

UNIVERSIDADE DE LISBOA
FACULDADE DE BELAS-ARTES



LISBOA

UNIVERSIDADE
DE LISBOA

Virtual Avatars

Creating expressive embodied characters for Virtual Reality

Diogo José Martins dos Santos

Dissertação

Mestrado em Arte Multimédia

Especialização em Artes Interativas

Dissertação orientada pela Professora Mónica Mendes

2023

Declaração de autoria

Eu Diogo José Martins dos Santos, declaro que a presente dissertação intitulada “Virtual Avatars: Creating expressive embodied characters for Virtual Reality”, é o resultado da minha investigação pessoal e independente. O conteúdo é original e todas as fontes consultadas estão devidamente mencionadas na bibliografia ou outras listagens de fontes documentais, tal como todas as citações diretas ou indiretas têm devida indicação ao longo do trabalho segundo as normas académicas.

O Candidato

Lisboa, 30/10/2023

Abstract

This dissertation presents a record of the experiences gained through the creation of avatars for Virtual Reality. As such, it reflects on that information through research on the history behind the concept of avatars, their context within Virtual Reality, various methods of translating real-world information into the virtual, and how that may affect the user's perception of their own virtual body and identity.

In the second part of this project, with the newly researched knowledge in mind, I present "Morph", an avatar I created as part of this research considering the goals to maximize avatar creativity and enhance user expressiveness. The report of the detailed creation process carried out and corresponding documentation goes through various steps on how the avatar was developed, from planning and concepting to various 3D modeling steps and, finally, applying it to the game engine.

I conclude with considerations and reflections on the project itself and the research that supported it, including observations on the current situation of avatars, as well as a reflection on their future potential.

Keywords

Virtual Reality, Avatar, Presence, 3D Modeling, Expressivity.

Resumo alargado

Esta dissertação de mestrado apresenta as experiências adquiridas através da criação de avatares para Realidade Virtual. A documentação é sustentada com uma investigação sobre a história do conceito de “avatar”, bem como com a sua contextualização no âmbito da Realidade Virtual. Apresenta, também, várias formas de converter aspectos da vida real para o mundo virtual e como ambos podem afetar o utilizador que incorpora o avatar e a sua identidade visual. Tendo como base a informação adquirida, idealizo e concebo, no desenvolvimento desta dissertação, um avatar original intitulado “Morph”, que vem exemplificar a investigação efetuada. A criação do avatar “Morph” envolveu a documentação do processo em si que contém toda a evolução da sua produção, desde o seu planeamento e conceito, até à sua versão finalizada e pronta para ser utilizada no ambiente de Realidade Virtual. Durante este processo, os objetivos da criação do avatar foram maximizar a sua criatividade e a expressividade do utilizador que o incorpora.

Numa fase introdutória, procuro ir às origens do termo “Avatar”, proveniente do hinduísmo, cujo significado está puramente relacionado com encarnações divinas dos deuses hindus (Parrinder, 1982, pp. 14-21). Na atualidade, e similarmente à sua origem hindu, a etimologia de avatar é enfatizada pela personificação de ícones de identidade. Encontrados em várias histórias e, maioritariamente, em plataformas de redes sociais online, estes avatares modernos têm a função de aparência visual para a auto expressão e manifestação de identidade do utilizador. Devido às características anónimas inerentes das interações online de determinadas plataformas, os avatares podem também definir-se como personalidades alternativas adaptadas por indivíduos que se queiram separar da sua identidade real (Thomson & Greenwood, 2020, p. 349).

Uma nova identidade, uma nova cara. O rosto humano desempenha um papel vital na identidade e na comunicação, tanto em ambientes reais como virtuais, auxiliando conexões sociais ao expressar emoções (Belting, 2017, pp. 106-107). A tendência do cérebro humano para percecionar rostos e formas específicas, mesmo em padrões visuais aleatórios ou ambíguos (pareidolia), ajuda a simplificar a comunicação em espaços virtuais e a reconhecer emoções numa personagem não-humana, num avatar. Porém, para se tirar maior proveito desta relação entre comunicação e identidade em ambientes virtuais, a tecnologia de captura de movimento facial deverá estar o mais otimizada possível. O seu papel fundamental em converter a informação real do rosto do utilizador para expressões realistas em avatares, oferece nuances difíceis de replicar manualmente. Para além do rastreamento da cara,

também o rastreamento do corpo se destaca como uma tecnologia de extrema importância para a expressividade no meio da Realidade Virtual, devido à capacidade de o próprio corpo exprimir comunicação auxiliar, a chamada linguagem corporal. Tudo culmina numa experiência sensorial de estar fisicamente presente num espaço virtual – a “presença” (Skarbez, Smith, & Whitton, 2021, pp. 4-7). Contudo, quando o nível de presença é alto o suficiente, é possível que o utilizador comece a experienciar sensações táteis dentro da Realidade Virtual, como apertos de mão ou abraços, mesmo quando não haja contato físico presente, um fenómeno com o nome de “Phantom Touch” (Pilacinski, Metzler, & Klaes, 2023, p. 24-26[b]).

No contexto da sua aparência virtual, podemos categorizar os avatares realçando duas finalidades: a capacidade de replicar alguém específico ou a de criar uma nova personalidade única. No entanto, esta dissertação foca-se na criação de um avatar com a sua própria aparência e personalidade, isto porque, quando é atribuída a avatares a sua própria personalidade, diferente da do utilizador, existe a tendência para que esta influencie o comportamento deste último (Slater & Sanchez-Vives, 2014, p. 24).

O desenvolvimento desta dissertação é inicialmente suportado pelo processo de criação de um avatar, através do planeamento dos seus objetivos e características, nomeadamente, a atribuição de uma personalidade única e a sua elevada capacidade de expressão. Para além destes, foi escolhido e aplicado o antropomorfismo devido, não só às suas semelhanças, mas também, antagonicamente, às suas diferenças relativamente ao aspeto humano, o que permitiu a experimentação e análise destas em contexto de Realidade Virtual.

A consequente fase de conceito envolveu a criação da identidade do avatar, a partir da identificação do seu objetivo, ancorado na investigação realizada, e por inspiração em avatares preexistentes, de forma a compreender que aspetos são mais importantes para os utilizadores. Seguidamente, são criados esboços com o foco de dar vida ao avatar, infundindo-lhe vitalidade e dinamismo. Estes abrangem várias perspectivas e servem de referência para o desenvolvimento da personagem e, simultaneamente, para o desenvolvimento do modelo 3D na próxima fase, a estruturação. Nesta, o modelo 3D do avatar é construído com base nos esboços das vistas frontal e lateral. Começa-se com uma forma simples que depois é extrudida e manipulada para corresponder à forma geral prevista nos esboços. Este processo assegura que as proporções e a estrutura do modelo se alinham com os esboços iniciais e criam uma base sólida para o desenvolvimento posterior.

Na fase de escultura, a seguinte à estruturação, o processo assemelha-se em grande parte a escultura física, uma vez que é adicionado ou removido material digital para moldar a forma final do avatar. Esta fase envolve o refinamento gradual do modelo 3D, utilizando ferramentas de escultura digital para capturar detalhes mais pormenorizados e texturas que melhoram a aparência visual e o realismo do avatar. A fase subsequente de modelação “hard-surface” baseia-se nos fundamentos estabelecidos na fase de escultura anterior, com o objetivo de criar formas precisas e perfeitas. Não obstante, é importante notar que o nível de detalhe alcançado a este ponto pode exigir uma otimização adicional para renderização em tempo real em ambientes de Realidade Virtual. Por este motivo, é necessária a fase de retopologia, na qual o modelo 3D passa de uma versão altamente detalhada para uma versão mais otimizada, adequada a uma renderização eficiente em tempo real. Esta etapa essencial envolve a reformulação da estrutura do modelo para reduzir o número de polígonos, preservando a forma e as características do avatar.

Na fase seguinte, a chamada edição UV, o modelo 3D é preparado para o mapeamento de texturas através da criação de “seams”, ou costuras. Estas delineiam as arestas e permitem que cada parte do avatar seja achatada numa forma bidimensional, garantindo a aplicação precisa de texturas e materiais. Em consequência, a criação das texturas consiste em transferir os detalhes do modelo esculpido para a versão otimizada utilizando a técnica “normal baking”. Trata-se de uma técnica que projeta os pormenores na superfície do avatar para obter uma representação visualmente detalhada, permitindo, desta forma, a renderização em tempo real. Ainda na fase de edição UV, é dado foco à seleção de cores e materiais, adicionando texturas que vão desde o realismo da pele aos reflexos metálicos. Esta etapa confere vitalidade ao modelo e captura a estética desejada dos conceitos originais.

Após criação das texturas, as fases de “rigging” e “weight-painting” centram-se em possibilitar que o avatar se mova naturalmente e reproduza os movimentos do utilizador. O “rigging” estabelece uma estrutura esquelética que define as articulações e os sistemas de controlo, assegurando movimentos realistas no ambiente virtual. Em seguida, o “weight-painting” atribui influências desse esqueleto a diferentes partes do avatar, o que permite um controlo preciso da deformação para movimentos fluidos e naturais.

Com o modelo pronto, e de forma a incorporá-lo na envolvente virtual, é necessária a sua configuração no Unity, o software onde é finalizada a integração do avatar no VRChat, o jogo que suporta o avatar criado e respetivos ambientes virtuais. Esta fase envolve tarefas como a importação de modelos, a configuração de materiais, a implementação da lógica e o carregamento do avatar no jogo. Posteriormente, já perto da conclusão do projeto,

aprofunda-se as nuances da criação de formas de mistura de rastreamento facial para o modelo. Embora não estritamente necessárias para um avatar se tornar apelativo, a integração destas formas de mistura aumentam profundamente as capacidades expressivas do avatar. O objetivo principal é criar um avatar que ultrapasse a mera representação digital e, concomitantemente, promover o sentido de presença e ligação ao utilizador.

