Appearance

Rachel McDonnell and Bilge Mutlu

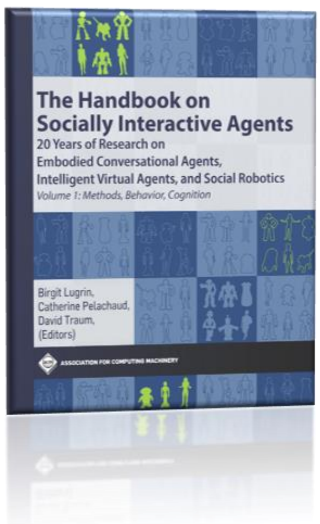

Author note:

This is a preprint. The final article is published in "The Handbook on Socially Interactive Agents" by ACM books.

Citation information:

McDonnell, R., and Mutlu, B. (2021) Appearance. In B. Lugrin, C. Pelachaud, D. Traum (Eds.), *Handbook on Socially Interactive Agents – 20 Years of Research on Embodied Conversational Agents, Intelligent Virtual Agents, and Social Robotics*, Volume 1: Methods, Behavior, Cognition (pp. 107-146). ACM books.

DOI of the final chapter: 10.1145/3477322.3477327

DOI of volume 1 of the handbook: 10.1145/3477322

Correspondence concerning this chapter should be addressed to Rachel McDonnell & Bilge Mutlu, ramcdonn@scss.tcd.ie, bilge@cs.wisc.edu

4 Appearance

Rachel McDonnell, Bilge Mutlu

4.1 Why appearance?

One might question why we need appearance for the metaphor to work, as voice assistants can effectively express characteristics of a metaphor solely through behavior. We argue that, although disembodied agents can effectively serve as computer-based assistants in specific scenarios of use, for example, involving driving and visually impaired users, appearance provides a "locus of attention" [Cassell 2001] for the cognitive and interactive faculties of the user of the system. Additionally, human communication mechanisms, such as mutual gaze, turn-taking, body orientation, necessitate the presence of appropriate visual cues to properly function, making appearance a necessity for agent design. Studies of human-human, human-agent, and human-robot interaction provide strong evidence that such mechanisms work more effectively when parties provide appearance-based cues. The mere presence of a form of embodiment in interacting with an agent improves social outcomes, such as motivation [Mumm and Mutlu 2011]. As the scale and modality of appearance get closer to that of the metaphor, these outcomes further improve; human-scale and physical agents have more perceived presence [Kiesler et al. 2008] and persuasive ability [Bainbridge et al. 2011] than scaled-down and virtual agents.

4.2 History

Agents with virtual and physical embodiments follow different historical trajectories. Virtual agents, also called embodied conversational agents, a term coined by Cassell [2000], are "computer-generated cartoonlike characters that demonstrate many of the same properties as humans in face-to-face conversation, including the ability to produce and respond to verbal and nonverbal communication." Early visions for virtual agents involved characters involved re-played recordings of human performers, such as the intelligent personal agent included in the Knowledge Navigator concept developed by Apple in 1987 [Colligan 2011 (accessed June 30, 2020] (Figure 4.1). First implementations of virtual agents were stylized nonhuman or human characters that were generated through 3D modeling and rendering and were embedded within virtual environments. An example of early nonhuman characters included Herman the Bug, an animated pedagogical agent embedded within a virtual learning environment [Lester et al. 1997]. Another early example is Rea, a real-estate agent that followed a stylized humanlike design and appeared within a simulated home environment

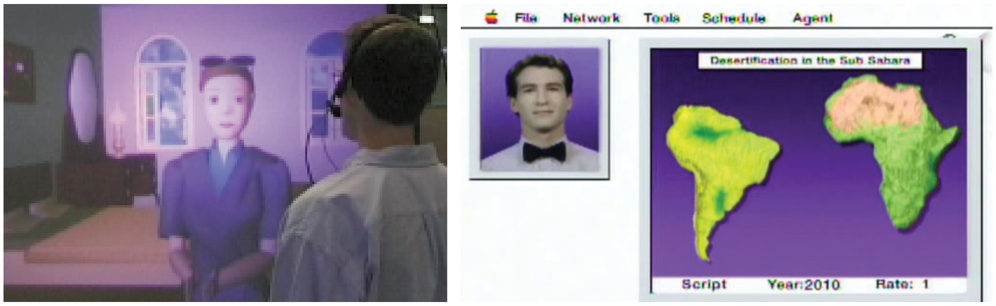

Figure 4.1 Early examples of virtual embodiments. *Left:* The Rea real-estate agent [Cassell 2000]; *Right:* the personal assistant envisioned for Knowledge Navigator [Sculley 1989].

[Cassell 2000]. Although these examples represent agents that are controlled and visualized by computer systems, the design of such nonhuman and human characters have a long history in shadow puppetry, dating back to the first millennium BC [Orr 1974]. These characters were designed for storytelling and entertainment, and the character designs reflected historical or cultural figures as well as characters developed with backstories. The design of the characters also include stylizations and ornamentations that reflect their ethnic and cultural context, such as the character Karagöz that followed a stylized human design with clothing and storyline from the 16-19th century Ottoman Empire [Scarce 1983].

The design of agents with robotic embodiments date back to mechanical humanoid automata designed as early as the 10th century BC [Hamet and Tremblay 2017]. As did with

Figure 4.2 Early physical agents. *Left:* Mechanical Turk automata by Joseph Racknitz (1789), image courtesy of Humboldt University Library; *Right:* a tea-serving Karakuri puppet, Karakuri ningyo (c) 2016 Donostia/San Sebastian.

virtual characters and shadow puppetry, the physical appearance of these early automata also followed stylized humanlike forms. Examples, shown in Figure 4.2, include the design of the Mechanical Turk, a covertly human-controlled chess-playing machine that integrated a humanoid chess player on a wooden chest where the human operator hid [Simon et al. 1999]. Karakuri puppets, mechanical automata designed in the 17-19th century Japan to be used, for example, to ceremonially serve tea, followed a stylized humanlike appearance and traditional Japanese clothing [Yokota 2009]. Although the appearance of robotic agents has overwhelmingly followed a human form with some level of stylization, robotic agents also commonly follow nonhuman morphologies. Examples of nonhuman appearances include the doglike robot Aibo designed by Sony in 1999 [Pransky 2001], a robotic seal designed for therapy in assisted living settings [Wada et al. 2005], and Keepon, a robot whose appearance resembled that of a chick [Kozima et al. 2009]. Finally, robots have also been envisioned as cartoonish characters that blend features from different sources, such as the design of the WALL·E robot by Pixar, a trash compactor with features that suggested humanlike eyes and arms [Whitley 2012].

In the 1960s, the field of computer graphics and animation started to gain momentum, and by the 1970s most of the building blocks of 3D computer animation were laid, such as surface shading by Gouraud [1971] and Phong [1975] and texture mapping by Catmull [1974]. It was not long until computer generated characters began to appear in feature-films such as *Futureworld* (1979, Richard T. Heffron), which was first to showcase a computer animated hand and face, with both wireframe and 3D shading, while the well-known film *Tron* (1982, Steven Lisberger) followed soon after with a whole 15 minutes of computer generated content. Fully animated characters also started to appear in other areas such as music videos (e.g., Mick Jagger's Hard Woman).

Ten years later, the technology was developed even further and adopted in films such as *Terminator 2: Judgment Day* (1991, James Cameron) *The Lawnmower Man* (1992, Brett Leonard), and *Jurassic Park* (1993, Steven Spielberg). This was the start of 3D animation receiving a widespread commercial success and it was not long until Pixar Animation Studios released the first entirely computer-animated feature-length film *Toy Story* (1995, John Lasseter). Toy Story was a massive success, largely due to the use of appealing cartoon-characters with plastic appearance, which computer graphics shading was perfectly suited to at that time.

In the 2000s, more technology was being developed to support the growing industry and Pixar's *Monsters Inc.* (2001, Pete Docter) showed impressive results with simulated fur depicting the subtle secondary motion on the coats of the monster characters. *The Lord of the Rings: The Fellowship of the Ring* (2001, Peter Jackson) pushed new boundaries with realistic crowd simulation, while in the same year *Final Fantasy: The Spirits Within* (2001, Hironobu Sakaguchi) attempted to create the first photo-realistic virtual humans. While the near-lifelike appearance of the characters in the film was well received, some commentators felt the character renderings appeared unintentionally creepy. Films *The Polar Express* (2004, Robert Zemeckis) and *Beowulf* (2007, Robert Zemeckis) marked further milestones in photorealism, but again received poor audience reactions. Photorealistic rendering was used more successfully for fantasy creatures such as the character Gollum from *The Lord of the Rings: The Fellowship of the Ring*, the first full CGI character in a live-action movie. The actor that drove the movements of Gollum (Andy Serkis) even went on to win the first performance-capture Oscar for his acting in later films. Similar success was achieved with the photorealistic fantasy Navi characters in *Avatar* (2009, James Cameron).

More recent advancements in 3D scanning, Deep learning, and performance capture have allowed actors to play realistic-depictions of their younger selves (*Bladerunner 2049* (2017, Denis Villeneuve), *The Irishman* (2019, Martin Scorsese), *Gemini Man* (2019, Ang Lee)) or even to play virtual roles after they have passed-away (Peter Cushing in *Star Wars: Rogue one* (2017, Gareth Edwards) and Paul Walker *Fast and Furious 7* (2015, James Wan)).

In the 1980 and 90s, there was also a shift towards interactive media such as games, where real-time animation was employed. This posed new challenges for character creation due to the additional requirements of character responsiveness and agency.

Game characters were thus less visually complex than film characters of the time due to the higher computation cost. The first attempts in the 1980s were in the form of simple 2D sprites such as Pac-Man (Namco), Sonic the Hedgehog (Sega), and Mario (Nintendo). With the advent of home console systems and consumer-level graphics processing units, there was a shift from 2D to 3D in games such as *Quake*, *The Legend of Zelda: Ocarina of Time*, *Tomb Raider*, and *Star Wars Jedi Knight: Dark Forces II*. Characters started to appear more sophisticated and used texture mapping techniques for materials and linear blend skinning for animation.

In the 2000s, many games utilized cut scenes of cinematic sequences which could achieve higher photo-realism and conversation while disabling the interactive element of the game (e.g., *LA Noire*, *Heavy Rain*, etc.). Nowadays, with real-time raytracing available in game engines, there is no longer a need for photorealism to be restricted to cut-scenes, and we are seeing incredibly realistic depictions of humans and environments in real-time (e.g., *Detroit: Become Human* and *Hellblade: Senua's Sacrifice*).

Throughout the years, the graphics and game components have developed rapidly, allowing progressively more realistic depictions every year, though characters with advanced facial animation and conversational capabilities are rarely seen. In commercial games, conversing with non-player characters (NPCs) is usually achieved by selecting predefined conversation texts on the screen, to progress the conversation. There is scope in the future for truly conversational NPCs. Additionally, as virtual reality becomes ever more immersive, we could be about to see the next evolution for the media with higher levels of realism, conversational capabilities and social presence with NPCs.

4.3 Design

4.3.1 What is appearance?

When we say “appearance” for agents, we refer to the virtual or physical embodiment that users can experience using their visual faculties. Most agents, from simple static visual representations that accompany chatbots to human surrogates, follow a metaphoric design, that is, the design of the agent takes inspiration or reference from a familiar an existing or envisioned biological entity (e.g., a human, a dog, a grasshopper) or hybrid entity (e.g., a “trash can” in appearance, but a cartoonish human in behavior). The expression of a metaphor involves two key dimensions: appearance and behavior. Metaphoric designs can follow consistent or inconsistent implementations across these two dimensions. For example, an agent that follows the metaphor of a dog and appears and behaves like a dog involves a consistent implementation, whereas a dog that speaks involves an inconsistent implementation, integrating dog-like appearance with human-like behavior. The power of agents as a family of computer interfaces comes from metaphoric design, which jumpstarts user mental models and expectations of the system using a familiar representation. For example, a computer system that uses speech as the mode of user interaction and follows a humanlike agent metaphor signals to the user that the system is capable of human mechanisms of communication, such as speech. Similarly, a robot designed to follow the metaphor of a maid or a butler is expected to be competent in household work.

A common approach to designing the appearance of agents is *metaphorical design*, where the design follows a well-known metaphor to elicit familiarity and jumpstart user mental models of the agent’s capabilities. For example, a virtual agent designed to review hospital discharge procedures with patients followed the metaphor of a nurse, appearing on the screen as a nurse in scrubs [Bickmore et al. 2009]. The design of most agents follow a singular metaphor, such as the ASIMO humanoid robot designed to appear as an astronaut wearing a spacesuit [Sakagami et al. 2002], although some designs blend multiple metaphors [Deng et al. 2019], such as the MiRo robot, which integrates multiple animal features chosen to improve perceptions of its friendliness and feelings of companionship [Prescott et al. 2017]. Metaphorical design provides not only morphological features for the design of the agent, but it also provides additional behavioral and physical features such as clothing and environmental context to further support the expression of the metaphor. An example of such features is the design of Valerie the Roboceptionist, a receptionist robot situated in a receptionist’s cubicle, equipped with a backstory that was consistent with the design of the character, and dressed in clothing that was consistent with the backstory and the metaphor that the agent’s design followed [Gockley et al. 2005]. Figure 4.3 illustrates examples of metaphorical design: the Paro, the Keepon and the iCat robots that followed the metaphors of a seal, a chick and a cat, respectively.

