nießner, Levi Valgaerts, Marc Stamminger, and Christian Theobalt. 2015. Real-time expression transfer for facial reenactment. *ACM Trans. Graph.* 34, 6 (2015), 183–1.
- [311] Justus Thies, Michael Zollhöfer, Marc Stamminger, Christian Theobalt, and Matthias Nießner. 2018. FaceVR: Real-Time Gaze-Aware Facial Reenactment in Virtual Reality. *ACM Trans. Graph. (TOG)* 37, 2 (2018), 1–15.
- [312] Lynn M. Thomas. 2020. *Beneath the surface: A transnational history of skin lighteners*. Duke University Press.
- [313] Krista A. Thompson. 2015. *Shine: The visual economy of light in African diasporic aesthetic practice*. Duke University Press.
- [314] Sonia Thompson. 2020. 3 Important Lessons In Inclusive Marketing From Annie Liebovitz' Controversial Simone Biles Vogue Covers. *Forbes* (2020).
- [315] Nils Thürey, Chris Wojtan, Markus Gross, and Greg Turk. 2010. A multiscale approach to mesh-based surface tension flows. *ACM Trans. Graph. (TOG)* 29, 4 (2010), 1–10.
- [316] Marc Thyng, Christopher Evart, Toby Jones, and Aleka McAdams. 2017. The art and technology of hair simulation in Disney's Moana. In *ACM SIGGRAPH Talks*. 1–2.
- [317] Kenneth E. Torrance and Ephraim M. Sparrow. 1967. Theory for off-specular reflection from roughened surfaces. *Journal of the Optical Society of America* 57, 9 (1967), 1105–1114.
- [318] Vilna Bashi Treitler. 2013. *The ethnic project: Transforming racial fiction into ethnic factions*. Stanford University Press.
- [319] Norimichi Tsumura, Nobutoshi Ojima, Kayoko Sato, Mitsuhiro Shiraiishi, Hideto Shimizu, Hirohide Nabeshima, Syuuichi Akazaki, Kimihiko Hori, and Yoichi Miyake. 2003. Image-based skin color and texture analysis/synthesis by extracting hemoglobin and melanin information in the skin. *ACM Trans. Graph. (TOG)* 22, 3 (2003), 770–779.
- [320] Levi Valgaerts, Chenglei Wu, Andrés Bruhn, Hans-Peter Seidel, and Christian Theobalt. 2012. Lightweight binocular facial performance capture under uncontrolled lighting. *ACM Trans. Graph.* 31, 6 (2012), 187–1.
- [321] Gul Varol, Javier Romero, Xavier Martin, Naureen Mahmood, Michael J. Black, Ivan Laptev, and Cordelia Schmid. 2017. Learning from synthetic humans. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 109–117.
- [322] Edgar Velázquez-Armendáriz, Shuang Zhao, Miloš Hašan, Bruce Walter, and Kavita Bala. 2009. Automatic bounding of programmable shaders for efficient global illumination. *ACM Trans. Graph. (TOG)* 28, 5 (2009), 1–9.
- [323] Delio Vicini, Vladlen Koltun, and Wenzel Jakob. 2019. A learned shape-adaptive subsurface scattering model. *ACM Trans. Graph. (TOG)* 38, 4 (2019), 1–15.
- [324] Alexander Vilesov, Pradyumna Chari, Adnan Armouti, Anirudh Bindiganavale, Harish, Kimaya Kulkarni, Ananya Deoghare, Laleh Jalilian, and Achuta Kadambi. 2022. Blending camera and 77 GHz radar sensing for equitable, robust plethysmography. *ACM Trans. Graph.* 41, 4 (2022), 36–1.
- [325] Felix von Luschan. 1897. *Beiträge zur Völkerkunde der deutschen Schutzgebiete*. D. Reimer.
- [326] Andre Walker. 1997. *Andre Talks Hair*. Simon & Schuster.