A iteração final do projeto marca o ponto culminante da criação do avatar. À medida que diferentes aspetos do avatar são afinados e aperfeiçoados para garantir que este funciona sem falhas, é importante também que diferentes utilizadores testem o avatar e apresentem críticas. Este ensaio permite corrigir imperfeições, não só observadas pelo artista, mas também por outros utilizadores, aumentando a qualidade geral do avatar e concluindo, deste modo, o processo da sua criação.

Concluo a dissertação com considerações e reflexões sobre o projeto em si e a investigação que o sustentou, incluindo observações em torno da atual conjuntura de avatares em ambiente de Realidade Virtual, assim como uma ponderação acerca do potencial desta indústria no futuro.

Palavras-chave:

Realidade Virtual, Avatar, Presença, Modelação 3D, Expressividade.

Index

Abstract.....	2
Resumo alargado.....	3
Introduction.....	7
The concept of “Avatar”.....	8
Importance of the face.....	12
Origins of facial motion capture.....	16
VR face-tracking.....	19
Video games and Standards.....	22
Human replication.....	23
Avatar self-identity.....	27
Full-body tracking.....	30
Presence and “Phantom Touch”.....	32
Avatar expression and embodiment.....	35
The false mainstream.....	37
Avatar creation.....	39
Introduction and Planning.....	39
Concepting.....	42
Structuring.....	49
Sculpting.....	51
Hard-Surface Modeling.....	57
Retopology.....	62
UV Editing.....	71
Texture Creation.....	74
Rigging and Weight-Painting.....	80
Unity Setup.....	83
Face-Tracking.....	86
Iteration.....	91
Closing Thoughts.....	97
Findings, statements and implications.....	97
Limitations.....	99
What was, what is, and what will be?.....	99
Bibliography.....	101

Introduction

This dissertation started with a personal interest in avatars in the context of virtual reality. They have the ability to abstain from real-life context, to offer an escapatory medium in which the person holds the power to change their identity or even to explore aspects of it that they've never discovered before.

The investigation holds broader implications for understanding user-avatar behavior in contemporary society beyond its role as a personal journey. This exploration adds significance in an era where avatars are assuming increasing relevance. Avatars wield a remarkable influence within virtual social interactions and environments, and their potential to shape and impact user behavior cannot be underestimated. As avatars continue to permeate various aspects of our digital lives, the desire for customization and personalization becomes a prevailing trend. Individuals increasingly seek avatars that resonate with their unique sense of self. However, to understand the consequences of avatars in this context, it is important to first explore their history and where the original idea started.

The concept of “Avatar”

In the modern context, an avatar serves as a visual representation of a person's identity. Social media platforms like *Facebook*, *Discord*, and *Reddit* employ avatars as personal icons or profile pictures, consisting of small images portraying users on each platform. These images play a crucial role in person-to-person communication environments, allowing individuals to convey their identity visually and instantly without relying on words. Avatars facilitate quick recognition and enable conversants to gain initial insights into each other. Consequently, users place great importance on their avatars, aiming to represent their online selves uniquely. The desire for a distinct and personalized avatar stems from the need for self-expression and the aspiration to stand out amidst the digital landscape. Individuals capture attention, assert their presence, and establish a recognizable online persona with a unique avatar. These become visual signatures that symbolize individuality, integral to online interactions and social connections. Hence, users carefully select and customize their avatars to create an image that reflects their personal identity and resonates with their online community (Thomson & Greenwood, 2020, p. 349).



Figure 1a. Santos, D. (2023). *A previously created icon to emphasize my personal online persona on multiple social media platforms*. Retrieved January 11, 2023.

However, due to the inherent nature of online communication platforms, users can remain anonymous, concealing certain aspects of their identity or adopting an alternative persona altogether. In this context, the profile icon becomes a tool for further cultivating this constructed personality, a manifestation of a persona that may deviate from the user's true identity, specifically designed for use in a new context or environment. This notion of persona creation is closely tied to the modern understanding of the term "Avatar", which traces its origins back to the Indo-Pakistani religion of Hinduism¹.

The word (*avatāra*, pronounced [əvəˈtɑːrə]), which in Sanskrit (written अवतार) means “descent”. It signifies the material appearance/incarnation of a powerful deity, goddess, or spirit on earth. In other words, it is associated with the acts of various Gods taking a physical form to perform a particular task, usually to bring *dharma* and to enforce righteous behavior and social order (Parrinder, 1982, pp. 14-21). It is possible to notice similarities between the meaning of avatar in a religious context throughout its history and the meaning it mostly has in a modern context. A higher-form being controls a lower-form vessel for a specific purpose in both. Just as in Hinduism, the higher being is a God, and the lower form is a physical vessel; in the modern context, the higher being is a physical person, and the lower form is a conceptual persona, enabling individuals to transcend their physical limitations and explore different aspects of their personality, creativity, and social interactions. While it is also compared to the word “incarnation,” the meaning differs from Hinduism as an incarnation in Christianity implies a new flesh that is imperfect. In contrast, by the Hindu religion's definition, the avatar is a perfect form/vessel. It is ideally suited to allow a God to perform its noble task (Parrinder, 1982, p. 267).

¹ Merriam-Webster. (n.d.). Avatar. In Merriam-Webster.com dictionary. Retrieved October 26, 2023, from <https://www.merriam-webster.com/dictionary/avatar>




Figure 1b. *Avatar - Various Gods in their avatar form.* (2023, January 5). In Wikipedia. <https://en.wikipedia.org/wiki/Avatar>

More recently, the term "avatar" has taken on a new meaning within the context of computer graphics, which is particularly relevant to this research. In computer graphics, an avatar refers to a graphical representation of the user, serving as a distinct visual identity used to represent an individual in the real world. Comparable to a social media platform avatar, a computer graphics avatar also functions as a persona through which users express themselves. However, in this case, the avatar is a virtual character rendered on a screen, representing the user within a digital environment and enabling interactions with other virtual personas. For instance, in online social games like *Second Life*² and *VRChat*³, users can choose their avatars - their virtual appearances visible to others within that specific environment. Through these avatars, individuals engage in conversations and interactions with others. These avatars act as the user's visual representation, allowing them to navigate and participate in the virtual world. Like online social media platforms, users can remain anonymous while embodying their chosen avatar and engaging with other individuals' avatars within the virtual environment. This anonymity and ease of expression allows individuals to curate their digital presence. Exploring *VRChat*, I personally found out that when people select an avatar that aligns with their desired identity, it allows them to experiment with different facets of themselves. In other words, avatars, in this way, serve as a gateway for users to engage with the virtual environment and establish connections with others in a unique and immersive way. The ability to assume an alternate visual identity while interacting within these virtual spaces adds an element of freedom, creativity, and anonymity, enabling users to navigate and shape their virtual experiences according to their unique preferences and desires.

² Rosedale, P. (2003). *Second Life* (PC version) [Video game]. Linden Labs. <https://secondlife.com>

³ Gaylor, G., & Joudrey, J (2014). *VRChat* (PC version) [Video game]. VRChat Inc. <https://hello.vrchat.com>

Nevertheless, the contemporary use of the word avatar to mean “online virtual bodies” was first popularized by Neal Stephenson in his initially published 1992 cyberpunk novel *Snow Crash*⁴, which gained significant popularity upon its initial publication in 1992. This influential book explored the concept of Virtual Reality (VR) as a means to escape from a challenging physical reality. Within the narrative, Stephenson introduced the term "Avatar" and coined the term "*Metaverse*" to describe the virtual reality realm depicted in the story. The term quickly gained recognition and has since become widely used when referring to present-day virtual reality applications. Stephenson's novel provided a vision of a future where individuals could assume digital personas and engage in immersive virtual experiences. By introducing the term "Avatar" in the context of online virtual bodies, Stephenson tapped into the growing fascination with the potential of virtual reality and the desire for individuals to transcend the limitations of their physical existence. The concept of avatars, as presented in "Snow Crash," resonated with readers and played a pivotal role in shaping the discourse around virtual reality technology (Gerhard, Moore & Hobbs, 2004, pp. 453-480).

Furthermore, introducing the Metaverse in the same book further solidified Stephenson's impact on the virtual reality landscape. The term encompasses the idea of a vast virtual universe, a collective space where users can interact, explore, and create. The notion of the Metaverse has captured the imagination of technologists, researchers, and enthusiasts, serving as a conceptual framework for developing virtual reality platforms and applications. Today, the term has had its surge of popularity thanks to Mark Zuckerberg's reference in Meta's 2021 presentation⁵ when discussing the potential of virtual reality and the immersive digital experiences it can offer; however, in my personal opinion, its main purpose seemed only to appeal to invest and shareholder due to the introduction and focus on NFTs⁶ and their questionable nature.

Stephenson's groundbreaking novel popularized the modern version of the term "avatar" and laid the foundation for envisioning virtual reality as a transformative medium. His imaginative portrayal of the Metaverse has influenced subsequent works of literature, media, and technology, shaping the trajectory of virtual reality development and exploration. The cultural significance of "Snow Crash" in introducing these terms demonstrates science

⁴Stephenson, N. (2022). *Snow crash*. DEL REY. Initially published in 1992.

⁵ An excerpt of the livestream. Meta Platforms, Inc. (2021, October 29). Introducing Meta. YouTube. https://www.youtube.com/watch?v=pjNI9K1D_xo

⁶ The acronym for non-fungible token. Merriam-Webster. (n.d.). NFT. In Merriam-Webster.com dictionary. Retrieved October 29, 2023, from <https://www.merriam-webster.com/dictionary/NFT>

fiction literature's influence in shaping our understanding and expectations of emerging technologies (Gerhard, Moore & Hobbs, 2004, pp. 453-480). In this context, the concept of avatars has evolved to become an integral part of our online presence, enabling individuals to shape and project their desired personas.