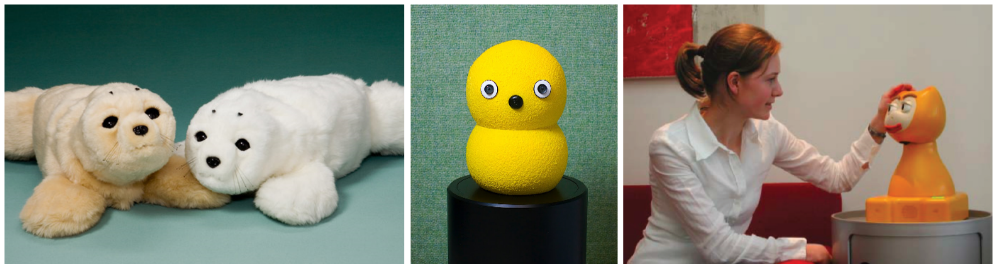

Figure 4.3 Example metaphors used in the design of robotic agents. *Left to right:* PARO Therapeutic Robot (c) 2014 PARO Robots U.S.; the Keepon robot that followed the metaphor of a chick (c) 2007 BeatBots LLC, [Kozima et al. 2009]; the iCat robot designed to follow the metaphor of a cat [van Breemen et al. 2005].

Virtual agents are also designed to follow different metaphors, most frequently of instructors or experts. For example, a digital double replica of a real doctor [Dai and MacDorman 2018] was found to be effective at delivering cues of warmth and competence (Figure 4.4). More importantly, the virtual doctor's recommendations also significantly influenced the decisions of participants in the same manner as the real doctor, implying effectiveness at persuasion.

In an educational context, a study on learning outcomes found that a human lecturer is preferable, but that robotic and virtual agents may be viable alternatives if designed properly [Li et al. 2016]. It was also shown that having a stereotypically knowledgeable appearance of the pedagogical agent influenced learning [Veletsianos 2010].

Virtual agents have also been used extensively as assistants. For example, as a navigation assistant in a crash-landing scenario in a study by Torre et al. [2018, 2019], where they had to persuade participants to accept their recommendations about items required for survival. Participants explicitly preferred interacting with a cartoon-like agent than a photorealistic one, and were more inclined to accept the cartoon-agents suggestions. Note that the photorealistic agent was rated low on attractiveness, and since persuasion and attractiveness have been linked in previous work (e.g., Suzanne R. Pallak and Koch [1983]) it may be the case that a more attractive virtual human may have been more persuasive.

Another study compared digital avatars, humans and humanoid robots to determine the influence of appearance on trust and identifying expert advice [Pan and Steed 2016]. They found that participants were less likely to choose advice from the avatar, irrespective of whether or not the avatar was an expert. In contrast, experts represented by the robot or by a real person were identified reliably.

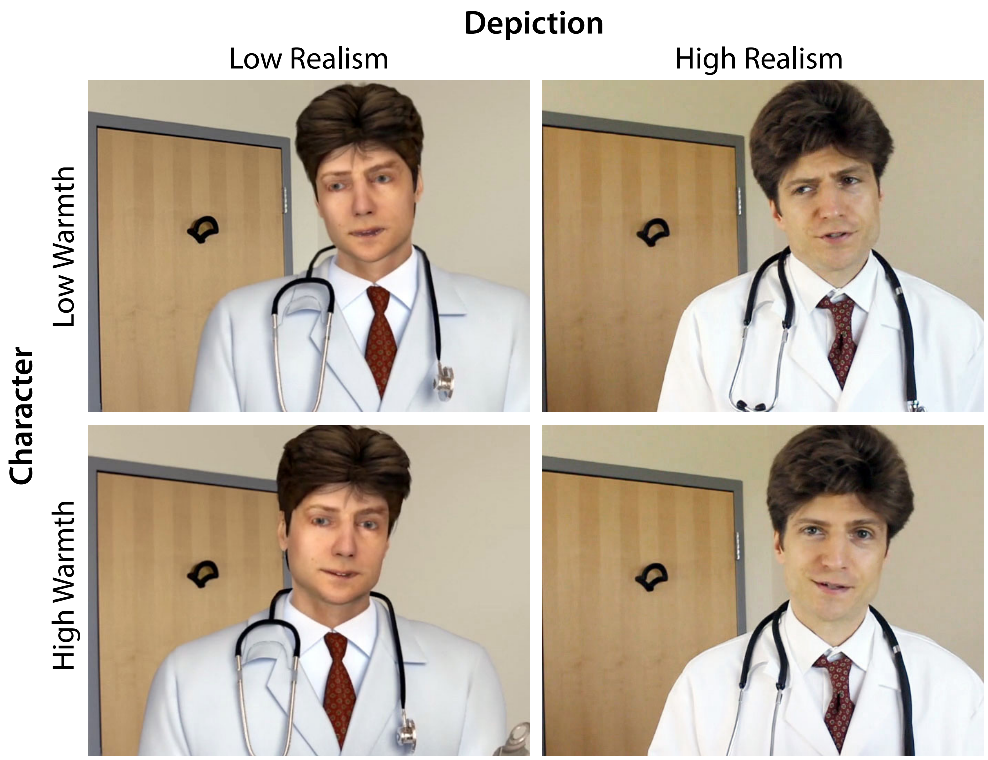

Figure 4.4 An example of agent design by replicating human experts [Dai and MacDorman 2018].

4.3.2 Modalities

Appearance can be expressed in graphical, virtual, video-mediated, physical, and hybrid modalities (Figure 4.5). Agents in graphical modalities are static or dynamic two-dimensional representations, such as a photo, drawing, or animation of a character. For example, "Laura," a virtual nurse designed to support low-literacy patients appeared as a two-dimensional rendering [Bickmore et al. 2009]. Virtual embodiments usually involve three-dimensional simulations that are rendered in real time or replays of rendered animations. An example virtual embodiment is MACH, a virtual interview coach that is rendered in real-time in a virtual environment and presented on a two-dimensional display [Hoque et al. 2013]. Such representations can also be presented in virtual-reality and mixed-reality modalities [Garau et al. 2005], which provide the user with a more immersive experience of the agent's embodiment. Agents with a physical appearance involve a robotic embodiment, such as the Robovie robot designed as a shopping mall assistant [Iwamura et al. 2011] or the Geminoid designed to serve as a human surrogate [Nishio et al. 2007]. Users of agents with physical embodiments can also experience the appearance of the agent over video [Kiesler et al. 2008].

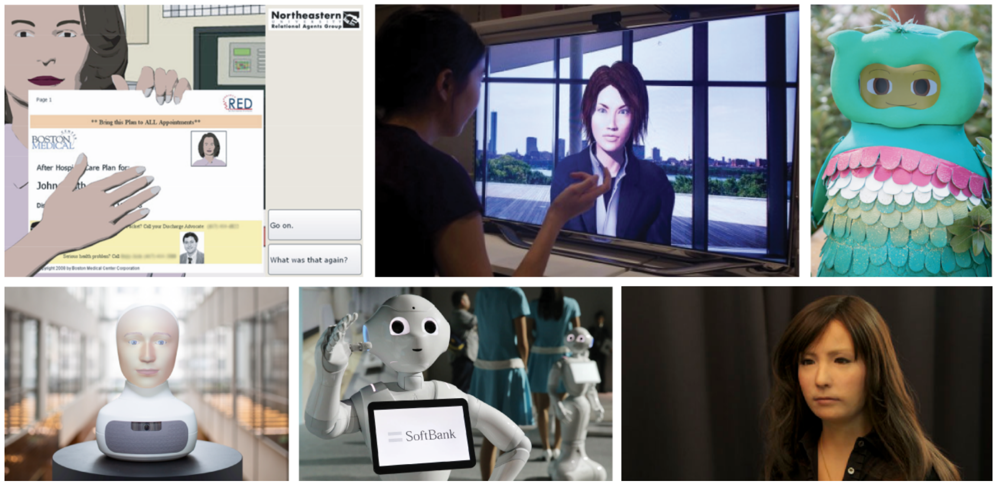

Figure 4.5 Modalities in which agents are expressed. *Left to right, top to bottom:* the nurse agent Laura rendered as a graphical agent [Bickmore et al. 2009]; the MACH virtual interview coach [Hoque et al. 2013]; the hybrid robot Spritebot with a physical body and a graphical face [Deng et al. 2019]; the hybrid FurHat robot with a physical head and a projected face (c) 2021 Furhat Robotics; the Pepper physical robot (c) 2021 SoftBank Robotics; the Geminoid F android robot [Watanabe et al. 2015].

Finally, hybrid embodiments bring physical and graphical or virtual features together, such as a graphical face appearing on a physical body or graphical features that are projected on the surface of a physical body. Example of hybrid appearances include the FurHat robot [Al Moubayed et al. 2012] or Valerie/Tank, a receptionist robot [Lee et al. 2010].

The modality in which an agent is presented affects user perceptions of and experience with the agent. A large body of literature has aimed to compare interaction outcomes across different modalities toward testing the "embodiment hypothesis:" that physical embodiment has a measurable effect on user performance and perceptions in interactions with an agent. This body of work shows that, in general, users respond more favorably to agents with stronger embodiments and human-scale sizes. In this context, "strong" embodiment refers to modalities that elicit a strong sense of presence, such as physical or hybrid modalities, and "weak" embodiment describes modalities such as graphical or virtual that may not elicit a sense of presence at such an extent. Deng et al. [2019] systematically analyzed 65 studies that compared virtual and physical agents in measures of perceptions of the agent and task performance. The analysis showed that 78.5% of these studies involved improvements in at least one of these categories of measures, consistent with the embodiment hypothesis, 15.4% involved no change, and 6.1% involved worsening in at least one of the categories of measures. Among the studies included in this analysis, the most comprehensive comparison was performed by Kiesler et al. [2008], who compared a collocated robot, a lifesize video projection of a remote robot, a lifesize projection of the virtual version of the robot, and the virtual robot on a computer screen. The measured interaction outcomes generally decreased in this order, the participants responding to the robot more favorably than the virtual agent and the collocated robot more than the projected robot.

The modality in which the agent is presented not only affects user interaction with the agent, but it also presents different sets of affordances. For example, even if the behaviors of a virtual character and a physical robot are controlled by the same algorithm, the behaviors demonstrated by the agents might look very different due to the differences inherent in the modalities. Unlike virtual characters, physical robots are subject to mechanical limitations and bound by the physical properties of the real world, which might affect the speed with which the agent displays a desired behavior (unbounded in virtual characters, bounded by actuator performance in robots), the sounds that the agent makes (e.g., sound artifacts produced by robots executing motion), the detail with which agent features can be fabricated (bound by modeling and rendering limitations in virtual characters and by physical fabrication limitations in robots), and so on. Physical robots and hybrid agents afford touch interactions and offer texture and material hardness as additional cues. The scale in which the agent is presented is another factor that affects affordances and interaction outcomes. Across all modalities, the closer the agent is presented to human scale, the more likely the agent will support human communication mechanisms. For example, a robot that is expected to be hugged by users must have a size that affords hugging.

4.3.3 Agent Construction

An important factor that shapes agent appearance is how agents are constructed, which due to historical as well as practical reasons varies based on the modality of the agent. For example, physical agents are constructed using processes and practices from *industrial design*, and their designs are affected by factors such as manufacturing limitations, product safety, and material choice. On the other hand, the construction of virtual characters borrows processes and practices from *animated filmmaking* and *game design*, and their designs are affected by factors including character backstory, the environment in which the agent will be presented, and the mechanisms with which the agent interacts with its environment, the user, and the user's environment. The paragraphs below outline some of these processes and practices.

4.3.3.1 Construction of virtual characters

Virtual characters have fewer constraints in terms of design than robots, and can be programmed to take on a multitude of different appearances, using a variety of modelling, and rendering techniques. For modelling, virtual characters are typically visualised in 3D using a mesh of consecutive planar polygons which approximate the surface of the human's body.