- [327] Cong Wang, Di Kang, Yan-Pei Cao, Linchao Bao, Ying Shan, and Song-Hai Zhang. 2023. Neural point-based volumetric avatar: Surface-guided neural points for efficient and photorealistic volumetric head avatar. In *SIGGRAPH Asia 2023 Conference Papers*. 1–12.
- [328] Daoye Wang, Prashanth Chandran, Gaspard Zoss, Derek Bradley, and Paulo Gotardo. 2022. Morf: Morphable radiance fields for multiview neural head modeling. In *ACM SIGGRAPH 2022 Conference Proceedings*. 1–9.
- [329] Lvdi Wang, Yizhou Yu, Kun Zhou, and Baining Guo. 2009. Example-based hair geometry synthesis. In *Proceedings of SIGGRAPH*. 1–9.
- [330] Tianyu Wang, Jiong Chen, Dongping Li, Xiaowei Liu, Huamin Wang, and Kun Zhou. 2023. Fast GPU-based two-way continuous collision handling. *ACM Trans. Graph.* 42, 5 (2023), 1–15.
- [331] Zian Wang, Tianchang Shen, Merlin Nimier-David, Nicholas Sharp, Jun Gao, Alexander Keller, Sanja Fidler, Thomas Müller, and Zan Gojcic. 2023. Adaptive Shells for Efficient Neural Radiance Field Rendering. *ACM Trans. Graph.* (2023).
- [332] Zhibo Wang, Xin Yu, Ming Lu, Quan Wang, Chen Qian, and Feng Xu. 2020. Single image portrait relighting via explicit multiple reflectance channel modeling. *ACM Trans. Graph. (TOG)* 39, 6 (2020), 1–13.
- [333] Alan Warburton. 2017. *Goodbye Uncanny Valley*. <https://alanwarburton.co.uk/goodbye-uncanny-valley>
- [334] Keith Waters. 1987. A muscle model for animation three-dimensional facial expression. *Proceedings of SIGGRAPH* 21, 4 (1987), 17–24.
- [335] Yichen Wei, Eyal Ofek, Long Quan, and Heung-Yeung Shum. 2005. Modeling hair from multiple views. *ACM Trans. Graph. (TOG)* 24, 3 (2005), 816–820.
- [336] Andreas Wenger, Andrew Gardner, Chris Tchou, Jonas Unger, Tim Hawkins, and Paul Debevec. 2005. Performance relighting and reflectance transformation with time-multiplexed illumination. *ACM Trans. Graph. (TOG)* 24, 3 (2005), 756–764.
- [337] Tim Weyrich, Wojciech Matusik, Hanspeter Pfister, Bernd Bickel, Craig Donner, Chien Tu, Janet McAndless, Jinho Lee, Addy Ngan, Henrik Wann Jensen, et al. 2006. Analysis of human faces using a measurement-based skin reflectance model. *ACM Trans. Graph. (TOG)* 25, 3 (2006), 1013–1024.
- [338] Chris Wiggum, Michael Fong, Britta Wilson, Mara MacMahon, Jaclyn Simon, Sean Muriithi, Bryn Imagire, Ian Megibben, and Junyi Ling. 2021. Pixar’s soul and the search for authenticity. In *ACM SIGGRAPH 2021 Production Sessions*. 1–1.
- [339] Lance Williams. 1990. Performance-driven facial animation. In *Proceedings of SIGGRAPH*. 235–242.
- [340] Sebastian Winberg, Gaspard Zoss, Prashanth Chandran, Paulo Gotardo, and Derek Bradley. 2022. Facial hair tracking for high fidelity performance capture. *ACM Trans. Graph. (TOG)* 41, 4 (2022), 1–12.