Importance of the face

Just as avatars in virtual reality represent identity in a summarized, visual form, the face, as the center of physical expression, holds immense importance in capturing the nuances of emotions and facilitating connections in both the real and virtual worlds. The face has always been the center of physical expression merely because it is the place that naturally gathers attention when communication is expected. Moreover, because of its complexity, it is possible to communicate a wide range of emotions simply by moving different sets of muscles. Facial expression is so ingrained in human communication that it is done naturally, without the person thinking about it (Belting, 2017, p. 9).




Figure 2a. Ekman, P. (1967). *A man from New Guinea asked to make different expressions based on different situations*. Paul Ekman Group. Retrieved January 12, 2023, from <https://www.paulekman.com/resources/universal-facial-expressions>

The human face is an extraordinary instrument that empowers us to convey thoughts and feelings, comprehend the emotions of others, and successfully engage in intricate social exchanges. When exploring the significance of the human face in communication, we unravel its remarkable abilities encompassing the expressive conveyance of emotions, the facilitation of nuanced nonverbal cues, the manifestation of personal identity, and the cultivation of empathetic bonds. Through these multifaceted dimensions, the human face becomes an indispensable channel through which we establish connections, understand one another, and forge meaningful relationships. Our facial muscles enable us to display emotions like happiness, sadness, anger, fear, surprise, and disgust. These expressions are almost universally recognized and provide immediate insights into an individual's emotional state.

The face is a visual display of our internal feelings, allowing others to empathize, respond, and adapt their communication accordingly. For example, a smile can signal happiness or friendliness. At the same time, a furrowed brow can indicate concern or confusion. Facial expressions enhance the clarity and depth of emotional communication, facilitating a deeper understanding of one another (Belting, 2017, pp. 106-107).

Beyond emotions, the human face is instrumental in communicating nonverbal cues. Facial gestures, eye contact, and microexpressions complement verbal messages and provide additional information about the speaker's intentions, attitudes, and sincerity. For instance, eye contact can convey interest, attentiveness, and trustworthiness, while raised eyebrows may indicate surprise or disbelief. Nonverbal cues from the face enrich the communication process, allowing for a more comprehensive and nuanced understanding of the message. Moreover, the face serves as a vehicle for expressing personal identity and establishing social connections. Our facial features, including distinguishing characteristics and expressions, contribute to our individuality and uniqueness. The face is integral to presenting ourselves to others, shaping their perceptions, judgments, and initial impressions. Recognizing familiar faces allows us to establish and maintain social relationships, fostering a sense of belonging and connection. Facial appearance is essential to our self-image and influences how others perceive and receive us in various social contexts (Belting, 2017, pp. 12-13).

The human face is also vital in fostering empathy and building interpersonal connections. When we see someone's face, we can pick up on subtle emotional cues and nonverbal signals, leading to a greater understanding of their thoughts and feelings. The ability to read facial expressions and empathize with others' emotions contributes to successful social interactions, conflict resolution, and building trust. Face-to-face communication allows immediate feedback and adjustment, creating an environment conducive to effective and meaningful exchanges. Although there is a complex and wide range of facial expressions, there is a degree of universality. Despite different facial expression responses in different cultures, the same facial structure produces roughly similar results for possible communication worldwide. Even though a New Guinea tribe member might react with different facial expressions than a New York citizen, these are similar enough for essential communication (Ekman, 1972, pp. 210-211). This primary evidence alludes to the evolutionary advantage of having a developed facial expression structure, and

its certain universality aspects could also explain why humans see faces where there are not any present (i.e. pareidolia⁷).




Figure 2b. Wardle, S. (2022). *Various examples of pareidolia*. SciNews. Retrieved January 13, 2023, from <https://doi.org/10.1073/pnas.2117413119>

Pareidolia, a fascinating phenomenon of human perception, refers to perceiving meaningful patterns or faces in random or ambiguous stimuli. It is a universal aspect of human cognition and has intrigued researchers and psychologists for many years. Understanding pareidolia and its significance in recognizing virtual faces can shed light on various aspects of human perception, cognition, and technology design. The human brain can recognize and interpret facial features, even when presented with incomplete or distorted visual stimuli. This cognitive process plays a crucial role in pareidolia, as individuals often perceive faces in everyday objects, such as clouds, tree bark, or electrical outlets, even determining each one as either male or female (*Figure 2b*). In virtual environments and technology, recognizing virtual faces becomes increasingly essential. That is, it enables avatars' facial expressions and nonverbal actions to be identified even if said avatars do not resemble a strictly human form. With this human characteristic, communication between avatars becomes more straightforward when the extra steps of nonverbal communication and emotional hints are added.

Recognizing virtual faces enables users to establish a sense of presence and engagement in virtual environments. When virtual characters display facial expressions,

⁷ Merriam-Webster. (n.d.). Pareidolia. In Merriam-Webster.com dictionary. Retrieved October 26, 2023, from <https://www.merriam-webster.com/dictionary/pareidolia>

emotions, and communicative gestures, users can establish a deeper connection and better understand the character's intentions and messages. Moreover, recognizing virtual faces enhances the user experience and facilitates effective human-computer interaction. Studies have shown that users respond more positively to virtual characters exhibiting facial expressions and emotional cues. Adding facial expressions can enable more intuitive and natural interactions with virtual characters, promoting greater engagement (Trappl, Payr, & Petta, 2003, pp. 189-211). Designers must be mindful of cultural differences and context when creating virtual faces to ensure they are universally recognizable to a degree and appropriately convey intended emotions or messages.

The remarkable ability of humans to distinguish between subtly different faces allows each individual face to become unique, serving as a visual marker of a person's identity. This inherent uniqueness facilitates easier connections and relationships with others, enabling the recognition of individuals more easily. The importance of facial identity extends beyond physical recognition in real-life interactions. In virtual and digital spaces, the face takes on additional significance. Virtual avatars often center around the face to establish personal identity. The face becomes a focal point for creating a digital persona, just like the icon is for social media, representing oneself in virtual environments, social media, and online communities. Users can express their individuality and establish a distinct digital identity by customizing an avatar's facial features. Moreover, visual identity plays a crucial role in character building, even if it may lead to preconceived notions about personality. In various contexts, such as movies, literature, or video games, characters are often associated with specific facial features and expressions. These associations can shape the audience's perception of the character's traits, emotions, and motivations (Lebowitz, & Klug, 2017, p. 81). For instance, a character depicted with a stern facial expression and wrinkles might be perceived as serious or wise. In contrast, a character with a broad smile and youthful appearance might be seen as cheerful and approachable.

The power of visual identity is not limited to actual human faces. In a phenomenon known as *anthropomorphism*, similar to pareidolia, humans ascribe human-like characteristics and personalities to non-human entities, including inanimate objects, animals, or even virtual characters.⁸ When presented with images or illustrations that depict objects with facial attributes or emotional expressions, individuals tend to anthropomorphize these entities, projecting human-like qualities onto them. Personas created through

⁸ Merriam-Webster. (n.d.). Anthropomorphism. In Merriam-Webster.com dictionary. Retrieved October 26, 2023, from <https://www.merriam-webster.com/dictionary/anthropomorphism>

anthropomorphism mimic the human traits previously mentioned, such as emotions, intentions, and social behaviors. This tendency to personify non-human entities based on facial cues further underscores the importance of the human face in communication and cognition. Taking advantage of these pareidolia preconceptions makes creating avatars with somewhat predefined personalities possible by modifying their appearance. For example, a small and slender avatar will likely be associated with fragility; in the opposite case, a tall and muscular avatar will be seen as a solid, full-bodied, and strong character. These characteristics, so-called archetypes, help define a character's personality traits with their visual appearance (Lebowitz, & Klug, 2017, pp. 81-85). However, relying too much on archetypal features to construct a personality runs a risk of increasing predictability, which, in the context of virtual avatars, might either be a negative or a positive aspect.

Origins of facial motion capture

For an avatar's facial expressions to be the most human-like, it must take advantage of human microexpressions' inherent randomness. Therefore, facial motion capture needs to track the user's face. This technology converts the movements of a person's face into a digital database using cameras or laser scanners. Because the actions of real people influence the data in the animation, its final result will usually look more nuanced and realistic compared to an animation created manually from scratch because it is replicating the intricacies of human microexpressions.




Figure 3a. Robinson, P., & Baltrušaitis, T. (2012). *Examples of tracked faces*. University of Cambridge. Retrieved January 15, 2023, from <https://www.cl.cam.ac.uk/research/rainbow/emotions/facetrack>

This technology can be traced back to the early days of computer vision and image processing. Tracking and recognizing human faces in digital images and videos has been a significant area of research and development for several decades, driven by the desire to create more interactive and immersive experiences in various fields, such as computer graphics, human-computer interaction, surveillance, and augmented reality. The earliest attempts at face-tracking were found in the 1960s and 1970s, with researchers exploring techniques to identify and locate facial features manually (Nilsson, 2009, p. 172-173). However, these early methods were limited in their accuracy and efficiency due to the past's computational constraints and the task's complexity.