Polygons are very simple building blocks, and so can be used to describe many different shapes. They are also very quick to render on graphics hardware. The construction of 3D models is an established industry with many sophisticated packages available for model-building (e.g., 3D Studio Max, Maya, Blender, Houdini, etc.). Creating detailed 3D virtual characters using these packages is a highly skilled and labour-intensive task primarily due to the fact that 3D models are created using a 2D display and a high level of geometric detail is required to create convincing virtual characters. Generating 3D data for virtual characters can also be accomplished by scanning real people using a range of techniques such as photogrammetry, structured light scanning or laser scanning. Photogrammetry is a type of scanning whereby a collection of still photographs from regular DSLR cameras taken from various angles is all that is required to create a 3D model. Software then analyzes the photographs, matching characteristic points of the object on the images. This creates a point cloud of vertices which can later be converted into a mesh. It is the most commonly used tool nowadays for scanning humans in the visual effects industry, where the number and quality of cameras used in the rig contribute to the accuracy of the recovered mesh.

3D scanning can also be performed using sophisticated 3D scanning devices to project structured patterns of light or lasers onto the surface of the human to reproduce a 3D model that is a copy of the original.

For more stylized characters, artists can sculpt characters out of clay and then use one of the mentioned forms of 3D scanning to gather the data onto the computer.

Professional grade 3D scanners are expensive, but there are also more affordable, consumer-grade technologies such as depth-sensor based 3D scanning (e.g., Microsoft Kinect) and low-cost photogrammetry, which use regular cameras, but results are generally of lower quality and suitable only for low fidelity non-player characters. In the industry, there are a number of rapid character creation products that only require a single photo and create a virtual human within seconds on a tablet or phone [Didimo 2019, itSeez3D 2020, Loom.ai 2020, Pinscreen 2019]. These methods are improving in quality and speed with recent advancements in computer vision and deep learning [Hu et al. 2017, Nagano et al. 2018, Saito et al. 2017, Thies et al. 2016, Yamaguchi et al. 2018].

Once a 3D representation of a human character has been created, a number of different techniques can be utilised in order to add detail and realism. A wide variety of render styles from photorealistic to non-photorealistic can be achieved using rasterization for local illumination or ray-tracing for more realistic global illumination [Marschner and Shirley 2016]. While the rasterizer is the current standard for real-time, recent GPU optimization allows for ray-tracing in real-time, and we expect to see much higher realism in virtual characters in the future with global-illumination.

Besides the underlying rendering approach, there are many other methods for adding realism such as texture mapping, and approximating the surface reflectance through shading [Masson 2007]. Diffuse texture mapping enhances the character by adding image-based

Figure 4.6 *Left:* Wireframe render of a character with no texture mapping, *Center:* diffuse textures applied, *Right:* high quality rendering including normal maps, specular map, subsurface scattering, global illumination, etc.

information to its geometry, while entailing only a small increase in computation. The basic idea is to map the colour of the image or 'texture' onto the corresponding colour of an object at each pixel [Catmull 1974] which adds the illusion of detail to the model, such as clothing material and skin colour.

In order to add colour detail to virtual characters, diffuse texture maps are used which define the color of diffused light (Figure 4.6). Additionally, there are situations where surfaces are not smooth and roughness needs to be added if it is not present in the geometry. For example, skin is not a smooth surface as it has imperfections such as pores and wrinkles. These details are best added using normal maps which perturb the surface normals to add detail or displacement maps which add geometric detail.

In modern computer graphics, surface properties are governed by shaders, the code snippets describing how a surface should react to incident light. Many physically-based shaders have been developed to produce realistic materials with different Bidirectional Reflectance Distribution Functions (BRDF) [Nicodemus et al. 1992] (the function that relates the incident to the reflected light). More recently, with the rapid advancements in graphics hardware, more complex shading effects approximating a wide range of BRDFs can now be achieved in real-time. For example, subsurface light transport in translucent materials [Jensen et al. 2001] for realistic scattering of light on the skin was once a technique only used in off-line high-end visual effects, but real-time methods [Jimenez et al. 2009, 2010] are now used to enhance the realism in real-time.

Hair for interactive virtual characters has traditionally been modelled using card-based rendering, where images of chunks of the hair are mapped onto large flat sheets, to approximate the shape of a much larger number of individual hairs. Later advancements allowed for modelling each individual hair which dramatically improves realism. For rendering of hair, physically-based fiber reflectance models are used, based on a combination of an anisotropic specular and a diffuse component [Kajiya and Kay 1989]. More recently, the scattering distribution of the hair fiber is split into different lobes based on the number of internal reflections within the fiber [Marschner et al. 2003].

The use of physically-based simulations is ubiquitous for realism in virtual clothing, where fast mass-spring models [Liu et al. 2013] or more complex implicitly integrated continuum techniques [Baraff and Witkin 1998] are used in the state-of-the-art. Implementing realistic cloth and hair dynamics in real-time applications still represents a significant challenge for developers since simulation dynamics need to be solved at run-time, and are required to be fast and stable. Based on this, depictions of stiff clothing and hair with little secondary-motion effects are still commonplace for interactive virtual characters across a range of applications from video games to virtual assistants.

4.3.3.2 Industrial design of robots

The paragraphs above have discussed design approaches, e.g., metaphorical design, to and the resources used, e.g., facial features, for the development of agent appearance. Another factor that significantly affects agent appearance is the industrial design of physical agents or the physical platforms in which virtual or hybrid agents are presented, including form, material use, scale, color choice, and so on. Although there are no systematic studies of how these factors affect agent appearance or how they must be designed to maximize user experience, the HRI literature includes reports of the design process for the appearance of specific robot platforms. For example, Lee et al. [2009] described the design process for Snackbot, a robot designed to deliver snacks in an academic building, including the form of the housing of the robotic hardware and the snack tray that the robot would carry; the material and colors used to construct the housing and the tray; the height of the robot; and the expressive features of the head and face of the robot. Another example is the design of the Simon humanoid robot, where the research team explored the proportions that the robot's head and body should follow, the placement of the eyes on the head, facial features that would achieve the appearance of a "friendly doll," and the interplay between the design of the housing and structural or mechanical elements of the robot's head [Diana and Thomaz 2011]. Hegel et al. [2010] documented and reported on the industrial design of the social robot Flobi, which included an exploration of the design of the robot's head to follow a "baby face" schema; effective color combinations of the robot's face, hair, lips, and eyebrows; and how blushing on the robot's cheeks could be achieved using LEDs placed behind the surface of the face. A final example is the design of Kip1, a peripheral robotic conversation companion, involving form and material exploration through sketches and mock-ups [Hoffman et al. 2015]. Figure 4.7 illustrates the sketches and mock-ups generated in the industrial design of some of these examples.

In all of the examples discussed above, the research team engaged professional industrial designers or members of the research team with training in industrial design as well as an

Figure 4.7 Sketches and models generated during the industrial design of the Snackbot [Lee et al. 2009] (top-left), Simon [Diana and Thomaz 2011] (top-right), and Kip1 [Hoffman et al. 2015] (bottom) robots.

iterative design process. The literature does not include any discussion of such considerations for virtual characters, and characters designed for research and commercial use all utilize existing display platforms, such as mobile phones, tablet computers, computer monitors, large displays, or virtual- or mixed-reality environments. Overall, there is a great need for systematic research on the industrial design of the appearance of agents, including the effects of the physical design of the agent itself and the environment within which virtual agents are presented on user interaction and experience.

4.4 Features

The design approaches described above draw on a rich space of features, shaped by the metaphor followed by the design (e.g., humanlike features included in the design of a virtual human), functional requirements of the agent (e.g., light displays placed on physical robots to convey the agent's status), and/or aesthetic and experiential goals of the design (e.g., material, color, and texture choices for a robot). The paragraphs below provide an overview of this space, focusing on facial and bodily features as well as features that communicate demographic characteristics of virtual and physical agent embodiments.

4.4.1 Facial features

The face of an agent serves as the primary interface between the agent and its user, and facial features make up a substantial portion of the design space for agents. Even when designs lack anthropomorphic or zoomorphic faces, people attribute facial features to them, highlighting the importance of faces in the perception of non-living objects [Kühn et al. 2014]. Designers of virtual and physical agents draw on this human propensity and create faces that can display conversational cues, express affect, and communicate direction of attention.

In order to convey a true feeling of life in a character, the appearance of the eye is highly important. Rendering techniques such as adding specular and reflection maps can be very useful for this purpose to increase the appearance of wetness and to reflect the environment. Additionally, more advanced techniques such as ambient occlusion allow for soft shadowing, and refraction to replicate the refraction of light that passes through the eyeball, which is filled with fluid. Creating the geometry of the eye is a difficult task, due to the complexity of the surface but there exist special photogrammetry rigs for capturing the visible parts of the eye—the white sclera, the transparent cornea, and the non-rigidly deforming colored iris [Bérard et al. 2014]. Computer generated eyes used in computer graphics applications are typically gross approximations of the actual geometry and material of a real eye. This is also true for facial expressions, which typically take a simple approach of linearly blending pre-generated expression meshes (blendshapes) to create new expressions and motion [Anjyo 2018]. However, little is known about how these approximations affect user perception of the appearance of virtual characters.

Similar to the studies on real humans, virtual humans with narrow eyes have been rated as more aggressive and less trustworthy for both abstract creatures [Ferstl et al. 2017] and more realistic depictions [Ferstl and McDonnell 2018] (Figure 4.8). It should be noted that for realistic eye size alterations, the size of the eyes themselves should not be scaled as this will be quickly perceived as eerie and artificial [Wang et al. 2013]. Instead, the shape of the eyelids can be changed as protruding eyes appear larger, whereas hooded and monolid eyes appear smaller.

In contrast to human face studies, wider faces were not judged as less trustworthy, and were perceived as less aggressive compared to narrow faces for realistic [Wang et al. 2013] and abstract virtual characters [Ferstl et al. 2017], even when a particularly masculine rather than a babyface appearance was presented [Ferstl and McDonnell 2018]. The results of these studies support the notion that virtual faces are perceived differently from real human faces. A potential explanation could be the tendency of villains in animated movies to be portrayed with narrow, long, sharp facial features (e.g., Captain Hook in *Peter Pan* (Clyde Geronimi, 1953), Scar in *The Lion King* (Roger Allers, 1994), Maleficent in *Sleeping Beauty* (Clyde Geronimi, 1959)). This tendency could influence the perception of computer-generated characters towards automatic association of narrow faces with dangerous characters.

Other work has addressed the perception of rather unusual facial proportions for realistic characters and their influence on perceived appeal. Seyama and Nagayama [2007] studied eye size by morphing between photographs of real people and dolls, and found that characters were judged as unpleasant if the eyes had strong deviations from their original size. Participants were more sensitive to the alterations for real faces than for artificial faces. Several studies confirmed that altering facial parts lowers perceived appeal, especially for humanlike characters. Green et al. [2008] demonstrated that not only proportions, but also the placement of facial parts may negatively affect the perceived appeal. The measured effect was greater for the humanlike and more attractive faces. Additionally, it has been demonstrated that a mismatch of realism between facial parts negatively affects appeal [Burleigh et al. 2013, MacDorman et al. 2009].

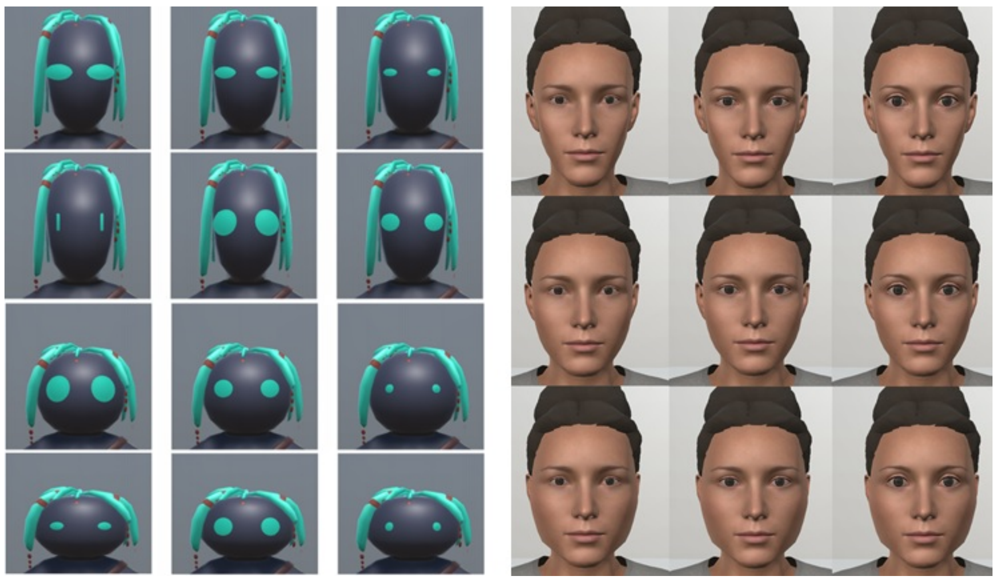

Figure 4.8 *Left:* Examples of eye and head shape manipulations on abstract characters (based on [Ferstl et al. 2017]), *Right:* More subtle facial feature manipulations on realistic virtual characters (adapted from [Ferstl and McDonnell 2018]).