- [341] Erroll Wood, Tadas Baltrušaitis, Charlie Hewitt, Sebastian Dziadzio, Thomas J Cashman, and Jamie Shotton. 2021. Fake it till you make it: face analysis in the wild using synthetic data alone. In *Proceedings of the IEEE/CVF international conference on computer vision*. 3681–3691.
- [342] Computer Graphics World. 2015. *Highlights from SIGGRAPH 2015*. <https://www.cgw.com/Press-Center/Siggraph/2015/Highlights-from-SIGGRAPH-2015.aspx>
- [343] Franz J Wortmann, Gabriele Wortmann, and Therakanya Sripho. 2020. Why is hair curly?—Deductions from the structure and the biomechanics of the mature hair shaft. *Experimental dermatology* 29, 3 (2020), 366–372.
- [344] Chenglei Wu, Derek Bradley, Markus Gross, and Thabo Beeler. 2016. An anatomically-constrained local deformation model for monocular face capture. *ACM Trans. Graph. (TOG)* 35, 4 (2016), 1–12.
- [345] Chenglei Wu, Takaaki Shiratori, and Yaser Sheikh. 2018. Deep incremental learning for efficient high-fidelity face tracking. *ACM Trans. Graph. (TOG)* 37, 6 (2018), 1–12.
- [346] Haomiao Wu, Alvin Shi, A.M. Darke, and Theodore Kim. 2024. Curly-Cue: Geometric Methods for Highly Coiled Hair. In *Proceedings of SIGGRAPH Asia (Tokyo, Japan) (SA '24)*. Association for Computing Machinery, New York, NY, USA, Article 112, 11 pages. doi:10.1145/3680528.3687641
- [347] Mengqi Xia, Bruce Walter, Christophe Hery, Olivier Maury, Eric Michielssen, and Steve Marschner. 2023. A practical wave optics reflection model for hair and fur. *ACM Trans. Graph. (TOG)* 42, 4 (2023), 1–15.
- [348] Mengqi Xia, Bruce Walter, Eric Michielssen, David Bindel, and Steve Marschner. 2020. A wave optics based fiber scattering model. *ACM Trans. Graph. (TOG)* 39, 6 (2020), 1–16.
- [349] Chufeng Xiao, Deng Yu, Xiaoguang Han, Youyi Zheng, and Hongbo Fu. 2021. SketchHairSalon: deep sketch-based hair image synthesis. *ACM Trans. Graph. (TOG)* 40, 6 (2021), 1–16.
- [350] Minshan Xie, Chengze Li, Xueting Liu, and Tien-Tsin Wong. 2020. Manga filling style conversion with screentone variational autoencoder. *ACM Trans. Graph. (TOG)* 39, 6 (2020), 1–15.
- [351] Bing Xu, Junfei Zhang, Rui Wang, Kun Xu, Yong-Liang Yang, Chuan Li, and Rui Tang. 2019. Adversarial Monte Carlo denoising with conditioned auxiliary feature modulation. *ACM Trans. Graph.* 38, 6 (2019), 224–1.
- [352] Kun Xu, Li-Qian Ma, Bo Ren, Rui Wang, and Shi-Min Hu. 2011. Interactive hair rendering and appearance editing under environment lighting. *ACM Trans. Graph. (TOG)* 30, 6 (2011), 1–10.
- [353] Yuelang Xu, Lizhen Wang, Xiaochen Zhao, Hongwen Zhang, and Yebin Liu. 2023. Avatarmav: Fast 3d head avatar reconstruction using motion-aware neural voxels. In *ACM SIGGRAPH 2023 Conference Proceedings*. 1–10.
- [354] Zexiang Xu, Hsiang-Tao Wu, Lvdi Wang, Changxi Zheng, Xin Tong, and Yue Qi. 2014. Dynamic hair capture using spacetime optimization. *ACM Trans. Graph. (TOG)* 33, 6 (2014), 1–11.