Lance Williams published one of the first papers discussing performance-driven animation in 1990.⁹ He describes “a means of acquiring the expressions of real faces and applying them to computer-generated faces” (Williams, 1990, p. 237). In this paper, the method was to record the data and then play it back in a digital medium such as computer graphics. However, this process must be real-time in virtual reality avatar expressiveness, which was impossible then. It was not until the late 1990s that significant progress was made in the field of face-tracking with one of the key milestones during this period being the

⁹ Williams, L. (1990). Performance-driven facial animation. *Proceedings of the 17th Annual Conference on Computer Graphics and Interactive Techniques*. <https://doi.org/10.1145/97879.97906>

Active Appearance Model (AAM) developed by Cootes, Edwards, and Taylor in 1998.¹⁰ AAM is a statistical model that represents the appearance of a face and its variations in shape and texture. It allows for tracking facial landmarks and estimating facial expressions, providing a more robust and flexible approach to face-tracking (Edwards, Taylor, & Cootes, 1998, p. 6). Another notable advancement in face-tracking came with the emergence of the Viola-Jones algorithm in 2001.¹¹ Viola and Jones introduced a real-time face-detection framework with a machine-learning approach called the *AdaBoost* algorithm. This algorithm trained a cascade of simple classifiers to detect faces in an image by efficiently evaluating many potential features (Viola & Jones, 2001, pp. 7-8). The Viola-Jones algorithm revolutionized face detection and became a cornerstone for many subsequent face-tracking systems. In recent years, the rapid progress in computer vision and machine learning, along with the availability of powerful hardware, has led to significant advancements in face-tracking technology. Deep learning techniques, particularly *Convolutional Neural Networks*¹² (CNNs), have played a crucial role in improving the accuracy and robustness of face detection and recognition systems. CNNs can automatically learn hierarchical representations of facial features, enabling more effective face-tracking in various conditions and scenarios.

The widespread adoption of face-tracking technology can be attributed to its applications in numerous fields. Face-tracking has been extensively used in the entertainment industry for facial animation in movies, video games, and, as I will explore in this dissertation, virtual reality experiences, allowing for realistic and expressive virtual characters. Moreover, face-tracking has become a fundamental component in developing *Augmented Reality*¹³ (AR) applications. AR systems can overlay virtual objects onto a user's face by accurately tracking and mapping facial features in real-time, creating interactive and immersive experiences.

¹⁰ Edwards, G. J., Taylor, C. J., & Cootes, T. F. (1998). Interpreting face images using active appearance models. Proceedings Third IEEE International Conference on Automatic Face and Gesture Recognition. <https://doi.org/10.1109/afgr.1998.670965>

¹¹ Viola, P., & Jones, M. (2001). Rapid object detection using a boosted cascade of Simple features. Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001. <https://doi.org/10.1109/cvpr.2001.990517>

¹² Awati, R. (2023, April 24). What are convolutional neural networks?: Definition from TechTarget. Enterprise AI. <https://www.techtarget.com/searchenterpriseai/definition/convolutional-neural-network>

¹³ A website explaining the definition of AR and its difference between other types of extra realities. Microsoft Corporation. (2022). What is augmented reality (AR): Microsoft Dynamics 365. What is Augmented Reality (AR) | Microsoft Dynamics 365. <https://dynamics.microsoft.com/en-us/mixed-reality/guides/what-is-augmented-reality-ar/#:~:text=What%20is%20AR%3F,sensory%20stimuli%20via%20holographic%20technology>

Nowadays, there are two main ways to track the human face and transmit the data for computer graphics: *Marker-based*¹⁴ and *Markerless*¹⁵ or *Image-based* tracking. The former uses a camera system that looks at and tracks physical trackers attached to the user's face. It was more widely used in movie production but has since fallen out of use because of the discomfort of having physical points attached to the face. The latter method only uses cameras/laser systems to capture the face. This system is the one that received the most development as it currently is an iteration of the previously mentioned developments. By looking at contrast points, such as the nostrils/lip/eye corners, it uses a neural network and artificial intelligence to detect which muscles are in use. As a result of its ease of use, this method is the most widely used one for Virtual Reality as it provides “good enough” fidelity and does not cause any kind of physical discomfort as it does not need to make contact with the face.




Figure 3b. Weta Digital. (2017). *Marker-based face-tracking example*. FXGuide. Retrieved January 15, 2023, from <https://www.fxguide.com/xf/featured/weta-digital-facets-honoured-at-the-sci-tech-awards/>

&

Figure 3c. Dynamixyz Team. (2020). *Markerless face-tracking example*. 80LV. Retrieved January 15, 2023, from <https://80.lv/articles/dynamixyz-insights-into-facial-motion-capture-technology/>

VR face-tracking

It is possible to apply the previous techniques to avatars in a virtual environment, which has been done before in Virtual Reality experiences such as *NeosVR*¹⁶, *VRChat*, and, more recently, *Horizon Worlds*¹⁷. In these experiences, the system detects users' facial expressions and converts those into data, which, in turn, is transmitted to an avatar. In the

¹⁴ Weta Digital. (2017). Marker-based face-tracking example. FXGuide. Retrieved January 15, 2023, from <https://www.fxguide.com/xf/featured/weta-digital-facets-honoured-at-the-sci-tech-awards/>

¹⁵ Dynamixyz Team. (2020). Markerless face-tracking example. 80LV. Retrieved January 15, 2023, from <https://80.lv/articles/dynamixyz-insights-into-facial-motion-capture-technology/>

¹⁶ Mariancik, T. (2018). *NeosVR* (PC version) [Video game]. Solirax. <https://neos.com>.

¹⁷ Zuckerberg, M. (2021). *Horizon Worlds* (Quest version) [Video game]. Meta Platforms, Inc. <https://www.oculus.com/horizon-worlds>.

end, if the avatar is created with face-tracking in mind, it will look like it mimics the user's expressions and feelings.

As mentioned before, virtual reality completely immerses the user in its new, fabricated reality, a phenomenon called “immersion” or “presence” (Barfield & Furness, 1995, pp. 9-10), and derived initially from *telepresence*.¹⁸ This phenomenon can be increased by creating a virtual reality experience miming various real-life characteristics. Examples of such are engaging environments that encompass the entire perspective of the user and make physical sense, and, more importantly for this project, user-avatar interaction, one that makes the user feel it is embodying the avatar and taking its identity. Face-tracking technology is a great way to increase the feeling of presence because it increases the avatar’s response to the user’s actions, increasing user-avatar interaction. For example, by looking at a virtual mirror, the user will see the avatar following every expression, every smile and frown, and where it is looking. Behavior that can, in turn, increase the connection between the user and the avatar.



Figure 4a. HTC Corporation. (2021). *Example of a virtual reality headset with a face-tracker add-on*. Vive. Retrieved January 16, 2023, from <https://www.vive.com/eu/accessory/facial-tracker>

Due to Virtual Reality headsets occupying the face and obstructing the view from the outside, VR face-tracking has to be divided into two parts: eye-tracking inside the face gasket and lower face-tracking below the headset. The eyes can be associated with the “lower-hanging fruit” of facial expression.¹⁹ In other words, they are easier to implement

¹⁸ Mentioned by Barfield and Furness and originated in: Bejczy, A. K. (1980). Sensors, controls, and man-machine interface for Advanced teleoperation. *Science*, 208(4450), 1327–1335. <https://doi.org/10.1126/science.208.4450.1327>

¹⁹ VRChat developer discussing the relative simplicity of implementing eye-tracking over face-tracking. Tupper. (2022, November 18). Developer update - 17 November 2022. VRChat Ask. <https://ask.vrchat.com/t/developer-update-17-november-2022/14986/15>

because of their straightforward muscle structure (consisting of look direction, eyelid openness, and squinting). However, with only that information, it is possible to communicate intent and attention, where the user is looking at and where it is focused on, two of the most critical aspects of human expression. The lower face, in contrast, is much more challenging to implement because of the amount and complexity of muscles (jaw movement, lip movement, tongue movement, and others) and all the different face-tracking standards that exist, a topic that I will approach in the next chapter. Despite that, the lower face is still essential and is responsible for speech, reaction, and the bulk appearance of facial expressions. With it, more complex communication is possible.

Some higher-end headsets are sold with face-tracking built-in (e.g., *Meta Quest Pro*²⁰ or the upcoming *Apple Vision Pro*²¹) or require an extra eye-tracking and/or lower face-tracking module to give it the ability to track the individual's face (e.g., *Vive Pro Eye*²² with the *Vive Facial Tracker* or the upcoming *Vive XR Elite*²³ with the optional eye-tracking module and lower face-tracking module). These virtual reality headsets with face-tracking technology are often released with a relatively high price point, making them less accessible to the general public. This is primarily due to the advanced hardware and software components required to enable accurate and responsive face-tracking capabilities. Incorporating specialized sensors, cameras, depth sensors, and the necessary computational power contributes to the overall cost of these headsets.

However, the increasing number of releases and popularity in the market suggests a promising future for the accessibility and affordability of virtual reality headsets with face-tracking. As the technology continues to evolve and gain traction, manufacturers will likely invest in research and development, leading to advancements in hardware miniaturization, cost reduction, and increased efficiency, as seen with the smartphone industry.²⁴ As with any emerging technology, the initial high price can be attributed to limited production volume, research costs, and the need to recoup investments. Over time, as demand grows and production scales up, economies of scale come into play, resulting in reduced

²⁰ Meta Platforms Inc. (2023, October 25). *Meta Quest Pro product page*. Meta. <https://www.meta.com/quest/quest-pro/>

²¹ Apple Inc. (2023). *Apple Vision pro product page*. Apple. <https://www.apple.com/apple-vision-pro/>

²² HTC Corporation. (2019). *Vive Pro Eye product page*. Vive. <https://www.vive.com/au/product/vive-pro-eye/overview/>

²³ HTC Corporation. (2023). *Vive XR Elite product page*. Vive. <https://www.vive.com/us/product/vive-xr-elite/overview/>

²⁴ Blog that displays various statistics on smartphone usage and popularity. Turner, A. (2023, October 1). 3.86 billion more phones than people in the world!. BankMyCell. <https://www.bankmycell.com/blog/how-many-phones-are-in-the-world>

manufacturing costs. Additionally, advancements in manufacturing processes and component integration are expected to drive down the overall price of these headsets. Furthermore, as competition intensifies, companies are likely to vie for a larger market share by offering more affordable options with face-tracking capabilities. This healthy competition can drive innovation, efficiency, and cost reduction, making virtual reality headsets with face-tracking increasingly accessible to a wider range of consumers. Moreover, the demand for face-tracking features will likely increase as the technology becomes more widespread and integrated into mainstream applications.