Prior work in HRI includes a large body of literature on the facial features of robotic agents. A number of studies aimed to characterize the design space for robot faces. Blow et al. [2006a] characterized this space as varying across the dimensions of abstraction, from low to high abstraction, and realism, from realistic to iconic, borrowing from literature on the design of cartoon faces [McCloud 1993]. DiSalvo et al. [2002] carried out an analysis of 48 robots and conducted an exploratory survey that resulted in a number of design recommendations to improve human perceptions of humanlike robots: (1) the head and the eye space should be wide; (2) facial features should dominate the face with minimal space for a forehead and a chin, (3) the design should include eyes with sufficient complexity; (4) the addition of a nose, a mouth, and eyelids improve perceptions of humanlikeness; and (5) the head should include a skin or a casing that core the electromechanical components. A similar analysis was carried out by Kalegina et al. [2018] of 157 rendered robot faces—physical robots that

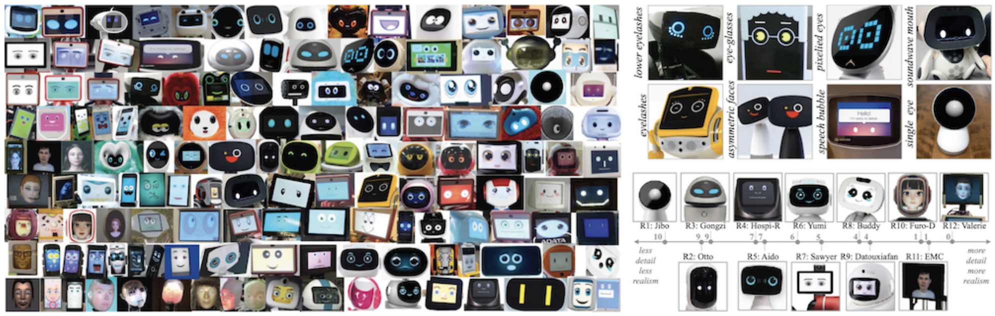

Figure 4.9 The 157 faces analyzed by Kalegina et al. [2018] (left), their analysis of facial features used in the design of the robot faces (right-top), and the spectrum of facial realism (right-bottom). *Copyright Information:* Images included in this paper under ACM guidelines on Fair Use

are equipped with a screen-based face and facial features that are virtually rendered on the screen—who coded the faces for 76 different features and conducted a survey to understand how each feature affected user perceptions of the robot (Figure 4.14). The study found that faces with no pupils and no mouth were consistently ranked as being unfriendly, machinelike, and unlikable; those with pink or cartoon-style cheeks were perceived as being feminine; and faces with detailed blue eyes were found to be friendly and trustworthy. Survey participants also expressed preferences for robots with specific facial features for specific contexts of use, e.g., selecting robots with no pupils and no mouth for security work and faces with detailed blue eyes for entertainment applications. Consistently, Goetz et al. [2003] argued that there is not a universally preferred design for the facial features of a robot, but that people prefer appearances that match the robot's task. They varied the robot's appearance across three stylistic dimensions—human vs. machine, youth vs. adult, and male vs. female—and found that user preferences for facial features presented in these styles depended on the robot's task. In a follow-up study, Powers and Kiesler [2006] showed that the length of the robot's chin and the fundamental frequency of its voice predicted whether participants expressed interest in following advice from the robot.

The literature also includes reports of the process for the design and development of faces for several robot platforms. For example, the design of the iCub social robot primarily involved the mechanical replication of human anatomical mechanisms to achieve realistic eye and head movements and the design of the rest of the face to follow a "toy-like" appearance [Beira et al. 2006]. The design specifications for the face of the KASPAR social robot included a sufficiently expressive but minimal design, an iconic overall design (as opposed to a realistic one), a humanlike appearance, and the ability to express autonomy, communicate attention, and display projected expressions [Blow et al. 2006b, Dautenhahn et al. 2009]. The design of the humanoid robot HUBO integrated an abstract body with the overall appearance of an astronaut and an highly humanlike face using elastomer-based materials that appeared and moved similar to human skin and a 28-degree-of-freedom mechanism to achieve humanlike facial movements [Oh et al. 2006]. The faces of robots including the Flobi [Lütkebohle et al. 2010], Melvin [Shayganfar et al. 2012], and iCat [van Breemen 2004] featured pairs of flexible actuators that served as the robot's lips and pairs of eyebrows to express emotion. As discussed earlier, the design of the face of the Flobi robot, shown in Figure 4.14, additionally included sophisticated mechanisms for emotion expression, such as lights placed behind the cheeks to enable the appearance of blushing. These reports illustrate how different facial features come together in the design of different robot systems and point to specific examples in the design space of facial features for robots.

4.4.2 Bodily features

While the face serves as the primary interface for human-agent interaction, the remainder of an agent's body also contributes to the appearance of the agent. The design space for an agent's body primarily includes several bodily features, how these features come together structurally, and how they are represented.

A virtual agent's body can be presented in a range of different styles, from low-detailed stick-figures or point-light displays to photorealistic bodies or anthropomorphised creatures, and there have been some studies aimed at investigating the effect of the body representation on perception of the agent's appearance and actions. Most studies apply motion captured animations to a virtual character and map the motion onto a range of bodies and assess if the different bodies change the meaning of the motion. Typically, factors such as emotion, gender, and biological motion are chosen since these have all been shown to be identifiable solely through motion cues (e.g.,[Cutting and Kozlowski 1977, Johansson 1973, Kozlowski and Cutting 1977]) thus allowing the contribution of the bodies appearance to be assessed.

Beginning with a study by Hodgins et al. [1998], the amount of detail in a virtual character's representation has been studied to investigate the effect on perception. Their study found that viewers' perception of motion characteristics is affected by the geometric model used for rendering. They observed higher sensitivity to changes in motion when applied to a polygonal model, than a stick figure. Chaminade et al. [2007] also found an effect on motion perception, where character anthropomorphism decreased the tendency to report their motion as biological, while another study found that emotions were perceived as less intense on characters with lower geometric detail [McDonnell et al. 2009b].

Body shape has also been investigated where it was found that a virtual character's body does not affect recognition of body emotions, even for extreme characters, such as a zombie with decomposing flesh [McDonnell et al. 2009b] (Figure 4.10). Fleming et al. [2016] evaluated the appeal and realism of female body shapes, which were created as morphs between a realistic character and stylized versions following design principles of major computer animation studios. Surprisingly, the most appealing characters were in-between

Figure 4.10 Different structural and material representations for agent body [McDonnell et al. 2009c].

morphs, where 33% morphs had the highest scores for realism and appeal and 66% morphs were rated as equally appealing, but less realistic (Figure 4.11).

The perception of sex of a virtual character's walking motion has also been found to be affected by body shape. Adding stereotypical indicators of sex to the body shapes of male and female characters influences sex perception. Exaggerated female body shapes influenced sex judgements more than exaggerated male shapes [McDonnell et al. 2009a].

In virtual reality, embodiment of virtual characters is where the user is positioned virtually inside the body of a virtual avatar, where they have agency over that virtual body. The character model used for the virtual avatar can affect the behaviour of the user, from becoming more confident when embodied in a taller avatar, more friendly as an attractive avatar [Yee and Bailenson 2009], to reducing implicit racial bias by embodying an avatar of a different race [Banakou et al. 2016]. This powerful effect is referred to as the Proteus effect [Yee and Bailenson 2007] (named after the Greek god known for his ability to take on many different physical forms). The use of self-avatars or virtual doppelgangers has also been shown to affect outcomes, with generally a positive influence on aspects such as cognitive load [Steed et al. 2016], pain modulation [Romano et al. 2014] and embodiment [Fribourg et al. 2020, Kilteni et al. 2012]. These effects describe to some extent the dynamism of interactions between users and avatars.

The design of a physical robot's body is shaped by a number of factors, including the metaphor that the design follows, the functional requirements of the robot, and environmental constraints that the design must consider. The first factor, the design metaphor, might dictate how the body of the robot is structured and the features that are articulated in the design.

For example, the Paro robot [Wada and Shibata 2007] follows the metaphor of a baby seal, and the design of the robot's body roughly follows the form of a seal, including fore and hind flippers. The functional requirement of the robot might include specific forms of mobility, such as holonomic movement, climbing stairs, or movement across rough terrain, or prehensile manipulation involving a single arm or two arms. Depending on such design requirements, the design of the body of a robot might follow a humanoid design including humanlike limbs attached to a torso, such as the ASIMO robot [Sakagami et al. 2002], or a single arm attached on a mobile base, such as the Fetch robot [Wise et al. 2016]. Finally, the environment that the robot is designed for can dictate the bodily features of the robot, such as requiring that a robot that crawls into tight spaces has a low profile and limbs that can be tucked away, such as a Packbot robot [Yamauchi 2004] used in search-and-rescue scenarios.

In addition to bodily features borrowed from the design metaphor, such as the hind flippers of a seal or the legs of a human, the design of physical robots also utilize features that facilitate specific functions. These functions include communication, and features that support communication include lights that communicate the robot's affective states using different colors [Bethel and Murphy 2007] or light arrays that convey information about the robot's direction of motion using light patterns [Szafir et al. 2015]. Features of a robot's body may also support transferring items, such as a tray that the Snackbot robot held to carry food items [Lee et al. 2009] and the different configurations of carts that hospital delivery robots pull to transport materials [Ozkil et al. 2009].

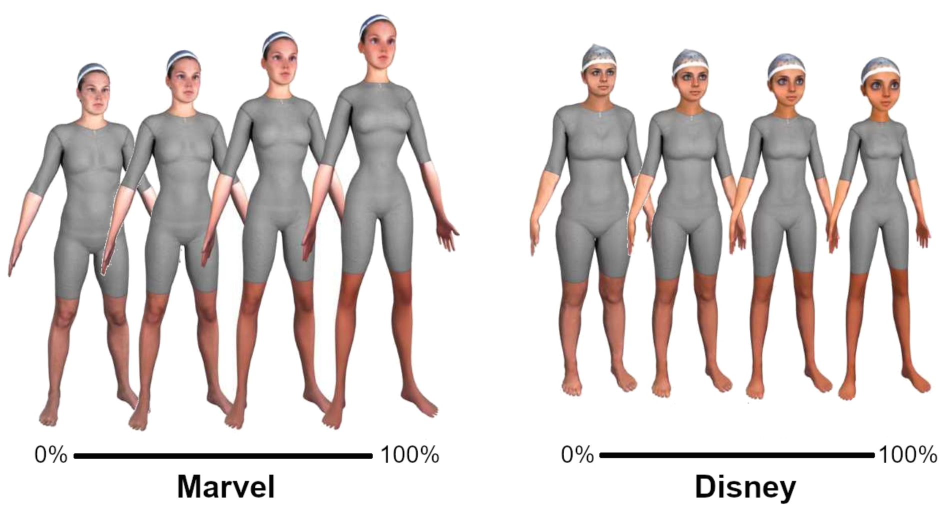

Figure 4.11 Stylization applied at different levels (33%, 66%, 100%) to captured performer body (0%) in Marvel and Disney styles (image based on Fleming et al. [2016]).

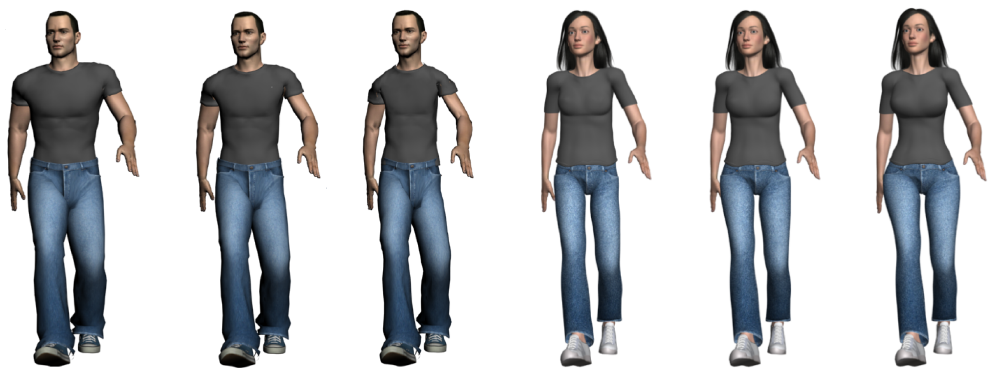

Figure 4.12 Six body shapes with indicators of gender [McDonnell et al. 2009a].

An agent's body can also include bodily features, such as clothing or furniture, designed to support the agent's character or backstory or eventually improve user experience with the agent. For example, the Roboreceptionist robot was placed in a booth that resembled an information booth and wore clothes that were consistent with the gender and the backstory of its character [Gockley et al. 2005]. The Geminoid robot, a highly realistic android developed to serve as a robotic surrogate to support remote communication, was constructed to resemble its creator and dressed in similar fashion [Nishio et al. 2007]. Figure 4.13 illustrates examples of bodily features that support specific functions, such as a tray, and that support the agent's character, such as clothing.