- [355] Shugo Yamaguchi, Shunsuke Saito, Koki Nagano, Yajie Zhao, Weikai Chen, Kyle Olszewski, Shigeo Morishima, and Hao Li. 2018. High-fidelity facial reflectance and geometry inference from an unconstrained image. *ACM Trans. Graph. (TOG)* 37, 4 (2018), 1–14.
- [356] Ling-Qi Yan, Henrik Wann Jensen, and Ravi Ramamoorthi. 2017. An efficient and practical near and far field fur reflectance model. *ACM Trans. Graph. (TOG)* 36, 4 (2017), 1–13.
- [357] Ling-Qi Yan, Weilun Sun, Henrik Wann Jensen, and Ravi Ramamoorthi. 2017. A BSSRDF model for efficient rendering of fur with global illumination. *ACM Trans. Graph. (TOG)* 36, 6 (2017), 1–13.
- [358] Ling-Qi Yan, Chi-Wei Tseng, Henrik Wann Jensen, and Ravi Ramamoorthi. 2015. Physically-accurate fur reflectance: Modeling, measurement and rendering. *ACM Trans. Graph. (TOG)* 34, 6 (2015), 1–13.
- [359] Haotian Yang, Mingwu Zheng, Wanquan Feng, Haibin Huang, Yu-Kun Lai, Pengfei Wan, Zhongyuan Wang, and Chongyang Ma. 2023. Towards practical capture of high-fidelity relightable avatars. In *SIGGRAPH Asia 2023 Conference Papers*. 1–11.
- [360] Haotian Yang, Mingwu Zheng, Chongyang Ma, Yu-Kun Lai, Pengfei Wan, and Haibin Huang. 2024. Vmmm: A volumetric relightable morphable head model. In *Proceedings of SIGGRAPH*. 1–11.
- [361] Lingchen Yang, Gaspard Zoss, Prashanth Chandran, Markus Gross, Barbara Solenthaler, Eftychios Sifakis, and Derek Bradley. 2024. Learning a Generalized Physical Face Model From Data. *ACM Trans. Graph. (TOG)* 43, 4 (2024), 1–14.
- [362] Yu-Ying Yeh, Koki Nagano, Sameh Khamis, Jan Kautz, Ming-Yu Liu, and Ting-Chun Wang. 2022. Learning to relight portrait images via a virtual light stage and synthetic-to-real adaptation. *ACM Trans. Graph. (TOG)* 41, 6 (2022), 1–21.
- [363] Kanamori Yoshihiro. 2018. Relighting humans: occlusion-aware inverse rendering for full-body human images. *ACM Trans. Graph.* 37 (2018), 270–1.
- [364] Genevieve Yue. 2020. *Girl Head: Feminism and Film Materiality*. Fordham University Press.
- [365] Cem Yuksel, Scott Schaefer, and John Keyser. 2009. Hair meshes. *ACM Trans. Graph.* 28, 5 (Dec. 2009), 1–7. doi:10.1145/1618452.1618512
- [366] Tizian Zeltner and Wenzel Jakob. 2018. The layer laboratory: a calculus for additive and subtractive composition of anisotropic surface reflectance. *ACM Trans. Graph. (TOG)* 37, 4 (2018), 1–14.
- [367] Longwen Zhang, Qiwei Qiu, Hongyang Lin, Qixuan Zhang, Cheng Shi, Wei Yang, Ye Shi, Sibe Yang, Lan Xu, and Jingyi Yu. 2023. DreamFace: Progressive Generation of Animatable 3D Faces under Text Guidance. *ACM Trans. Graph. (TOG)* 42, 4 (2023), 1–16.
- [368] Longwen Zhang, Chuxiao Zeng, Qixuan Zhang, Hongyang Lin, Ruixiang Cao, Wei Yang, Lan Xu, and Jingyi Yu. 2022. Video-driven neural physically-based facial asset for production. *ACM Trans. Graph. (TOG)* 41, 6 (2022), 1–16.