Video games and Standards

3D-modeled characters have technologically evolved since the early days of computer graphics. As such, they have become much more dynamic, moveable, and easier to work with. One of these technologies (so-called *Morph target animation*, *per-vertex animation*, *shape interpolation*, *shape keys*, or *blendshapes*) is currently the backbone behind the face-tracking ability of 3D computerized characters (Ravikumar, 2017, pp. 17-18).




Figure 5a. Ravikumar, S. (2017). *Examples of different blendshapes applied to a singular model*. Research Gate. Retrieved January 17, 2023, from https://www.researchgate.net/figure/Blendshapes-from-our-Blendshape-model-for-facial-animation_fig4_329019919

By deforming the model, the system records that deformed shape and stores it along with its information. This deformation is applied afterward by an animation that can choose any state between its original form and its new, deformed shape. For example, a character's eyelid can have a blendshape that reshapes it to a closed state. Afterward, an animation can access that information and smoothly transition from a character with its eyes opened to the same character with one eyelid closed. Since the face has multiple muscles that can move independently, it is necessary to create numerous blendshapes that also move independently. This way, each blendshape can be animated separately, which, in turn, simulates accurate

face-tracking (Ravikumar, 2017, p. 18). This system allows the avatar to follow the user's face smoothly. For example, if the user slowly blinks one eye, the avatar can also slowly blink the same eye. If the user moves their jaw down and then to the left, the avatar will most likely move its jaw down using one blendshape and then activate another blendshape on top of the existing one to move the jaw to the left, resulting in a combination of both blendshapes, mimicking the user's jaw pose.

However, if every single muscle is copied using a blendshape, the final model will be unoptimized and hard to compute in real-time due to the high number of muscles/blendshapes. To get around this issue, some muscles have to be grouped, muscles that would naturally move together on their own; therefore, different standards have been created that categorize different muscle groups. The most popular of these standards are Apple's *ARKit*, utilized in existing virtual reality applications, Apple's *Memoji*²⁵ system, and, more recently, in some streaming applications like *VTubing*²⁶, and Vive's *SRanipal*²⁷, utilized only in some virtual reality applications.

Human replication

Avatars with faces, in the context of this dissertation, can be categorized into two due to the main feature they are trying to achieve: human replication or character identity. The former tries to mimic the user's physical appearance, and the latter tries to be distinctive and separate from the user. In this chapter, I will go into more detail about avatars whose primary purpose is to replicate a specific human face:

²⁵ Apple Inc. (2022, September 9). *Memoji*. Apple Support. Retrieved January 17, 2023, from <https://support.apple.com/en-us/HT208986#:~:text=How%20to%20create%20your%20Memoji,and%20tap%20the%20New%20Memoji>

²⁶ Saito, J. (2022, June 25). *Vtubing for Beginners*. Streamlabs. Retrieved January 17, 2023, from <https://streamlabs.com/content-hub/post/vtubing-for-beginners>

²⁷ *SRanipal Getting Started Steps - Developer Resources*. (2019). <https://developer.vive.com/resources/vive-sense/eye-and-facial-tracking-sdk/tutorials/sranipal-getting-started-steps/?site=kr>




Figure 6a. Nichols, G. (2019). *ObEN co-founder Adam Zheng and his avatar*. ZDNet. Retrieved January 16, 2023, from <https://www.zdnet.com/article/wanted-realistic-avatars-for-virtual-reality-meetings>

In the realm of computer graphics and digital representation, *photogrammetry* has emerged as a groundbreaking technique that revolutionizes the process of scanning reality and reconstructing into the virtual,²⁸ which can also be used for scanning people's faces and reconstructing those scans into realistic avatars.²⁹ This innovative technology opens up new frontiers in virtual environments, entertainment industries, human-computer interaction, and, most importantly, allows for highly detailed and lifelike digital representations of individuals. At its core is the process of extracting three-dimensional information from two-dimensional photographs. It combines computer vision, image processing, and mathematical algorithms to reconstruct accurate and detailed 3D models from a collection of photographs. Traditionally, photogrammetry has been employed in various fields, such as surveying, architecture, and archaeology, to capture and analyze real-world objects and landscapes. However, recent technological advancements have expanded its applications, creating highly realistic digital representations of human faces.³⁰ This process of scanning people's faces into realistic avatars begins with the capture of multiple high-resolution photographs from various angles. These photographs serve as the input for the photogrammetry software, which analyzes the visual data and reconstructs a 3D model of the face. The software matches common features and

²⁸ Various definitions of photogrammetry within different fields. ScienceDirect. (2021). Photogrammetry. Photogrammetry - an overview | ScienceDirect Topics. <https://www.sciencedirect.com/topics/agricultural-and-biological-sciences/photogrammetry>

²⁹ Nichols, G. (2019). *ObEN co-founder Adam Zheng and his avatar*. ZDNet. Retrieved January 16, 2023, from <https://www.zdnet.com/article/wanted-realistic-avatars-for-virtual-reality-meetings>

³⁰ Summers, N. (2019, July 19). *The real-time motion capture behind "hellblade."* Engadget. <https://www.engadget.com/2017-08-08-ninja-theory-hellblade-motion-capture-demo-video.html>

points across the photographs to calculate the depth and geometry of the face, creating a highly accurate representation of the subject's facial structure. Additionally, texture mapping techniques are employed to apply the captured photographs to the 3D model, ensuring that the avatar retains the realistic appearance of the individual. The advancements in hardware and software technologies have played a vital role in enabling photogrammetry for facial scanning. High-resolution cameras, depth sensors, and specialized scanning equipment capture intricate facial details precisely. Moreover, powerful computational algorithms and software packages have made the reconstruction process more efficient and accessible to more users.

The applications of photogrammetric avatars are diverse and far-reaching. In the realm of virtual reality, these realistic avatars enhance the sense of presence and immersion, enabling users to embody digital representations that closely resemble their physical selves. This technology has also found extensive use in the entertainment industry, where actors' faces can be scanned and used to create digital characters for films, video games, and animated content, usually called *performance capture*³¹. Furthermore, photogrammetric avatars have the potential to revolutionize teleconferencing and social interactions, allowing for more engaging and expressive communication in digital spaces.

In addition to the high level of realism achieved through photogrammetry, the expressive abilities of digital avatars can be further enhanced using the methods mentioned before. Mapping the user's facial movements onto the avatar's face can exhibit a wide range of emotions and convey subtle nuances, adding a layer of authenticity and immersion to virtual interactions beyond the already realistic model can provide by itself. As discussed before in the face-tracking chapter, this mapping process allows the avatar to mirror the user's expressions and gestures in real-time, creating a more expressive and responsive digital representation. The combination of photogrammetry and face-tracking technology introduces a new dimension to the communication potential of digital avatars, enabling users to convey emotions and nonverbal cues in virtual environments with remarkable realism.

However, it is fairly easy for this replication process to go wrong. Since the individual has grown up with their appearance and is very familiar with what it looks like, they will likely take the similarities between their appearance and the avatar's appearance for granted, which makes all of the subtle differences more pronounced. In other words, the individual

³¹ A video demonstrating the performance capture technology utilized for the video-game Death Stranding. Kojima Productions. (2019, November 15). Death stranding – performance capture – 4K. YouTube. <https://www.youtube.com/watch?v=-tYEhXxVEwo>

will focus on all the disparities between their perceived face and the avatar's face, triggering a sense of unease whenever these disparities are too significant. This unfortunate side-effect makes avatars with the sole goal of replicating reality quickly fall into the *Uncanny Valley*³². It is essential to acknowledge that despite the advancements in creating realistic avatars through photogrammetry, there is a potential challenge associated with the uncanny valley phenomenon. The uncanny valley refers to the discomfort or unease experienced by individuals when encountering human-like entities that appear almost but not entirely human. This phenomenon arises when the level of realism in a digital representation is not convincing enough to the point where it becomes indistinguishable from reality (Mori, MacDorman, & Kageki, 2012, pp. 98-100).

The uncanny valley effect becomes particularly pronounced when attempting to replicate human faces since individuals are highly accustomed to the appearance of their own faces and the subtle nuances of human facial expressions. Our brains are finely tuned to detect the slightest discrepancies in facial features, proportions, and movements. As a result, if the photogrammetric avatars fall short of accurately capturing these nuances, users may experience a sense of eeriness or discomfort when interacting with their digital counterparts. Researchers and developers continuously strive to improve photogrammetric avatars' fidelity and realism to overcome this challenge. They focus on capturing intricate details such as skin texture, facial hair, and microexpressions to create more convincing representations. The user's familiarity with the appearance of their own face and general human facial features adds an additional layer of complexity.