Figure 4.13 Bodily features that support specific functions, such as a tray that the robot uses to delivery snacks [Lee et al. 2009] (left) and light arrays that a flying robot uses to communicate direction [Szafir et al. 2015] (left-center), and that support the agent's character, such as a booth and clothing for a receptionist robot [Lee et al. 2010] (right-center) and clothing for a surrogate robot [Watanabe et al. 2015] (right).

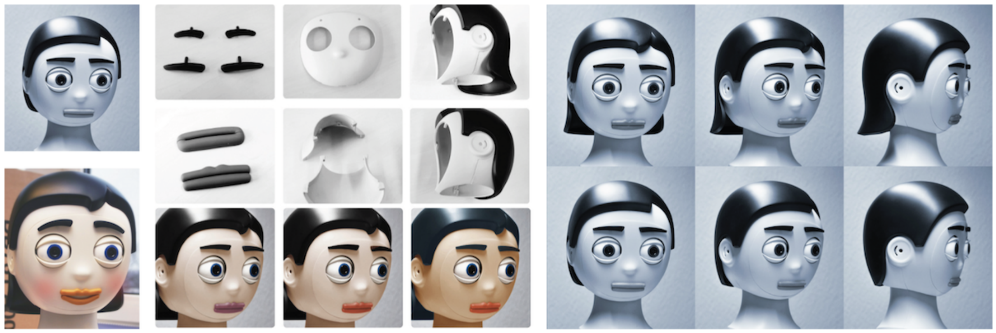

Figure 4.14 Facial features of the Flobi robot that provide the robot with different demographic characteristics. *Left:* neutral male (top) and smiling female (bottom) faces; *Center:* the physical parts that represent facial features; *Right:* different hair and lip styles. Adapted from Lütkebohle et al. [2010].

4.4.3 Features expressing demographic characteristics

Agent appearance communicates other attributes of the character of the agent, such as gender, age, race, and ethnicity. Virtual agents are usually designed as distinctive characters, such as the two female nurse characters, one middle-aged Caucasian and one middle-aged African American, designed by Bickmore et al. [2009] to match user patient demographics. Physical agents, on the other hand, are designed as characters with ambiguous features and interchangeable parts that highlight specific character attributes, such as the interchangeable hair and lips of the Flobi robot that communicate a male or female gender [Lütkebohle et al. 2010] (Figure 4.14).

A large body of research on human-agent interaction has shown such character attributes to significantly shape interaction outcomes. For example, Siegel et al. [2009] asked participants to make an optional donation to a robot that used pre-recorded male or female voices, which research has shown to be sufficient to trigger gender stereotypes [Nass et al. 1997], and found a significant interaction between robot and participant gender over the proportion of participants who donated any amount, e.g., men consistently donating more to a female robot. Eyssel and Hegel [2012] manipulated the gender of the Flobi robot by varying the robot's appearance via its interchangeable parts for hair and lips and found that participant perceptions of the male and female robots closely followed gender stereotypes. The male robot was perceived as having more agency and being more suitable for stereotypically male tasks (e.g., repair), and the female robot was perceived as being more communal and being more suitable for stereotypically female tasks (e.g., childcare).

The effect of stereotypes has also been studied for virtual characters, mostly in the context of embodiment in virtual reality. The Proteus Effect, as mentioned previously, has additionally shown that users conform to stereotypes associated with their avatar's appearance. For example, embodiment in female avatars made players more likely to conform to female-typed language norms [Palomares and Lee 2010] and made them more likely to engage in healing activities [Yee et al. 2011]. Interestingly, these effects were observed regardless of the actual gender of the player, indicating a tendency to conform to expectations associated with the virtual gender.

In other work, Zibrek et al. [2015] explored gender bias on different types of emotions applied on male and female virtual characters. They found that emotion biases gender perception according to gender stereotypes: an angry motion is seen as more male, while fear and sadness are seen as less male motions, and they observed a contrast effect where anger was seen as more male when viewed on a female model than when viewed on a male model. Similar effects were found for real humans [Hess et al. 2004], indicating that virtual humans follow similar stereotyping effects.

4.4.4 Realism, Appeal, Uncanny Valley

Metaphorical design involves the application of a familiar metaphor to the design of an agent, such a virtual human following the metaphor of a human. In practice, metaphors are applied at different levels of abstraction due technical limitations (e.g., inability to closely replicate the original metaphor) and design choices (e.g., stylization). Deng et al. [2019] argued that designs follow discrete metaphors (e.g., a "baby seal" metaphor) but the realism in which these metaphors are applied to vary along a spectrum of abstraction (e.g., a stylized or abstract household robot vs. a highly realistic robotic surrogate). The design choices of metaphor and abstraction result in differences in user perceptions of the agent and experience with it.

In the classic textbook "Disney Animation: The Illusion of Life," Thomas and Johnston [1995] use the term *appeal* to describe well designed, interesting and engaging characters. This is contrary to many face perception studies, which use the term appeal and attractiveness interchangeably. Appeal is an essential ingredient for virtual characters in video games and movies, as well as for avatars, agents, and robots, to ensure audience engagement and positive interactions. Creating highly detailed, photorealistic virtual characters does not necessarily produce appealing results [Geller 2008], and it is often the case that more stylized approximations evoke more positive audience responses and engagement [Zell et al. 2019]. However, additional factors are the context of the interaction and how appropriate the appearance is under the circumstances. For example, having a fun cartoon-appearance may be less appropriate for a more serious application such as a for a business meeting [Junuzovic et al. 2012] or medical training [Volante et al. 2016], etc. Perception of appeal of virtual characters is an ongoing area of research, with the ultimate goal to speed-up or automate the process of producing appealing characters, and avoid negative reactions from audiences.

The term Uncanny Valley (UV) is often used to describe the negative reactions that can occur towards virtual characters. It is a feeling of repulsion produced by artificial agents that

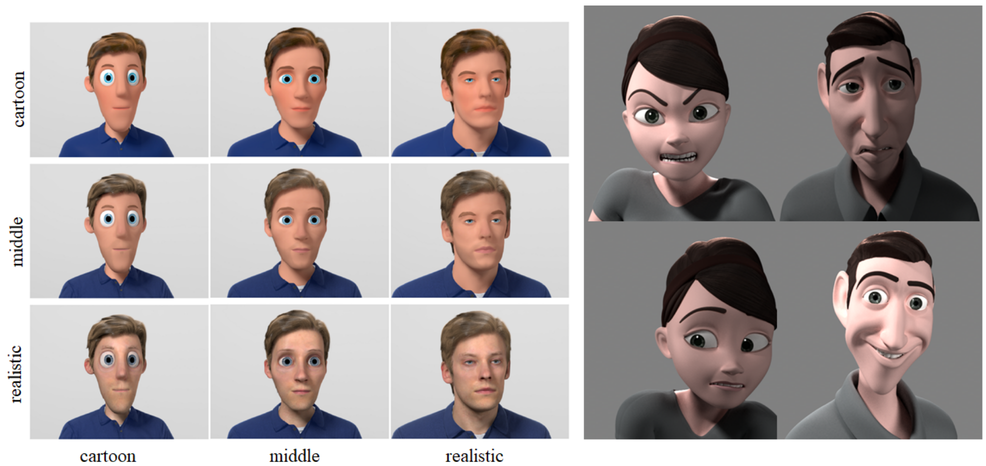

Figure 4.15 *Left:* Examples of manipulating material (y-axis) and shape (x-axis) to vary character realism and appeal, image based on Zell et al. [2015], *Right:* Examples of brightness and shadow alterations on cartoon characters displaying emotion which were shown to change the perceived intensity of emotion [Wisessing et al. 2020].

appear close to human-form but not quite real. This UV phenomenon was first hypothesized by in the 1970s by robotics professor Mori [1970]. Mori's original hypothesis states that as a robot's appearance becomes more human, humans evoke more positive and empathetic responses, until a point where the response quickly becomes strongly negative resulting in feelings of disgust, eeriness and even fear. Once the robot's appearance becomes less distinguishable from a human being, the emotional response becomes positive once again. This negative response has been attributed to many causes such as motion errors or lack of familiarity or a mismatch in realism between elements of character design. More recently, the UV hypothesis has been transferred to virtual humans in computer graphics, and has been explored directly in some studies [Bartneck et al. 2009, MacDorman et al. 2009]. Virtual faces in particular are difficult to reproduce as humans are very adept at perceiving, and recognising other faces and facial emotions.

As discussed previously, the appearance of a character can be separated into texture, materials, shape and lighting. Various studies have attempted to isolate these factors and independently examine the effect on appeal and UV.

Wallraven et al. [2007] studied the perceived realism, recognition, sincerity, and aesthetics of real and computer-generated facial expressions using 2D filters to provide brush, cartoon, and illustration styles and found that stylization caused differences in recognition accuracy and perceived sincerity of expressions. Additionally, their realistic computer-generated faces scored high aesthetic rankings, which is contrary to the UV theory. Pejsa et al. [2013] additionally found no effect on appeal or lifelikeness between a character with human proportions and one with stylized geometry including large eyes, while other studies found realistic and cartoon depictions to be equally appealing when expressing personality [Ruhland et al. 2015] and when a user had agency over their movements [Kokkinara and McDonnell 2015].

In order to investigate the effect of stylization in more detail, McDonnell et al. [2012] created a range of appearances from abstract to realistic by altering the rendering style (texture, material and lighting) of a realistically modelled male character while keeping the shape and motion constant. They analyzed subjective ratings of appeal and trustworthiness and found that the most realistic character was often rated as equally appealing or pleasant as the cartoon characters, and equally trustworthy in a truth-telling task. A drop in appeal occurred for characters in the middle of the scale (rated neither abstract nor realistic), which was attributed to the difficulty in categorizing these characters due to their uncommon appearance [Saygin et al. 2012]. Other studies of the UV that used still images generated by morphing between photographs and animated characters also found valleys in participant ratings of uncanniness for intermediate morphs [Green et al. 2008, Hanson 2005, Seyama and Nagayama 2007]. This idea was further developed in the categorization ambiguity hypothesis [Cheetham and Jancke 2013, Yamada et al. 2013], where it was shown that this response is more prominent when the morph is between a real human and an inanimate object or representation of a human. Studies focusing on neurocognitive mechanisms attribute negative evaluation to a competing visual-category representations during recognition [Ferrey et al. 2015].

This effect was also investigated in a study by Carter et al. [2013] where they created a realistic, cartoon, and robot female character and assessed subjective pleasantness ratings as well as analyzing eye-tracking as a psychophysiological measure. Contrary to the UV theory, they found higher ratings of unpleasantness for their cartoon than for their realistic character, and that fixations were affected by subjective perceptions of pleasantness.

Investigating yet more parameters, Zell et al. [2015] independently examined the dimensions of shape, texture, material and lighting, by creating a range of stimuli of characters with various levels of realism and stylization (Figure 4.15 (left)). Their study identified that the shape of the character's face is the main descriptor for realism, and material increases realism only for realistic shapes. Also, that strong mismatches in stylization between material and shape made characters unappealing and eerie, in particular abstract shapes with realistic materials were perceived as highly eerie, validating the design choices of some horror movies with living puppets. Finally, blurring or stylizing a realistic texture can achieve a make-up effect, increasing character appeal and attractiveness, without reducing realism. The opposite was found in a study on body stylization, where the stylization of body shape predicted appeal ratings rather than improvements to render quality [Fleming et al. 2016].