- [369] Longwen Zhang, Zijun Zhao, Xinzhou Cong, Qixuan Zhang, Shuqi Gu, Yuchong Gao, Rui Zheng, Wei Yang, Lan Xu, and Jingyi Yu. 2023. Hack: Learning a parametric head and neck model for high-fidelity animation. *ACM Trans. Graph. (TOG)* 42, 4 (2023), 1–20.
- [370] Meng Zhang, Menglei Chai, Hongzhi Wu, Hao Yang, and Kun Zhou. 2017. A data-driven approach to four-view image-based hair modeling. *ACM Trans. Graph.* 36, 4, Article 156 (July 2017), 11 pages.
- [371] Meng Zhang, Pan Wu, Hongzhi Wu, Yanlin Weng, Youyi Zheng, and Kun Zhou. 2018. Modeling hair from an RGB-D camera. *ACM Trans. Graph.* 37, 6, Article 205 (Dec. 2018), 10 pages. doi:10.1145/3272127.3275039
- [372] Xiuming Zhang, Sean Fanello, Yun-Ta Tsai, Tiancheng Sun, Tianfan Xue, Rohit Pandey, Sergio Orts-Escolano, Philip Davidson, Christoph Rhemann, Paul Debevec, et al. 2021. Neural light transport for relighting and view synthesis. *ACM Trans. Graph. (TOG)* 40, 1 (2021), 1–17.
- [373] Shuang Zhao, Ravi Ramamoorthi, and Kavita Bala. 2014. High-order similarity relations in radiative transfer. *ACM Trans. Graph. (TOG)* 33, 4 (2014), 1–12.
- [374] Xiaochen Zhao, Jingxiang Sun, Lizhen Wang, Jinli Suo, and Yebin Liu. 2024. Invertavatar: Incremental gan inversion for generalized head avatars. In *ACM SIGGRAPH 2024 Conference Papers*. 1–10.
- [375] Shaokun Zheng, Fengshi Zheng, Kun Xu, and Ling-Qi Yan. 2021. Ensemble denoising for Monte Carlo renderings. *ACM Trans. Graph. (TOG)* 40, 6 (2021), 1–17.
- [376] Yujian Zheng, Yuda Qiu, Leyang Jin, Chongyang Ma, Haibin Huang, Di Zhang, Pengfei Wan, and Xiaoguang Han. 2024. Towards Unified 3D Hair Reconstruction from Single-View Portraits. In *Proceedings of SIGGRAPH Asia (Tokyo, Japan) (SA '24)*. Association for Computing Machinery, New York, NY, USA, Article 114, 11 pages. doi:10.1145/3680528.3687597
- [377] Kun Zhou, Qiming Hou, Zhong Ren, Minmin Gong, Xin Sun, and Baining Guo. 2009. RenderAnts: Interactive Reyes Rendering on GPUs. *ACM Trans. Graph.* (2009).
- [378] Yuxiao Zhou, Menglei Chai, Alessandro Pepe, Markus Gross, and Thabo Beeler. 2023. Groomgen: A high-quality generative hair model using hierarchical latent representations. *ACM Trans. Graph. (TOG)* 42, 6 (2023), 1–16.
- [379] ChangAn Zhu and Chris Joslin. 2024. Fabrig: A Cloth-Simulated Transferable 3D Face Parameterization. In *SIGGRAPH Asia 2024 Conference Papers*. 1–10.
- [380] Junqiu Zhu, Sizhe Zhao, Lu Wang, Yanning Xu, and Ling-Qi Yan. 2022. Practical level-of-detail aggregation of fur appearance. *ACM Trans. Graph. (TOG)* 41, 4 (2022), 1–17.
- [381] Arno Zinke, Martin Rump, Tomás Lay, Andreas Weber, Anton Andriyenko, and Reinhard Klein. 2009. A practical approach for photometric acquisition of hair color. *ACM Trans. Graph. (TOG)* 28, 5 (2009), 1–9.