Figure 6b. Reynolds, E. (2015). *Example graph of the uncanny valley effect*. Wired. Condé Nast Britain. Retrieved January 17, 2023, from <https://www.wired.co.uk/article/uncanny-valley-creepy-robot>

³² A blog explaining the definition of Uncanny Valley. TechTarget. (2016, February 10). What is Uncanny Valley?: Definition from TechTarget. WhatIs.com. <https://www.techtarget.com/whatis/definition/uncanny-valley>

In aesthetics, the uncanny valley is a hypothesized relationship between an object's degree of resemblance to a human being and the emotional response to the object. The concept suggests that humanoid objects that imperfectly resemble actual human beings provoke uncanny or strangely familiar feelings of uneasiness and repulsion in observers. "Valley" denotes a dip in the human observer's affinity for the replica, a relation that otherwise increases with the replica's human likeness. This phenomenon also applies to VR, so for it to be avoided, avatars must either perfectly replicate the human appearance or create their own unique appearance and character, with the aim of the user not immediately identifying with them, but instead identifying them as something other than strictly human.³³

Avatar self-identity

Avatars can also acquire their self-identity, separate from the user. This, in turn, makes them more resistant to the uncanny valley effect since their appearance is so different that the user will not start by subconsciously looking for differences between the two, possibly looking for similarities instead. The same general guidelines of character design apply to the creation of avatars with their own distinct identity (Seegmiller, 2003, pp. 3-12):

- Achieving an unmistakable look and feel. It involves carefully considering various visual elements and their deliberate arrangement to convey essential information about the character while capturing the user's attention. From the character's physical appearance, such as their facial features, body proportions, and distinctive attire, to the use of color palettes, textures, and visual motifs, each element plays a crucial role in conveying the character's personality, background, and role within a fictional narrative. Using these visual cues, character designers can communicate key aspects of the character's identity, motivations, and emotions, providing a glimpse into their world and establishing a strong visual connection. Through the harmonious integration of these elements, the character's presence becomes visually striking and capable of evoking emotions, creating intrigue, and leaving a lasting impression;
- Crafting a unique and striking appearance. A character with a distinct visual identity stands out amidst a sea of familiar faces, leaving an indelible mark in the user's mind. By carefully considering these elements, character designers can create a visually arresting persona that instantly grabs attention and sparks curiosity. A unique

³³ Sparks, M. (2019, December 26). Metafocus: Avoiding the Uncanny Valley in VR & serious games : Learning solutions. The Learning Guild.
<https://www.learningguild.com/articles/metafocus-avoiding-the-uncanny-valley-in-vr-serious-games/>

appearance serves as a visual hook and conveys important aspects of the character's personality, backstory, and role within the narrative. Whether it be through unconventional shapes, intricate details, bold color combinations, or unexpected combinations of visual elements, a character with a distinctive appearance becomes a memorable icon, capable of resonating with the audience and leaving a lasting impact;

- Maintaining consistency between its personality traits and its visual traits. When designing a character, it is essential to consider how their visual elements align with their personality and fictional values. The visual traits should reflect and reinforce the character's core attributes, whether bold, energetic, mysterious, brooding, gentle, or compassionate. This consistency helps to create a cohesive and believable character, allowing the user to connect with and understand their motivations and behavior easily. For example, a character with a fierce and confident personality may exhibit sharp, angular features and bold, vibrant colors in their design, visually conveying their strong-willed nature. On the other hand, a character with a calm and nurturing disposition may have soft curves, warm colors, and gentle expressions, visually representing their compassionate nature. By carefully aligning the character's visual traits with their personality, character designers can enhance the audience's immersion in the story and strengthen the overall impact of the character's portrayal.



Figure 7a. Santos, D. (2022). *Example of Zuri - an Avatar I created with its own identity*. Gumroad. Retrieved January 16, 2023, from <https://gh0stt.gumroad.com/l/zuri>

Because of the visual disparity between the user and the avatar, the initial feeling of presence will arrive slowly compared to a hyper-realistic visual copy of the user. However, given enough time and interaction, it can come later as the user acclimates to their embodiment of the avatar. Surprisingly, the user can associate parts of their identity with the avatar, creating a more intense feeling of presence. This phenomenon usually grows with time as the user gets more and more used to the avatar. In this state, the user will start to roleplay as the character, potentially changing their behavior to match the character's:

- *“Because virtual reality is entirely programmed, the form or type of virtual body can be quite different from the participant's actual body, which can impact perception, attitudes and behavior. A dramatic example of this is when adults inhabiting a virtual child's body overestimate the size of objects and demonstrate implicit attitude and behavioral changes that seem more child-like.”* (Slater & Sanchez-Vives, 2014, p. 24).

This divergence between the participant's real and virtual bodies can significantly affect perception, attitudes, and behavior. This suggests that the virtual embodiment in a different body can influence cognitive processes, leading to altered perceptions, attitudes, and behaviors that align with the characteristics associated with the virtual body. Such findings shed light on the malleability of our self-perception and the potential of virtual reality to shape our psychological experiences and social interactions. It is a medium that lets someone try on and play with a different identity. The boundaries of their physical self blur, and they assume a new identity, free from the constraints of their everyday existence. This liberation allows individuals to explore and discover different aspects of themselves, adopt roles they may not typically embrace, and immerse themselves fully in the virtual experience.

Moreover, with a sufficient presence factor, virtual reality becomes a medium for people to express themselves. In certain instances, it can even facilitate a more natural exploration of one's gender identity, a significant and prevalent topic within *VRChat*. This exploration is exacerbated by the considerable presence of the queer community within the game, further highlighting its importance and impact.³⁴ However, for this sense of transformation and embodiment to reach its pinnacle, there needs to be a degree of correspondence between the user's facial expressions and body movements and those of the avatar (Slater & Sanchez-Vives, 2014, pp. 24-28).

³⁴ Straszfilms. (2021, February 3). Identity, gender, and Vrchat (why is everyone in VR an anime girl?). YouTube. https://www.youtube.com/watch?v=5v_DI7i4Bew



Figure 7b. Santos, D. (2023). *Face-Tracking Showcase*. YouTube. Retrieved July 10, 2023, from <https://www.youtube.com/watch?v=3xJ0opUrHZY>

To comprehensively demonstrate the concept of avatar identity and the dynamic nature of face-tracking, I have produced a video showcasing a diverse array of avatars with the functionality (*Figure 7b*). This visual presentation serves as a compelling illustration of the capabilities inherent in avatar technology, explicitly highlighting the ability of these virtual entities to mirror and emote facial expressions in direct response to the user's own. Of note, the video presented was created with the intent of showcasing different visually-styled avatars within the face-tracking context, demonstrating how each one, besides looking different, can also relay the user's facial expressions differently. It offers viewers a captivating glimpse into the seamless integration of user interaction and avatar embodiment and how different each avatar's final look can be from each other.

Full-body tracking

Full-body tracking (FBT) in virtual reality, similar to face-tracking, aims to enhance the immersive experience by capturing and translating a user's movements onto an avatar, except it tracks the rest of the user's body into the virtual environment instead of just the face. Unlike traditional VR setups that primarily track the movement of the head and hands, full-body tracking utilizes additional sensors or devices to monitor the user's movements across various body parts, such as the legs, torso, and feet, usually referred to as *trackers*³⁵.

³⁵ An example of SteamVR-based trackers created by TundraLabs using the lighthouse system, a system developed by Valve Corporation using precise laser timing information to triangulate a point in space. Tundra Tracker Bundle. Tundra Labs. (2021). <https://tundra-labs.com/products/tundra-tracker-1>

By incorporating FBT into a VR experience, virtual reality becomes a more encompassing and embodied experience, allowing users to see their whole body and observe its movements mirrored in their virtual avatar, behaving similarly to *motion capture*³⁶.

This level of immersion goes beyond mere hand gestures and head tilts, enabling users to engage in a more natural and intuitive manner within the virtual world. Users can see their body movements replicated in real-time within the virtual environment, creating a powerful sense of presence. This heightened presence amplifies the feeling of "being there," as users feel a stronger connection between their physical actions and their virtual representation (Slater & Sanchez-Vives, 2014, pp. 24-25). FBT allows users to walk, run, jump, crouch, and perform various physical movements accurately translated to their avatars. Moreover, FBT opens up opportunities for more interactive and engaging social experiences within virtual reality. Users can engage in activities such as dancing, sports, or collaborative tasks, where their movements and gestures contribute to the shared experience. In the previously shown video, FBT can also be seen contributing to the level of expression of the avatars.

In the virtual world, communication goes beyond words and facial expressions. Just as in the physical world, body language plays a crucial role in conveying meaning and fostering connections. However, in virtual reality, users can explore an expanded repertoire of body language that goes beyond what is possible in traditional communication mediums. Full-body tracking technology allows users to communicate nonverbally using gestures, postures, and movements. They can wave hello, stand in a slouched pose, indicating comfort and ease, or in a straightened-up pose, indicating attentiveness, directness, and possible discord. These actions allow users to express emotions, intentions, and reactions, creating a rich and dynamic social interaction within virtual environments. By embracing the subtleties of body language, users can form deeper connections and engage in multifaceted social interactions, transcending the limitations of verbal communication alone.

³⁶ “A technology for digitally recording specific movements of a person (such as an actor) and translating them into computer-animated images”. Merriam-Webster. (n.d.). Motion capture. In Merriam-Webster.com dictionary. Retrieved October 26, 2023, from <https://www.merriam-webster.com/dictionary/motion%20capture>



Figure 8a. Thrillseeker. (2021). Cheap Full Body Tracking for EVERYONE : SlimeVR Review. Youtube. Retrieved July 3, 2023, from <https://www.youtube.com/watch?v=O13L5MxOIHo>

While FBT has immense potential, it is essential to note that its current implementation does not perfectly replicate the body's movements. Nevertheless, from personal experience, the tracking accuracy is more than good enough to increase the level of immersion the user feels and, more importantly, to communicate non-verbally in the medium. As the technology evolves, there will possibly be advancements in sensor technology, such as improved accuracy and comfort, better pose AI estimation and affordability, making FBT more accessible to a broader audience.

Presence and “Phantom Touch”

In this medium, the feeling of presence can become so strong as to affect the user's physical perception of their environment and their own sense of identity, the action happening around them, as well as create fake physical sensations. When users enter a virtual environment, they strive to be fully present in that digital space, as if they are physically there. This feeling of presence, mentioned before, is established through a combination of sensory inputs, visual and auditory cues, and the user's cognitive engagement. As users navigate and interact within virtual worlds, the sense of presence intensifies, enabling them to forge meaningful connections, engage in social interactions, and even experience sensations like “phantom touch”. This recent phenomenon refers to the perceptual experience of feeling tactile sensations in a virtual reality environment, even though no physical contact or stimulation is actually occurring (Pilacinski, Metzler, & Klaes, 2023). In various social VR

platforms, users can interact with each other using virtual avatars and engage in various activities, including virtual handshakes, hugs, or even more intimate gestures. Phantom touch arises from the brain's ability to integrate visual and auditory cues with expectations and past experiences. When users engage in social interactions within VR, their brains may generate sensations that mimic the feeling of physical touch. This phenomenon occurs due to a combination of sensory cues the VR system provides, the virtual environment's immersive nature, and the individual's cognitive processes and psychological immersion.