More recently, Wisessing et al. [2020] carried out an in-depth analysis of the effect of lighting on appeal, particularly brightness and shadows, and found that increasing the brightness of the key-light or lessening the key-to-fill ratio (lighter shadows) increased the

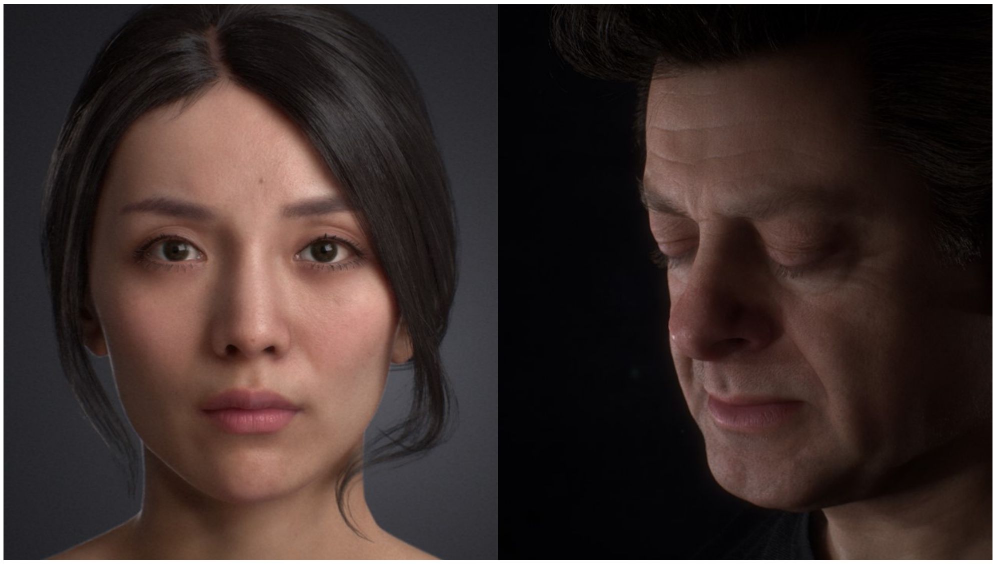

Figure 4.16 State-of-the-art real-time virtual humans in Unreal Engine 4 created by 3Lateral in collaboration with Cubic Motion, Epic Games, Tencent and Vicon. *Left:* Siren demo. *Right:* virtual replica of the actor Andy Serkis. With permission of Epic Games © 2020.

appeal ratings (Figure 4.15 (right)). They also found little effect of key-light brightness on eeriness but reported reduced eeriness as a consequence of lightening the shadows, which could be used to reduce UV effects of virtual characters. However, shadow lightening did not improve appeal for characters with realistic appearance, and thus key-light brightness alone should be used to enhance appeal for such characters.

Several studies in immersive VR have also examined the effect of character appearance on viewer responses, focusing on co-presence, i.e., the sense that one is present and engaged in an interpersonal space with the character [Biocca 1997, Garau et al. 2003]. While some evidence confirms the importance of realistic appearance [Nowak 2001, Zibrek and McDonnell 2019], others put less importance on it [Garau et al. 2003, Slater and Steed 2002]. On the other hand, a mismatch between the realism of behaviour and appearance has been often shown to lower the feeling of co-presence [Bailenson et al. 2005]. There are a number of reasons why mismatches may cause negative effects on the viewer. A mismatch between the physical and emotional states of a character violate expectations and thus can result in a breakdown in how users experience agents [Vinayagamoorthy et al. 2006].

4.5 Summary

Technical advancements are increasingly pushing the boundaries of how agents are designed and developed, the capabilities of these agents, and their use in human environments. The rapid development in real-time rendering technologies has enabled incredibly detailed, high-quality virtual character appearances (Figure 4.16), often reaching photorealism [Epic Games, Inc. 2018, Seymour et al. 2017]. Deep learning is also improving the ease and speed at which characters can be created, even from a single photograph [Yamaguchi et al. 2018]. Additionally, animation and behaviours are starting to become easier and less expensive to create, allowing virtual human technologies to be more accessible to a wider audience than ever before. With these advancements comes the increasing use of characters across different domains such as education, sales, therapy, entertainment, social media, and virtual and augmented reality.

New methods are also emerging for the construction of physical robots. Rapid fabrication methods, such as 3D printing, have led to the development of new robot morphologies, including 3D printable robots inspired by "origami" [Onal et al. 2014] and robots with soft skin that can change appearance and texture to communicate internal states to the user [Hu et al. 2018]. Mixed-reality technologies are also being utilized to facilitate human interaction with robots, including displaying cues that communicate the motion intent [Walker et al. 2018] and and the field of view [Hedayati et al. 2018] of the robot. Finally, robots are increasingly being integrated into human environments across different domains, including manufacturing [Sauppé and Mutlu 2015], education [Belpaeme et al. 2018, Michaelis and Mutlu 2018], food services [Jennings and Figliozzi 2019], hospitality [Tussyadiah and Park 2018], surveillance [Inbar and Meyer 2019], and healthcare [Miseikis et al. 2020, Mutlu and Forlizzi 2008]. As applications proliferate, we will gain a better understanding of how the design space for agent appearance is utilized to support each application domain, how the features of this space affect user perceptions of and experience with these agents, and how the appearance of robotic agents might be designed to support personalization, customization, and environmental fit.

In this chapter, we have shown that the choice of appearance can have implications for human interactions in a number of ways, including changes to the perception of personality, emotion, trust, and confidence. Studies have shown that the many factors that constitute the final appearance of an agent, such as the design metaphor, modality of representation, and methods of agent construction, including modelling, texturing, materials, and even lighting, have different effects on how people perceive and respond to it. This multidimensionality has the drawback that some factors might cancel each other out or amplify each other, leading to inconsistent conclusions. Additionally, more frequent exposure to agents and increasing technological sophistication may continuously change the way we perceive them, much like how we are becoming more and more sensitive to poor visual effects in movies [Tinwell et al.

2011]. The need for understanding the implications of different appearances of agents has therefore never been greater.

Bibliography

S. Al Moubayed, J. Beskow, G. Skantze, and B. Granström. 2012. Furhat: a back-projected human-like robot head for multiparty human-machine interaction. In *Cognitive behavioural systems*, pp. 114–130. Springer.

K. Anjyo. 2018. *Blendshape Facial Animation*, pp. 2145–2155. Springer International Publishing, Cham. ISBN 978-3-319-14418-4. https://doi.org/10.1007/978-3-319-14418-4_2. DOI: 10.1007/978-3-319-14418-4_2.

J. N. Bailenson, K. R. Swinth, C. L. Hoyt, S. Persky, A. Dimov, and J. Blascovich. 2005. The independent and interactive effects of embodied-agent appearance and behavior on self-report, cognitive, and behavioral markers of copresence in immersive virtual environments. *Presence*, 14(4): 379–393.

W. A. Bainbridge, J. W. Hart, E. S. Kim, and B. Scassellati. 2011. The benefits of interactions with physically present robots over video-displayed agents. *International Journal of Social Robotics*, 3(1): 41–52.

D. Banakou, P. D. Hanumanthu, and M. Slater. 2016. Virtual embodiment of white people in a black virtual body leads to a sustained reduction in their implicit racial bias. *Frontiers in Human Neuroscience*, 10: 601. ISSN 1662-5161. https://www.frontiersin.org/article/10.3389/fnhum.2016.00601. DOI: 10.3389/fnhum.2016.00601.

D. Baraff and A. Witkin. 1998. Large steps in cloth simulation. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '98, p. 43–54. Association for Computing Machinery, New York, NY, USA. ISBN 0897919998. https://doi.org/10.1145/280814.280821. DOI: 10.1145/280814.280821.

C. Bartneck, T. Kanda, H. Ishiguro, and N. Hagita. 2009. My robotic doppelgänger – a critical look at the uncanny valley. In *Proc. of Robot and Human Interactive Communication*, pp. 269–276.

R. Beira, M. Lopes, M. Praça, J. Santos-Victor, A. Bernardino, G. Metta, F. Becchi, and R. Saltarén. 2006. Design of the robot-cub (icub) head. In *Proceedings 2006 IEEE International Conference on Robotics and Automation, 2006. ICRA 2006.*, pp. 94–100. IEEE.

T. Belpaeme, J. Kennedy, A. Ramachandran, B. Scassellati, and F. Tanaka. 2018. Social robots for education: A review. *Science robotics*, 3(21).

P. Bérard, D. Bradley, M. Nitti, T. Beeler, and M. Gross. Nov. 2014. High-quality capture of eyes. *ACM Trans. Graph.*, 33(6). ISSN 0730-0301. https://doi.org/10.1145/2661229.2661285. DOI: 10.1145/2661229.2661285.

C. L. Bethel and R. R. Murphy. 2007. Survey of non-facial/non-verbal affective expressions for appearance-constrained robots. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 38(1): 83–92.

T. W. Bickmore, L. M. Pfeifer, and B. W. Jack. 2009. Taking the time to care: empowering low health literacy hospital patients with virtual nurse agents. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pp. 1265–1274.

F. Biocca. 1997. The cyborg's dilemma: Progressive embodiment in virtual environments [1]. *Journal of Computer-Mediated Communication*, 3(2): 0–0.

M. Blow, K. Dautenhahn, A. Appleby, C. L. Nehaniv, and D. Lee. 2006a. The art of designing robot faces: Dimensions for human-robot interaction. In *Proceedings of the 1st ACM SIGCHI/SIGART conference on Human-robot interaction*, pp. 331–332.

M. Blow, K. Dautenhahn, A. Appleby, C. L. Nehaniv, and D. C. Lee. 2006b. Perception of robot smiles and dimensions for human-robot interaction design. In *ROMAN 2006-The 15th IEEE International Symposium on Robot and Human Interactive Communication*, pp. 469–474. IEEE.

T. J. Burleigh, J. R. Schoenherr, and G. L. Lacroix. 2013. Does the Uncanny Valley exist? An empirical test of the relationship between eeriness and the human likeness of digitally created faces. *Computers in Human Behavior*, 29(3): 759–771.

E. J. Carter, M. Mahler, and J. K. Hodgins. 2013. Unpleasantness of animated characters increases viewer attention to faces. In *Proceedings of the ACM Symposium in Applied Perception*, pp. 35–40.

J. Cassell. 2000. Embodied conversational interface agents. *Communications of the ACM*, 43(4): 70–78.

J. Cassell. 2001. Embodied conversational agents: representation and intelligence in user interfaces. *AI magazine*, 22(4): 67–67.

E. Catmull. 1974. A subdivision algorithm for computer display of curved surfaces. *PhD thesis, Dept. of CS, University of Utah*.

T. Chaminade, J. Hodgins, and M. Kawato. 2007. Anthropomorphism influences perception of computer-animated characters' actions. *Social Cognitive and Affective Neuroscience*, 2(3).

M. Cheetham and L. Jancke. 2013. Perceptual and category processing of the uncanny valley hypothesis' dimension of human likeness: some methodological issues. *Journal of visualized experiments: JoVE*, (76).

B. Colligan. 2011 (accessed June 30, 2020). *How the Knowledge Navigator video came about*. http://www.dubberly.com/articles/how-the-knowledge-navigator-video-came-about.html.

J. E. Cutting and L. T. Kozlowski. 1977. Recognizing friends by their walk: Gait perception without familiarity cues. *Bulletin of the psychonomic society*, 9(5): 353–356.

Z. Dai and K. F. MacDorman. 2018. The doctor's digital double: how warmth, competence, and animation promote adherence intention. *PeerJ Computer Science*, 4: e168.

K. Dautenhahn, C. L. Nehaniv, M. L. Walters, B. Robins, H. Kose-Bagci, N. Assif, M. Blow, et al. 2009. Kaspar–a minimally expressive humanoid robot for human–robot interaction research. *Applied Bionics and Biomechanics*, 6(3, 4): 369–397.

E. Deng, B. Mutlu, and M. Mataric. 2019. Embodiment in socially interactive robots. *arXiv preprint arXiv:1912.00312*.

C. Diana and A. L. Thomaz. 2011. The shape of simon: creative design of a humanoid robot shell. In *CHI'11 Extended Abstracts on Human Factors in Computing Systems*, pp. 283–298.

Didimo, 2019. The breathtaking reality of your digital you. https://mydidimo.com/. https://mydidimo.com/.

C. F. DiSalvo, F. Gemperle, J. Forlizzi, and S. Kiesler. 2002. All robots are not created equal: the design and perception of humanoid robot heads. In *Proceedings of the 4th conference on Designing interactive systems: processes, practices, methods, and techniques*, pp. 321–326.

Epic Games, Inc., Mar 2018. Siren. https://www.3lateral.com/projects/siren.html. https://www.3lateral.com/projects/siren.html.

F. Eyssel and F. Hegel. 2012. (s) he's got the look: Gender stereotyping of robots 1. *Journal of Applied Social Psychology*, 42(9): 2213–2230.

A. E. Ferrey, T. J. Burleigh, and M. J. Fenske. 2015. Stimulus-category competition, inhibition, and affective devaluation: a novel account of the Uncanny Valley. *Frontiers in Psychology*, 6: 249.

Y. Ferstl and R. McDonnell. 2018. A perceptual study on the manipulation of facial features for trait portrayal in virtual agents. In *Proc. of Int. Conf. on Intelligent Virtual Agents (IVA)*, pp. 281–288. DOI: 10.1145/3267851.3267891.

Y. Ferstl, E. Kokkinara, and R. McDonnell. 2017. Facial features of non-player creatures can influence moral decisions in video games. *ACM Transaction on Applied Perception*, 15(1): 4:1–4:12. ISSN 1544-3558. DOI: 10.1145/3129561.

R. Fleming, B. J. Mohler, J. Romero, M. J. Black, and M. Breidt. 2016. Appealing female avatars from 3D body scans: Perceptual effects of stylization. In *Int. Conf. on Computer Graphics Theory and Applications (GRAPP)*.