- [382] Arno Zinke, Cem Yuksel, Andreas Weber, and John Keyser. 2008. Dual scattering approximation for fast multiple scattering in hair. *ACM Trans. Graph.* 27, 3 (2008), 1–10.
- [383] Gaspard Zoss, Thabo Beeler, Markus Gross, and Derek Bradley. 2019. Accurate markerless jaw tracking for facial performance capture. *ACM Trans. Graph. (TOG)* 38, 4 (2019), 1–8.
- [384] Gaspard Zoss, Derek Bradley, Pascal Bérard, and Thabo Beeler. 2018. An empirical rig for jaw animation. *ACM Trans. Graph. (TOG)* 37, 4 (2018), 1–12.
- [385] Gaspard Zoss, Prashanth Chandran, Eftychios Sifakis, Markus Gross, Paulo Gotardo, and Derek Bradley. 2022. Production-ready face re-aging for visual effects. *ACM Trans. Graph. (TOG)* 41, 6 (2022), 1–12.

A Methods

A.1 Data Processing

Through an agreement with our university library, we have been able to obtain the full text and metadata of every *SIGGRAPH*, *SIGGRAPH Asia*, and *ACM Transactions on Graphics* article from 1974 to 2024 from the ACM Digital Library. The data consisted of DTD BITS and JATS paper texts, programs, authors, affiliations and CCS classifications, among other data, for all papers in every issue of *ACM Transactions on Graphics* and the *SIGGRAPH* proceedings. Before analyzing the data, we performed three stages of manual data cleaning and labeling, which we refer to as *Splitting*, *Clustering*, and *Classifying*.

Splitting. The dataset schema allows for multiple listed affiliations per author, but we found that roughly one-fifth of affiliations were listed as a single entry despite referring to two or more institutions (e.g. “*Company A and University B*”). These compound affiliation strings use no consistent delimiter and are often indistinguishable from non-compound affiliations without institution-specific knowledge, precluding automated splitting.

Instead, we performed splitting with a two-stage semi-automated approach. First, the 6,610 unique affiliation strings from the papers dataset were categorized as either compound or non-compound. Each of the compound affiliation strings was then split by the human reviewer into its component parts, using an interface that highlighted likely conjunctions but that allowed the human labeler to split the string into arbitrary elements. All affiliation strings were split by the same human reviewer.

Clustering. Splitting reduced the total number of unique affiliation strings from 6,610 to 5,913. As many of these referred to the same institution under different names (e.g., “*EPFL*” and “*EPF Lausanne*”), we sought to cluster affiliation strings that refer to individual institutions. As an exhaustive performance of this task would require $N^2 > 3e7$ human string comparisons, several heuristics were used to produce approximate clusterings.

For each unique affiliation string, a human labeler was shown a list of key collisions using three different text collision-finding algorithms from the OpenRefine project (FINGERPRINT, NGRAMFINGERPRINT, and METAPHONE3 keyers). The human labeler also performed text searches of all affiliation records to select other results for clustering. After this process was completed, any overlapping clusters were manually reviewed and either combined or disentangled to produce a set of 1,653 disjoint clusters of affiliations.

Classifying Country Codes. For each individual affiliation string within the clusters, we then determine the approximate geographic location. Each post-split affiliation string was first geocoded using the Google Maps Geocoding API to assign an ISO 3166-1 country code. Each affiliation/location tag was then examined by a human reviewer, who verified the geocoding and made corrections where necessary. The Geocoding API returned a result for roughly 90% of the affiliations, and our human review pass found an approximately 2% error rate among the returned country codes. All reviewers were graphics researchers who had familiarity with the field and the publishing institutions. The results of this data analysis was used to generate Figure 1.

B Results

B.1 Skin Papers

B.1.1 Skin as an optical medium. The *Manuka* [103] paper describes the technical choices needed to “render high-quality skin by path-traced subsurface scattering.” The paper tags two images as “skin” renderings. One is of *Digital Emily* [4], a photorealistic rendering of a white actress, and one of pale, iridescent aliens from the film *Valerian and the City of Thousand Planets* (2017).

The *RenderMan* [64] paper describes its strategy for rendering “skin and other translucent materials” using subsurface scattering and shows images of Arnold Schwarzenegger from *Terminator Genisys* (2015), Peter Cushing from *Rogue One: A Star Wars Story* (2016) and Sean Young from *Blade Runner: 2049* (2017).

The *Hyperion* paper [48] draws an explicit equivalence between skin and snow in the section *Unifying Subsurface Scattering, from Snow to Skin*. The skin algorithm started as experiments for the snow monster Marshmallow in the short *Frozen Fever* (2015) and the snowscapes in *Olaf’s Frozen Adventure* (2017). With the success of these experiments, “path-traced subsurface scattering is now being used in all current productions on all materials from snow to skin.” The one example in the paper labeled as “skin” shows Ralph and Vanellope from *Ralph Breaks the Internet: Wreck-It Ralph 2* (2018). Both characters are white. Darker-skinned characters from *Moana* (2016) also appear in the paper, but as showcases for the engine’s denoising and ocean rendering capabilities.