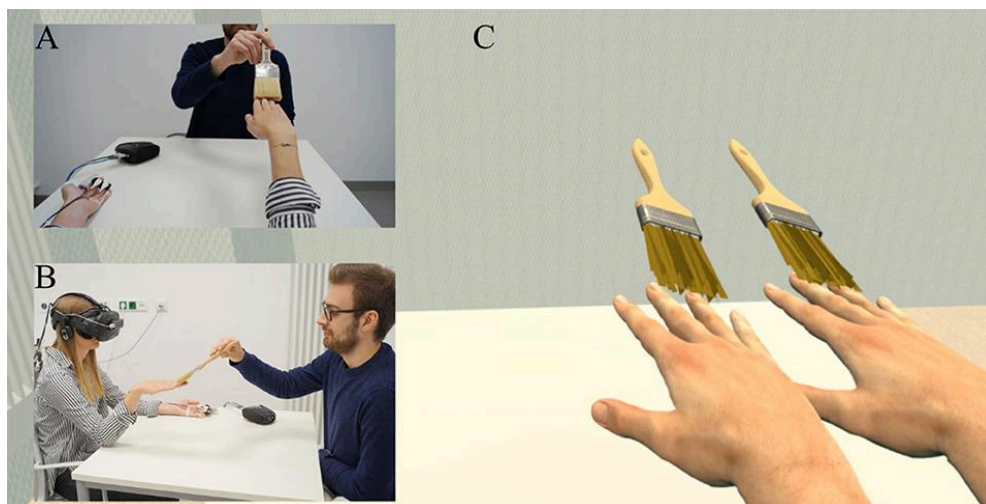


Figure 9a. Braun, N., Berisha, A., Anders, D., Kannen, K., Lux, S., & Philipsen, A. (2020). Experimental inducibility of supernumerary phantom limbs: A series of virtual reality experiments. *Frontiers in Virtual Reality*, 1. <https://doi.org/10.3389/frvir.2020.00012>

Although the term “phantom touch” is relatively recent, it takes inspiration from the previously known phenomenon “phantom limb sensation”, referring to the perception of sensations or movement in a limb or body part that is no longer present. It is commonly experienced by individuals who have undergone amputation or have lost a limb due to injury or illness. Despite the physical absence of the limb, the person may still feel sensations such as tingling, itching, pain, or even the ability to move the phantom limb. Phantom sensations are believed to occur due to the brain's ability to generate and maintain a representation of the body, known as the *body schema*³⁷. This neurological map lets us sense our body's position, movements, and sensations. When a limb is amputated, the brain may continue generating signals corresponding to the missing limb, leading to the perception of phantom sensations (Poor Zamany Nejatkermany et al., 2016).

³⁷ Mendoza, J. E. (1970, January 1). Body schema. SpringerLink. https://link.springer.com/referenceworkentry/10.1007/978-0-387-79948-3_713

Phantom touch, like phantom limb sensation, can vary in intensity and duration among individuals. Some people may experience occasional fleeting sensations, while others may have persistent and more pronounced sensations (Pilacinski, Metzler, & Klaes, 2023). The exact mechanisms behind phantom touch are still not fully understood. Still, it is believed to involve a complex interplay between the brain, the nervous system, and the individual's perception and interpretation of sensory information. This fascinating phenomenon highlights the intricate relationship between the body and the mind. It underscores the brain's remarkable ability to create and maintain a coherent perception of our bodies, even without physical sensory input. Users often describe experiencing phantom touch within virtual reality as feeling a sense of warmth, pressure, or the sensation of someone's presence without any physical contact. These sensations can create a heightened sense of presence and social connection within the virtual space, blurring the boundary between physical and digital realities. It adds another layer of realism and intimacy to the virtual interactions, contributing to a more engaging and immersive social experience. The concept of phantom touch underscores the transformative power of virtual reality in simulating sensory experiences and expanding the possibilities for social interaction in virtual environments (Alexdottir & Yang, 2022).

However, users can also embody avatars with physical attributes different from their own. This allows for exploring unique sensory experiences that extend beyond the boundaries of their physical bodies. For example, consider a user who adopts the avatar of an anthropomorphized rabbit, complete with long, expressive rabbit ears. When interacting with the virtual environment, including interactions with others, the user can vividly imagine the sensation of touch in the area of the rabbit ears; an example put forth from various conversations and from personal experiences inside *VRChat*. The fascinating imaginative phantom touch experience is made possible by the brain's capacity to integrate visual and auditory cues, past sensory memories, and learned associations. The brain constructs a perceptual representation of the virtual body, incorporating the presence of the rabbit ears into its body schema. As a result, when the virtual rabbit ears are touched or interacted with, the brain can generate the sensation of touch in that specific area, even though the user does not physically possess rabbit ears in real life.

This phenomenon exemplifies the extraordinary adaptability of the human brain and its ability to immerse us in alternative realities. By allowing users to inhabit avatars that deviate from their physical selves, VR opens up a realm of imaginative possibilities and

multisensory experiences. It enables users to engage in novel forms of embodiment, allowing them to perceive and interact with virtual body parts authentically and realistically.

Avatar expression and embodiment

Despite looking different from the user, avatars can still express themselves in all kinds of ways as, in a virtual medium, there are no limits to how an avatar can look. If an avatar is inexpressive, it will not react to the user's facial input, creating an abstract barrier between both. In this case, the user might not form as big of a connection with it as they would otherwise. However, with face-tracking, the avatar's expressions can follow the user's, further enhancing the relationship between both. The avatar will feel more natural compared to a non-face-tracked one.



Figure 10a. Santos, D. (2021). *Example of Belai - another Avatar I created with face-tracking capabilities*. Gumroad. Retrieved January 18, 2023, from <https://gh0stt.gumroad.com/l/belai>

Nevertheless, the avatar can take that input and exaggerate/modify its expressions, as face-tracking is just input. For example, if the avatar has animal ears, a frown from the user could subtly move down the ears along with its frown. This not only exaggerates the frowning of the avatar as it accentuates it but also adds personality and identity to the avatar itself. This behavior would, just like a self-identifying avatar, decrease presence at first because the user does not see the immediate connection between themselves and the avatar.

However, with time and interaction, the user could get acclimated to the changes, which increases the feeling of presence and, just like before, could associate parts of their identity with the avatar (Slater & Sanchez-Vives, 2014, p. 25).

Furthermore, a user has a personal preference for visual appearance. Just like one might choose different clothes from another, some might prefer different-looking avatars than others. In a real-life social setting, visual appearance generally defines identity and personality of each participant. In virtual reality, this is taken to the extreme as there is a more extensive and more accessible range of visual differences that there can be between users.

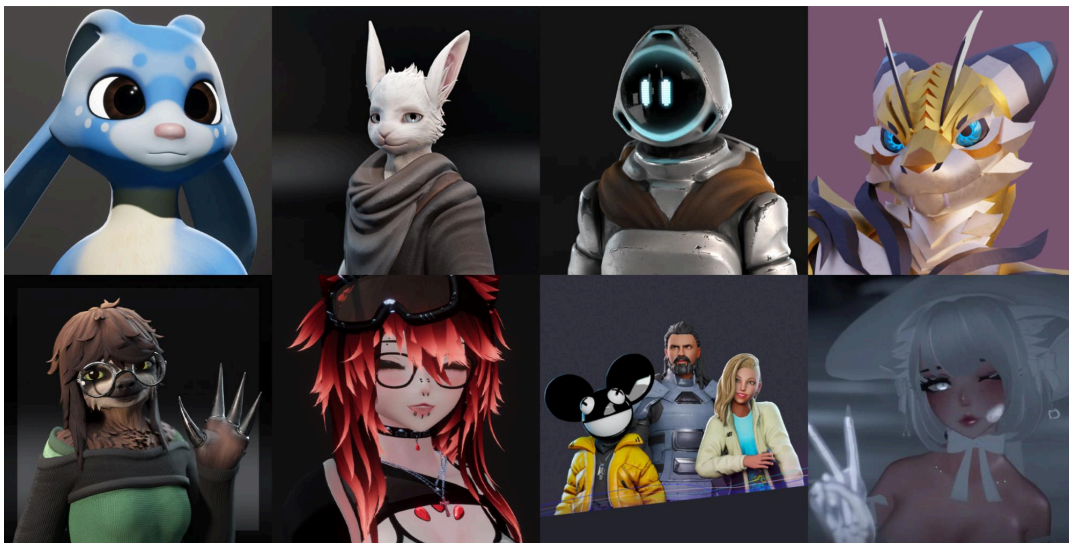


Figure 10b. Santos, D. (2023). *A collage of different avatars*. Retrieved January 18, 2023.