R. Fribourg, F. Argelaguet, A. Lécuyer, and L. Hoyet. 2020. Avatar and sense of embodiment: Studying the relative preference between appearance, control and point of view. *IEEE Transactions on Visualization and Computer Graphics*, 26(5): 2062–2072. DOI: 10.1109/TVCG.2020.2973077.

M. Garau, M. Slater, V. Vinayagamoorthy, A. Brogni, A. Steed, and M. A. Sasse. 2003. The impact of avatar realism and eye gaze control on perceived quality of communication in a shared immersive virtual environment. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pp. 529–536. ACM.

M. Garau, M. Slater, D.-P. Pertaub, and S. Razzaque. 2005. The responses of people to virtual humans in an immersive virtual environment. *Presence: Teleoperators & Virtual Environments*, 14(1): 104–116.

T. Geller. 2008. Overcoming the Uncanny Valley. *IEEE Computer Graphics and Applications*, 28(4): 11–17.

R. Gockley, A. Bruce, J. Forlizzi, M. Michalowski, A. Mundell, S. Rosenthal, B. Sellner, R. Simmons, K. Snipes, A. C. Schultz, et al. 2005. Designing robots for long-term social interaction. In *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 1338–1343. IEEE.

J. Goetz, S. Kiesler, and A. Powers. 2003. Matching robot appearance and behavior to tasks to improve human-robot cooperation. In *The 12th IEEE International Workshop on Robot and Human Interactive Communication, 2003. Proceedings. ROMAN 2003.*, pp. 55–60. Ieee.

H. Gouraud. 1971. Continuous shading of curved surfaces. *IEEE Transactions on Computers*, C-20(6): 623–629.

R. D. Green, K. F. MacDorman, C.-C. Ho, and S. Vasudevan. 2008. Sensitivity to the proportions of faces that vary in human likeness. *Computers in Human Behavior*, 24(5): 2456–2474.

P. Hamet and J. Tremblay. 2017. Artificial intelligence in medicine. *Metabolism*, 69: S36–S40.

D. Hanson. 2005. Expanding the aesthetics possibilities for humanlike robots. In *Proc. of IEEE Humanoid Robotics Conf., Special Session on the Uncanny Valley*.

H. Hedayati, M. Walker, and D. Szafir. 2018. Improving collocated robot teleoperation with augmented reality. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interac-*

tion, pp. 78–86.

F. Hegel, F. Eyssel, and B. Wrede. 2010. The social robot 'flobi': key concepts of industrial design. In *19th International Symposium in Robot and Human Interactive Communication*, pp. 107–112. IEEE.

U. Hess, R. B. Adams, and R. E. Kleck. 2004. Facial appearance, gender, and emotion expression. *Emotion*, 4(4): 378–388.

J. K. Hodgins, J. F. O'Brien, and J. Tumblin. Dec. 1998. Perception of human motion with different geometric models. *IEEE Transactions on Visualization and Computer Graphics*, 4(4): 101–113. http://graphics.cs.berkeley.edu/papers/Hodgins-PHM-1998-12/.

G. Hoffman, O. Zuckerman, G. Hirschberger, M. Luria, and T. Shani-Sherman. 2015. Design and evaluation of a peripheral robotic conversation companion. In *2015 10th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pp. 3–10. IEEE.

M. Hoque, M. Courgeon, J.-C. Martin, B. Mutlu, and R. W. Picard. 2013. Mach: My automated conversation coach. In *Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing*, pp. 697–706.

L. Hu, S. Saito, L. Wei, K. Nagano, J. Seo, J. Fursund, I. Sadeghi, C. Sun, Y.-C. Chen, and H. Li. Nov. 2017. Avatar digitization from a single image for real-time rendering. *ACM Transactions on Graphics (TOG)*, 36(6): 195:1–195:14. ISSN 0730-0301.

Y. Hu, Z. Zhao, A. Vimal, and G. Hoffman. 2018. Soft skin texture modulation for social robotics. In *2018 IEEE International Conference on Soft Robotics (RoboSoft)*, pp. 182–187. IEEE.

O. Inbar and J. Meyer. 2019. Politeness counts: Perceptions of peacekeeping robots. *IEEE Transactions on Human-Machine Systems*, 49(3): 232–240.

itSeez3D, 2020. Turn your mobile device into a powerful 3d scanner. https://itseez3d.com/. https://itseez3d.com/.

Y. Iwamura, M. Shiomi, T. Kanda, H. Ishiguro, and N. Hagita. 2011. Do elderly people prefer a conversational humanoid as a shopping assistant partner in supermarkets? In *Proceedings of the 6th international conference on Human-robot interaction*, pp. 449–456.

D. Jennings and M. Figliozzi. 2019. Study of sidewalk autonomous delivery robots and their potential impacts on freight efficiency and travel. *Transportation Research Record*, 2673(6): 317–326.

H. W. Jensen, S. R. Marschner, M. Levoy, and P. Hanrahan. 2001. A practical model for subsurface light transport. In *Proc. of SIGGRAPH*, pp. 511–518.

J. Jimenez, V. Sundstedt, and D. Gutierrez. 2009. Screen-space perceptual rendering of human skin. *ACM Transactions on Applied Perception*, 6(4): 23:1–23:15.

J. Jimenez, T. Scully, N. Barbosa, C. Donner, X. Alvarez, T. Vieira, P. Matts, V. Orvalho, D. Gutierrez, and T. Weyrich. 2010. A practical appearance model for dynamic facial color. *ACM Transactions on Graphics*, 29(6): 141:1–141:10.

G. Johansson. 1973. Visual perception of biological motion and a model for its analysis. *Perception & Psychophysics*, 14(2): 201–211.

S. Junuzovic, K. Inkpen, J. Tang, M. Sedlins, and K. Fisher. 10 2012. To see or not to see: A study comparing four-way avatar, video, and audio conferencing for work. pp. 31–34. DOI: 10.1145/2389176.2389181.

J. T. Kajiya and T. L. Kay. 1989. Rendering fur with three dimensional textures. *SIGGRAPH Computer Graphics*, 23(3): 271–280.

A. Kalegina, G. Schroeder, A. Allchin, K. Berlin, and M. Cakmak. 2018. Characterizing the design space of rendered robot faces. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, pp. 96–104.

S. Kiesler, A. Powers, S. R. Fussell, and C. Torrey. 2008. Anthropomorphic interactions with a robot and robot–like agent. *Social Cognition*, 26(2): 169–181.

K. Kilteni, R. Groten, and M. Slater. 11 2012. The sense of embodiment in virtual reality. *Presence Teleoperators and Virtual Environments*, 21. DOI: 10.1162/PRES_a_00124.

E. Kokkinara and R. McDonnell. 2015. Animation realism affects perceived character appeal of a self-virtual face. In *Proceedings of the 8th ACM SIGGRAPH Conference on Motion in Games*, pp. 221–226. Acm.

H. Kozima, M. P. Michalowski, and C. Nakagawa. 2009. Keepon. *International Journal of Social Robotics*, 1(1): 3–18.

L. T. Kozlowski and J. E. Cutting. 1977. Recognizing the sex of a walker from a dynamic point-light display. *Perception & Psychophysics*, 21(6): 575–580.

S. Kühn, T. R. Brick, B. C. Müller, and J. Gallinat. 2014. Is this car looking at you? how anthropomorphism predicts fusiform face area activation when seeing cars. *PloS one*, 9(12): e113885.

M. K. Lee, J. Forlizzi, P. E. Rybski, F. Crabbe, W. Chung, J. Finkle, E. Glaser, and S. Kiesler. 2009. The snackbot: documenting the design of a robot for long-term human-robot interaction. In *Proceedings of the 4th ACM/IEEE international conference on Human robot interaction*, pp. 7–14.

M. K. Lee, S. Kiesler, and J. Forlizzi. 2010. Receptionist or information kiosk: how do people talk with a robot? In *Proceedings of the 2010 ACM conference on Computer supported cooperative work*, pp. 31–40.

J. C. Lester, C. B. Callaway, B. Stone, and S. G. Towns. 1997. Mixed initiative problem solving with animated pedagogical agents. In *Workshop on Pedagogical Agents*, volume 19.

J. Li, R. Kizilcec, J. Bailenson, and W. Ju. 2016. Social robots and virtual agents as lecturers for video instruction. *Computers in Human Behavior*, 55: 1222 – 1230. DOI: https://doi.org/10.1016/j.chb.2015.04.005.

T. Liu, A. W. Bargteil, J. F. O'Brien, and L. Kavan. Nov. 2013. Fast simulation of mass-spring systems. *ACM Transactions on Graphics*, 32(6): 209:1–7. http://cg.cis.upenn.edu/publications/Liu-FMS. Proceedings of ACM SIGGRAPH Asia 2013, Hong Kong.

Loom.ai, 2020. 3d avatar platform for enterprise and developers. https://loomai.com/. https://loomai.com/.

I. Lütkebohle, F. Hegel, S. Schulz, M. Hackel, B. Wrede, S. Wachsmuth, and G. Sagerer. 2010. The bielefeld anthropomorphic robot head "flobi". In *2010 IEEE International Conference on Robotics and Automation*, pp. 3384–3391. IEEE.

K. F. MacDorman, R. D. Green, C.-C. Ho, and C. T. Koch. 2009. Too real for comfort? Uncanny responses to computer generated faces. *Computers in Human Behavior*, 25(3): 695–710.

S. Marschner and P. Shirley. 2016. *Fundamentals of Computer Graphics*. CRC Press.

S. R. Marschner, H. W. Jensen, M. Cammarano, S. Worley, and P. Hanrahan. 2003. Light scattering from human hair fibers. *ACM Transaction on Graphics*, 22(3): 780–791. ISSN 0730-0301.

T. Masson. 2007. *CG 101: A Computer Graphics Industry Reference*. Digital Fauxtography.

S. McCloud. 1993. Understanding comics: The invisible art. *Northampton, Mass*.

R. McDonnell, S. Jörg, J. K. Hodgins, F. Newell, and C. O'sullivan. 2009a. Evaluating the effect of motion and body shape on the perceived sex of virtual characters. *ACM Transactions on Applied Perception (TAP)*, 5(4): 20.

R. McDonnell, S. Jörg, J. McHugh, F. N. Newell, and C. O'Sullivan. 2009b. Investigating the role of body shape on the perception of emotion. *ACM Transactions on Applied Perception (TAP)*, 6(3): 14.

R. McDonnell, M. Larkin, B. Hernandez, I. Rudomin, , and C. O'Sullivan. 2009c. Eye-catching Crowds: saliency based selective variation. *ACM Transaction on Graphics*, 28(3): 55:1 – 55:10.

R. McDonnell, M. Breidt, and H. H. Bülthoff. 2012. Render me real? Investigating the effect of render style on the perception of animated virtual humans. *ACM Transaction on Graphics*, 31(4): 91:1–91:11.

J. E. Michaelis and B. Mutlu. 2018. Reading socially: Transforming the in-home reading experience with a learning-companion robot. *Science Robotics*, 3(21).

J. Miseikis, P. Caroni, P. Duchamp, A. Gasser, R. Marko, N. Miseikiene, F. Zwilling, C. de Castelbajac, L. Eicher, M. Fruh, et al. 2020. Lio–a personal robot assistant for human-robot interaction and care applications. *arXiv preprint arXiv:2006.09019*.

M. Mori. 1970. The uncanny valley. *Energy*, 7(4): 33 – 35.

J. Mumm and B. Mutlu. 2011. Designing motivational agents: The role of praise, social comparison, and embodiment in computer feedback. *Computers in Human Behavior*, 27(5): 1643–1650.

B. Mutlu and J. Forlizzi. 2008. Robots in organizations: the role of workflow, social, and environmental factors in human-robot interaction. In *2008 3rd ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pp. 287–294. IEEE.

K. Nagano, J. Seo, J. Xing, L. Wei, Z. Li, S. Saito, A. Agarwal, J. Fursund, H. Li, R. Roberts, et al. 2018. pagan: real-time avatars using dynamic textures. *ACM Trans. Graph.*, 37(6): 258–1.

C. Nass, Y. Moon, and N. Green. 1997. Are machines gender neutral? gender-stereotypic responses to computers with voices. *Journal of applied social psychology*, 27(10): 864–876.

F. E. Nicodemus, J. C. Richmond, J. J. Hsia, I. W. Ginsberg, and T. Limperis. 1992. Geometrical considerations and nomenclature for reflectance. 160: 4.

S. Nishio, H. Ishiguro, and N. Hagita. 2007. Geminoid: Teleoperated android of an existing person. *Humanoid robots: New developments*, 14: 343–352.

K. Nowak. 2001. The influence of anthropomorphism on social judgment in social virtual environments. In *Annual Convention of the International Communication Association, Washington, DC*.

J.-H. Oh, D. Hanson, W.-S. Kim, Y. Han, J.-Y. Kim, and I.-W. Park. 2006. Design of android type humanoid robot albert hubo. In *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 1428–1433. IEEE.