B.1.2 Skin as a layered material.

- A 2003 paper [319] on hemoglobin and melanin only demonstrates its algorithm on a pale-skinned East Asian faces. The authors are all based in Japan.
- A 2005 paper [91] on multi-layered translucent materials models skin using 6 different layers, and ends with a single rendering of a white face.
- A 2008 paper [92] proposes a “layered, heterogenous” model for skin, and renders test swatches of skin that include one instance of “African, skin type V”, in reference to the Fitzpatrick skin typing system [108] from §3. The final showcase renders of human hands and an ear are all of white skin.
- A 2010 paper [169] on dynamic facial color only demonstrates results on a white model. The data was collected from “one Caucasian female, 33 years old; three Caucasian males, 26, 33, and 35 years old”. A “skin color lookup texture” is proposed to alter the model’s melanin, but is never applied on a facial render.
- A 2015 paper [61] only showcases results on white models. The possibility of darker skin due to different melanin levels is mentioned but not demonstrated.
- A 2015 paper [245] on capturing the microstructure of skin, e.g. the forehead wrinkles that form under frowning, only shows results on white skin.

B.1.3 Skin as a Human Face.

Facial Capture. Of the 67 papers we examined:

- 18 contained only whites. [7, 21–23, 28, 29, 37, 43, 117, 128, 161, 268, 310, 311, 320, 334, 344, 384] The papers spanned the 1980s [334], 1990s [128], and 2010s [310]. The most recent appeared in 2023. [7]
- 10 contained only whites and East Asians. [50, 53, 63, 112, 121, 217, 253, 300, 379, 383] Of these, 7 were from the 2010s, and the most recent from 2024. [63].
- 3 contained only whites and South Asians, and appeared in 2018, 2022 and 2023. [127, 327, 340] One 2007 paper [36] contained an example where our labellers could not agree on a South Asian or Black label. One conflating factor is that the paper was written by European authors, implying that the subject was also European. Thus, the North American labellers had difficulty parsing the European grooming aesthetics.
- 3 contained only whites and East Asians, plus Barack Obama. [98, 116, 118]
- 3 contained only East Asians [15, 216, 368], and all were published in the 2020s.

- 2 contained whites, East Asians, and South Asians [241, 361], and both were published in the 2020s.
- 1 contained solely Black examples [339] and was published in 1990 by a white researcher.
- 26 contained skin tones ranging white to Black. All appeared after 2010 and 85% after 2020. [10, 11, 20, 46, 49, 51, 52, 60, 95, 107, 156, 184, 185, 200, 206, 227, 246, 271, 281, 295, 308, 328, 359, 367, 369, 385]

Facial Reflectance and Relighting. The initial 2000 work of Debevec et al. [78] on the topic prominently features a dark skinned model. Its proposal to use a metallic model [317] to model skin glossiness was later adopted by Weyrich et al. [337]. Immediately following the publication of the influential 2001 subsurface scattering paper [167], reflectance field papers from Debevec in 2002 and 2005 only featured white skinned final renders [79, 336] (though [79] shows Black skin in a didactic figure) before later including East Asian [264] and then South Asian and Black [122] renders in 2007 and 2008. These coincided with the addition of East and South Asians on the author lists. Overall, we found:

- 3 papers feature white-only final renders, the most recent from 2024. [199, 225, 336] 1 additional paper includes a Latina woman [132], though the highly diffuse render suggests white aesthetics.
- 2 papers show whites and East Asians [264, 363], and while one of these [363] shows a Black example in their training set, there are no Black examples in their final evaluation. 1 additional paper [250] shows East and South Asians, and 1 additional paper features white, East Asians, and South Asians. [33]
- 12 papers show final renders across a range of skin tones. One is from 2000 [78], two are from the 2010s [196, 355], and the rest (75%) are from the 2020s [35, 142, 226, 274, 293, 305, 360, 362, 372].