C. D. Onal, M. T. Tolley, R. J. Wood, and D. Rus. 2014. Origami-inspired printed robots. *IEEE/ASME transactions on mechatronics*, 20(5): 2214–2221.

I. C. Orr. 1974. Puppet theatre in asia. *Asian Folklore Studies*, pp. 69–84.

A. G. Ozkil, Z. Fan, S. Dawids, H. Aanes, J. K. Kristensen, and K. H. Christensen. 2009. Service robots for hospitals: A case study of transportation tasks in a hospital. In *2009 IEEE international conference on automation and logistics*, pp. 289–294. IEEE.

N. A. Palomares and E.-J. Lee. 2010. Virtual gender identity: The linguistic assimilation to gendered avatars in computer-mediated communication. *Journal of Language and Social Psychology*, 29(1): 5–23. DOI: 10.1177/0261927X09351675.

Y. Pan and A. Steed. 2016. A comparison of avatar-, video-, and robot-mediated interaction on users trust in expertise. *Frontiers in Robotics and AI*, 3: 12. DOI: 10.3389/frobt.2016.00012.

T. Pejsa, B. Mutlu, and M. Gleicher. 2013. Stylized and performative gaze for character animation. *Computer Graphics Forum*, 32(2pt2): 143–152. https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.12034. DOI: 10.1111/cgf.12034.

B. T. Phong. 1975. Illumination for computer generated pictures. *Communications of ACM*, 18(6): 311–317.

Pinscreen, 2019. The most advanced ai-driven personalized avatars. https://www.pinscreen.com/. https://www.pinscreen.com/.

A. Powers and S. Kiesler. 2006. The advisor robot: tracing people's mental model from a robot's physical attributes. In *Proceedings of the 1st ACM SIGCHI/SIGART conference on Human-robot interaction*, pp. 218–225.

J. Pransky. 2001. Aibo–the no. 1 selling service robot. *Industrial robot: An international journal*.

T. J. Prescott, B. Mitchinson, and S. Conran. 2017. Miro: An animal-like companion robot with a biomimetic brain-based control system. In *Proceedings of the Companion of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*, pp. 50–51.

D. Romano, C. Pfeiffer, A. Maravita, and O. Blanke. 3 2014. Illusory self-identification with an avatar reduces arousal responses to painful stimuli. *Behavioural Brain Research*, 261: 275–281. ISSN 01664328. DOI: 10.1016/j.bbr.2013.12.049.

K. Ruhland, K. Zibrek, and R. McDonnell. 2015. Perception of personality through eye gaze of realistic and cartoon models. In *Proc. of Symp. on Applied Perception*, pp. 19–23. ACM.

S. Saito, L. Wei, L. Hu, K. Nagano, and H. Li. 2017. Photorealistic facial texture inference using deep neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5144–5153.

Y. Sakagami, R. Watanabe, C. Aoyama, S. Matsunaga, N. Higaki, and K. Fujimura. 2002. The intelligent asimo: System overview and integration. In *IEEE/RSJ international conference on intelligent robots and systems*, volume 3, pp. 2478–2483. IEEE.

A. Sauppé and B. Mutlu. 2015. The social impact of a robot co-worker in industrial settings. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems*, pp. 3613–3622.

A. P. Saygin, T. Chaminade, H. Ishiguro, J. Driver, and C. Frith. 2012. The thing that should not be: Predictive coding and the Uncanny Valley in perceiving human and humanoid robot actions. *Social Cognitive Affective Neuroscience*, 7(4): 413–422.

J. Scarce. 1983. Karagoz shadow puppets of turkey.

J. Sculley. 1989. The relationship between business and higher education: A perspective on the 21st century. *Commun. ACM*, 32(9): 1056–1061. https://doi.org/10.1145/66451.66452. DOI: 10.1145/66451.66452.

J. Seyama and R. S. Nagayama. 2007. The Uncanny Valley: Effect of realism on the impression of artificial human faces. *Presence: Teleoperators and Virtual Environments*, 16(4): 337–351.

M. Seymour, C. Evans, and K. Libreri. 2017. Meet mike: epic avatars. In *ACM SIGGRAPH 2017 VR Village*, pp. 1–2.

M. Shayganfar, C. Rich, and C. L. Sidner. 2012. A design methodology for expressing emotion on robot faces. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 4577–4583. IEEE.

M. Siegel, C. Breazeal, and M. I. Norton. 2009. Persuasive robotics: The influence of robot gender on human behavior. In *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 2563–2568. IEEE.

S. Simon, C. William, and G. Jan. 1999. Enlightened automata. In *The Sciences in Enlightened Europe, Chicago and London*. The University of Chicago Press.

M. Slater and A. Steed. 2002. Meeting people virtually: Experiments in shared virtual environments. In *The Social Life of Avatars*, pp. 146–171. Springer.

A. Steed, Y. Pan, F. Zisch, and W. Steptoe. 2016. The impact of a self-avatar on cognitive load in immersive virtual reality. pp. 67–76. DOI: 10.1109/VR.2016.7504689.

E. M. Suzanne R. Pallak and J. Koch. 1983. Communicator attractiveness and expertise, emotional versus rational appeals, and persuasion: A heuristic versus systematic processing interpretation. *Social Cognition*, 2(2): 122–141.

D. Szafir, B. Mutlu, and T. Fong. 2015. Communicating directionality in flying robots. In *2015 10th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pp. 19–26. IEEE.

J. Thies, M. Zollhöfer, M. Stamminger, C. Theobalt, and M. Nießner. 2016. Face2face: Real-time face capture and reenactment of rgb videos. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2387–2395.

F. Thomas and O. Johnston. 1995. *The illusion of life: Disney animation*. Hyperion New York.

A. Tinwell, M. Grimshaw, D. A. Nabi, and A. Williams. 2011. Facial expression of emotion and perception of the Uncanny Valley in virtual characters. *Computers in Human Behavior*, 27(2): 741–749.

I. Torre, E. Carrigan, K. McCabe, R. McDonnell, and N. Harte. 2018. Survival at the museum: A cooperation experiment with emotionally expressive virtual characters. In *Proceedings of the 2018 on International Conference on Multimodal Interaction*, pp. 423–427. ACM.

I. Torre, E. Carrigan, R. McDonnell, K. Domijan, K. McCabe, and N. Harte. 2019. The effect of multimodal emotional expression and agent appearance on trust in human-agent interaction. In *Motion, Interaction and Games*, MIG '19. Association for Computing Machinery, New York, NY, USA. ISBN 9781450369947. https://doi.org/10.1145/3359566.3360065. DOI: 10.1145/3359566.3360065.

I. P. Tussyadiah and S. Park. 2018. Consumer evaluation of hotel service robots. In *Information and communication technologies in tourism 2018*, pp. 308–320. Springer.

A. van Breemen, X. Yan, and B. Meerbeek. 2005. icat: an animated user-interface robot with personality. In *Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems*, pp. 143–144.

A. J. van Breemen. 2004. Animation engine for believable interactive user-interface robots. In *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)(IEEE Cat. No. 04CH37566)*, volume 3, pp. 2873–2878. IEEE.

G. Veletsianos. Sept. 2010. Contextually relevant pedagogical agents: Visual appearance, stereotypes, and first impressions and their impact on learning. *Comput. Educ.*, 55(2): 576–585. ISSN 0360-1315. https://doi.org/10.1016/j.compedu.2010.02.019. DOI: 10.1016/j.compedu.2010.02.019.

V. Vinayagamoorthy, M. Gillies, A. Steed, E. Tanguy, X. Pan, C. Loscos, and M. Slater. 2006. Building Expression into Virtual Characters. In B. Wyvill and A. Wilkie, eds., *Eurographics 2006 - State of the Art Reports*. The Eurographics Association. DOI: 10.2312/egst.20061052.

M. Volante, S. V. Babu, H. Chaturvedi, N. Newsome, E. Ebrahimi, T. Roy, S. B. Daily, and T. Fasolino. 2016. Effects of virtual human appearance fidelity on emotion contagion in affective inter-personal simulations. *IEEE Transaction on Visualization and Computer Graphics*, 22(4): 1326–1335.

K. Wada and T. Shibata. 2007. Living with seal robots—its sociopsychological and physiological influences on the elderly at a care house. *IEEE transactions on robotics*, 23(5): 972–980.

K. Wada, T. Shibata, T. Saito, K. Sakamoto, and K. Tanie. 2005. Psychological and social effects of one year robot assisted activity on elderly people at a health service facility for the aged. In *Proceedings of the 2005 IEEE international conference on robotics and automation*, pp. 2785–2790. IEEE.

M. Walker, H. Hedayati, J. Lee, and D. Szafir. 2018. Communicating robot motion intent with augmented reality. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, pp. 316–324.

C. Wallraven, H. H. Bülthoff, D. W. Cunningham, J. Fischer, and D. Bartz. 2007. Evaluation of real-world and computer-generated stylized facial expressions. *ACM Transactions on Applied Perception*, 4(3).

Y. Wang, J. Geigel, and A. Herbert. 2013. Reading personality: Avatar vs. human faces. In *Humaine Association Conference on Affective Computing and Intelligent Interaction*, pp. 479–484. DOI: 10.1109/ACII.2013.85.

M. Watanabe, K. Ogawa, and H. Ishiguro. 2015. Can androids be salespeople in the real world? In *Proceedings of the 33rd annual ACM conference extended abstracts on human factors in computing systems*, pp. 781–788.

D. Whitley. 2012. *The idea of nature in Disney animation: From Snow White to WALL-E*. Ashgate Publishing, Ltd.

M. Wise, M. Ferguson, D. King, E. Diehr, and D. Dymesich. 2016. Fetch and freight: Standard platforms for service robot applications. In *Workshop on autonomous mobile service robots*.

P. Wisessing, K. Zibrek, D. W. Cunningham, J. Dingliana, and R. McDonnell. Apr. 2020. Enlighten me: Importance of brightness and shadow for character emotion and appeal. *ACM Trans. Graph.*, 39(3). ISSN 0730-0301. https://doi.org/10.1145/3383195. DOI: 10.1145/3383195.

Y. Yamada, T. Kawabe, and K. Ihaya. 01 2013. Categorization difficulty is associated with negative evaluation in the "uncanny valley" phenomenon. *Japanese Psychological Research*, 55: 20–32. DOI: 10.1111/j.1468-5884.2012.00538.x.

S. Yamaguchi, S. Saito, K. Nagano, Y. Zhao, W. Chen, K. Olszewski, S. Morishima, and H. Li. 2018. High-fidelity facial reflectance and geometry inference from an unconstrained image. *ACM Transactions on Graphics (TOG)*, 37(4): 1–14.

B. M. Yamauchi. 2004. Packbot: a versatile platform for military robotics. In *Unmanned ground vehicle technology VI*, volume 5422, pp. 228–237. International Society for Optics and Photonics.

N. Yee and J. Bailenson. 07 2007. The proteus effect: The effect of transformed self-representation on behavior. *Human Communication Research*, 33: 271 – 290. DOI: 10.1111/j.1468-2958.2007.00299.x.

N. Yee and J. N. Bailenson. 2009. The difference between being and seeing: The relative contribution of self-perception and priming to behavioral changes via digital self-representation. *Media Psychology*, 12(2): 195–209. DOI: 10.1080/15213260902849943.

N. Yee, N. Ducheneaut, M. Yao, and L. Nelson. 05 2011. Do men heal more when in drag? conflicting identity cues between user and avatar. pp. 773–776. DOI: 10.1145/1978942.1979054.

Y. Yokota. 2009. A historical overview of japanese clocks and karakuri. In *International Symposium on History of Machines and Mechanisms*, pp. 175–188. Springer.

E. Zell, C. Aliaga, A. Jarabo, K. Zibrek, D. Gutierrez, R. McDonnell, and M. Botsch. 2015. To stylize or not to stylize?: The effect of shape and material stylization on the perception of computer-generated faces. *ACM Transactions on Graphics*, 34(6): 184:1–184:12.

E. Zell, K. Zibrek, and R. McDonnell. 2019. Perception of virtual characters. In *ACM SIGGRAPH 2019 Courses*, SIGGRAPH '19. Association for Computing Machinery, New York, NY, USA. ISBN 9781450363075. https://doi.org/10.1145/3305366.3328101. DOI: 10.1145/3305366.3328101.

K. Zibrek and R. McDonnell. 2019. Social presence and place illusion are affected by photorealism in embodied vr. In *Motion, Interaction and Games*, MIG '19. Association for Computing Machinery, New York, NY, USA. ISBN 9781450369947. https://doi.org/10.1145/3359566.3360064. DOI: 10.1145/3359566.3360064.

K. Zibrek, L. Hoyet, K. Ruhland, and R. McDonnell. 2015. Exploring the effect of motion type and emotions on the perception of gender in virtual humans. *ACM Transactions on Applied Perception (TAP)*, 12(3): 11.