B.2 Hair Results

Hair Geometry. Within the category of curly hair, we observed subcategories of *European locks* [72], which are straight at the root but curly near the tip, and *pageant hair* [187], which is straight hair with a high-radius, low-frequency twist artificially inserted near the tip. Pageant hair style is politically charged in the United States, where it is also called *Big Washington Hair* [170] and *Republican Hair* [81]. We found the following across 61 hair geometry papers:

- 17 papers present results exclusively on straight hair. Of these, one reports failing on “frizzy” hair [193], and two showed textured hair in the training set but not the final results [46, 185].
- 2 papers present results on straight hair, plus Barack Obama. [95, 353]
- 10 papers present results on straight and wavy hair, with 2 from 2024. [33, 57, 115, 184, 199, 203, 259, 260, 354, 370]
- 2 papers present results on straight and wavy hair [97, 153], as well as European locks.
- 5 papers [58, 154, 188, 288, 376] present results on straight and wavy hair, plus pageant hair. Of these, 3 papers were from China-based researchers [58, 188, 376], and 2 from the U.S. [154, 288]
- 2 papers deal with braids [155, 349], which exclusively means straight (predominantly blonde) hair braided into ponytails or pigtails. Textured hair styles such as box braids, cornrows, or locs did not appear.
- 19 papers show examples of straight, wavy, and curly hair, but the definition of “curly” is *highly* inconsistent, and often corresponded to lightly wavy hair. [11, 52, 56, 123, 145, 152, 182, 209, 212, 279, 298, 304, 308, 329, 335, 360, 365, 371, 374]
- 2 papers show textured hair as failure cases, one involving box braids [288] and the other “frizzy” hair [56].
- 2 papers contain textured hair, one from 2023 [378] and the other from 2024. [346]

Hair Simulation. We found the following across 35 hair simulation papers:

- 15 papers exclusively showed straight hair. They variously dealt with hair, thin solids, and multiphysics simulations (interaction of cloth, rods, solid volumes, and fluids). [8, 25, 32, 62, 75, 85, 104, 105, 168, 197, 201, 215, 223, 231, 330]

- 4 papers showed exclusively straight and wavy hair. [59, 106, 150, 174]
- 2 papers showed exclusively straight and curly hair. [73, 296]
- 2 paper showed straight, wavy and curly hair. [30, 74]
- 1 paper showed exclusively European locks. [55]
- 2 papers from 2023 and 2024 [70, 151] showed straight, wavy, curly, and textured hair.

Hair Rendering. We found the following hair categories across 64 hair rendering papers:

- 11 papers [18, 143, 165, 171, 220, 237, 273, 287, 293, 377] showing exclusively straight hair. One of these papers [143] shows a “hairball”, (see below), but it is a ball of straight hair, not the noisy hair ball that other algorithms are tested on. Another [165] labels an example as “wavy”, but labelers agreed that under the current scheme, it should be classified as “straight”. The example showed a half-period oscillation.
- 6 papers [177, 256, 278, 347, 352, 382] on straight and wavy hair, and 4 [9, 336, 351, 375] on wavy hair only
- 2 papers [238, 348] on straight and curly hair, and 1 paper [207] on curly hair only
- 3 papers [244, 332, 381] on straight, wavy, and curly hair. One paper [332] had an instance (Figure 10, bottom row) of potentially textured hair, but the hair was deliberately covered with a wrapping to exclude it from evaluation. Only small portions peeked out of the covering that might qualify as textured.
- 8 papers [12, 34, 35, 142, 208, 257, 331, 362] that show straight, wavy, curly, and textured hair. All but one [208] were published after 2021.

There were 9 papers [14, 66, 110, 133, 149, 222, 236, 266, 322] showing renderings of “fur”, referring either to very short hairs, or hairs on an animal. An additional 5 papers [126, 356–358, 380] dealt specifically with the problem of “fur” rendering, whose appearance is considered distinct from human hair. Another 5 papers [17, 54, 87, 141, 195] showed a standard “hairball” example, which appears to be straight hair with noise applied.

Received 13 January 2026; revised 25 March 2026; accepted 16 April 2026