In order left-right, top-down:

- HIYU (2022). *Spirit VRChat Avatar*. from <https://hiyu.gumroad.com/l/Spirit>
- Santos, D. (2022). *Belai (VRChat Avatar)*. from <https://gh0stt.gumroad.com/l/belai>
- Santos, D. (2021). *Koru (VRChat Avatar)*. from <https://gh0stt.gumroad.com/l/tKwhb>
- Drakious (2022). *[VrChat] Robot Tiger Bee (PC & Quest)*. from <https://drakiousderg.gumroad.com/l/gchuj>
- Santos, D. (2022). *Zuri (VRChat Avatar)*. from <https://gh0stt.gumroad.com/l/zuri>.
- Googi (2022). *Ash 3.0 VRChat Avatar*. from <https://googii.gumroad.com/l/ash>
- Ready Player Me (2022) *Banner Image*. from <https://readyplayer.me>
- Rxsi (2021). *Hex*. from <https://rxsi.gumroad.com/l/Hex>

In conclusion, suppose an avatar is simultaneously created with face-tracking and proper body proportions in mind and liked aesthetically by the user. This user will likely feel the most presence embodying it. While varying from person to person, in this case, the user has a higher chance of merging parts of their personality with the avatar's, resulting in various behavioral changes depending on the avatar. However, this freedom of appearance, the escapism from real life through identity exploration, is not the public's perceived norm in social virtual reality, which I will explore in the next chapter.

The false mainstream

Why does the current state of virtual reality seem very corporate and why are there so few contemporary experiments with avatars? Recently, Mark Zuckerberg gave a popular press presentation regarding the future of his social media-based company, Facebook. During this presentation, he announced the company's shift in terms of goals and name. Meta, the new name presented, will now focus on developing the "Metaverse," a VR-based approach to social experiences. This announcement generated much discourse due to its significant change. The goals seem noble at first sight: to create the foundation for the future of social experiences and to expand it by setting guidelines. However, upon closer inspection, these goals and future benefit no one but Meta. The entire premise of an open future should be appraised with discernment simply because Meta is the one who wants to create it and, therefore, control it. This is the same company that previously controlled the social media landscape with Facebook and, later on, introduced monopolistic tactics of buying the competition to control the XR space.³⁸ In addition, monetization within the Metaverse has been hyped up based on NFTs, which were later discovered to be primarily used in scams.³⁹ However, despite its controversial nature, the presentation did well in marketing and bringing Meta's name into the public discourse, which was its ultimate goal.



Figure 11a. Meta. (2022, February 4). *Horizon Worlds' Avatar appearance and style*. Meta. <https://about.fb.com/news/2022/02/personal-boundary-horizon/>

³⁸ Kafka, P. (2021, November 11). Facebook is quietly buying up the metaverse. Vox. <https://www.vox.com/recode/22776461/facebook-meta-metaverse-monopoly>

³⁹ Olson, D. (2022, January 21). Line Goes Up – The Problem With NFTs. YouTube. https://www.youtube.com/watch?v=YO_xWvX1n9g

Regarding this thesis, the main issue with Meta's announcement is its current showcase of corporate, uncanny valley-looking avatars with the simple goal of mimicking a user's appearance, simplifying it to run on mobile processors. These Meta avatars never experiment with different ways of user expression and identity besides basic human facial shapes, hairstyles, clothing options, and a set of facial expressions copied and standardized between all avatars. As previously established in the Human Replication chapter, avatars made to simply replicate the human appearance risk falling into the uncanny valley. This showcase of “the next chapter of social connection” shifts the public’s perspective on what avatars are and limits what they can be, removing creativity in avatar appearance and, by extension, user identity. The presentation ends up seeming incredibly tone-deaf. Nevertheless, it increased Meta’s popularity at the cost of artificially inflating the popularity of their virtual avatars, souring the public's perception of avatars in general and possible social experiences in.

On the other hand, *VRChat*, the game used in this thesis to explore avatar expressivity, situates itself on the complete opposite of the spectrum. The established platform embraces user creativity by allowing entirely custom models to be imported into the game and thriving on this community-created content. This open system encourages users to explore different facets of their identity and experiment during social interactions. During a conversation within the game, users consciously wear avatars that represent specific aspects of themselves without necessarily looking human. This freedom of appearance allows for a greater breadth and, more importantly, a greater depth for exploring their identity – although, as in real life, avatar styles can also follow trends, like anime and furry, which are the most prevalent in *VRChat*. Like clothing choices, these trends may simply reflect a desire to visually align with others, which is also a valid aspect of identity. Nevertheless, the presence of these trends does not preclude the existence of more atypical and experimental styles.

There are multiple ways avatars can be experimental; previously mentioned were their style and form and the concept of their appearance. However, avatars can also experiment with different ways of expressivity. Face-tracking, in order to translate the user’s expressions onto the avatar’s face, is a relatively noticeable and efficient way to increase expressivity, but it can also be used in other aspects. For example, it is entirely possible to control the color of

an avatar based on the user's facial emotion, to control extra non-human limbs based on this information,⁴⁰ and much more.

Moreover, the experimental nature of expressivity does not have to stop at face-tracking or body-tracking. In fact, it is possible to extend this experimentation by integrating additional hardware. For instance, devices such as sports watches capable of tracking the user's heart rate can be integrated to transmit this data to the avatar. This data can then be utilized to adjust aspects like the avatar's breathing or the rhythm of its heartbeat, adding another layer of expression. Furthermore, more sophisticated equipment like *Brain-Computer Interfaces* (BCI) offer even more expressive potential. An illustrative example is provided by Twitter user *Rantis*⁴¹, who demonstrated using a BCI to control their avatar's ears based on simple brain activity such as focus, concentration, and excitement levels. This enables others to try to get a glimpse into his mental state during a conversation and proposes to unlock a novel and unconventional dimension of expressivity.

Avatar creation

Introduction and Planning

Having previously created an abundance of avatars and having a background in 3D modeling and character concepting,⁴² I bring a wealth of expertise to exploring the avatar creation process. This part serves as a comprehensive guide, going through the steps I usually take to create an expressive virtual reality avatar. My aim is not just to create a mere digital representation but to instill the avatar with a profound purpose - to become an extension of the user's self, a vessel through which they can immerse themselves fully in the virtual world. The avatar's goal will be expressiveness and to instill that “immersed” feeling, correctly displaying and translating the user's emotions and subtle communication cues. As such, this section will focus on relaying my personal experience, informed by the research previously mentioned in the first part of the dissertation, by other various avatars I have created in the past, and by my personal experience in social VR applications.

⁴⁰ Jerry's Mod. (2024, January 31). *A VRChat avatar that uses face-tracking to control the extra non-humanoid ears*. YouTube. https://youtu.be/beYb2nNk_KA?si=groy9CzJ5_lcJUq4

⁴¹ Rantis. (2024, January 15). Video of a user using a BCI to control the ears of their avatar. Twitter. <https://twitter.com/RantiMess/status/1746704510972580061>

⁴² A website showing my previously made avatars. Santos, D. (2020). *Ghostt on Gumroad*. Gumroad. <https://gh0stt.gumroad.com/>

The avatar creation process is an art form requiring careful consideration and attention to detail. It is not merely about designing a visually appealing character but about understanding the psychology of embodiment and presence. The avatar becomes a conduit for self-expression, allowing the user to project their identity and aspirations into the virtual, transcend their physical limitations through the avatar, and adopt a new virtual identity. The feeling of presence, a crucial aspect of the virtual reality experience, is heightened when the user “sees themselves” as the avatar. In other words, when they feel immersed enough so that there is a seamless connection between their actions and the avatar's movements (Slater & Sanchez-Vives, 2014, pp. 24-25), as previously mentioned in the introductory chapters with the context of face-tracking and full-body tracking. As I describe this process of avatar creation, I will delve into the details of designing a virtual entity that is visually captivating and resonates deeply with the user's sense of self. I will explore techniques, tools, and methodologies that enable breathing life into the avatar, infusing it with unique characteristics and a sense of authenticity.

At the core of avatar creation lies a crucial and foundational stage - thoughtful planning and figuring out what the avatar is trying to achieve. Within this broader framework, the initial step entails defining the avatar's purpose and identity, laying the groundwork for the creative process. A pivotal consideration at this juncture is to define whether the avatar is designed to faithfully replicate the user's physical appearance or embody a wholly separate and distinctive identity. While creating the example avatar for this dissertation, I will create one with its own identity and character, deviating from simply trying to replicate the user.

Deviation from a direct replication allows for more creative control, as discussed above, allowing the avatar to have distinct personality traits, emotions, and aesthetics. By unleashing artistic creativity, the avatar evolves into a captivating and visually compelling virtual entity, surpassing the confines of mere replication. This deliberate choice enables the exploration of diverse characteristics, traits, and aesthetics, elevating the avatar to a level of individuality that resonates deeply with its users. As a result, the avatar transcends its original form and takes on a life of its own, forging a deep and engaging connection between the user and their virtual representation. Through thoughtful planning and a deliberate shift away from replication, the avatar becomes a powerful and expressive medium, offering a unique avenue for artistic exploration and a visually arresting embodiment of individuality in the virtual realm.

Avatar self-identity constitutes a fascinating aspect that extends beyond mere visual representation. Delving into this realm provides a unique opportunity to explore how the

avatar's identity can influence and interact with the user's behavior. Drawing attention to the "*Transcending the self in immersive virtual reality*" paper mentioned in the Avatar self-identity chapter (p. 29), it becomes evident that an avatar's predefined social role and perceived identity have the potential to exert a subtle yet significant impact on the user's own sense of identity. While creating the avatar, one of its goals will be this subtle impact, allowing it to subtly change the user's perceived identity while they wear it. In addition, I aim to enhance the presence factor of the avatar by crafting a truly captivating and relatable virtual entity. To achieve this, the primary focus will be on creating an avatar with similar features to the human body, rendering it generally humanoid. One of the most pivotal aspects of this endeavor is matching the avatar's body structure to align closely with a human's. By ensuring that the rough skeletal structure mirrors a human's, I aim to amplify the user's sense of familiarity and embodiment, fostering a deeper connection and immersion within the virtual environment.

Furthermore, the avatar will have an expressive face, another crucial objective in the avatar's design. This aspect is important in the overall avatar design, allowing users to communicate and convey emotions effectively. By endowing the avatar with the capacity to reflect a diverse range of emotions, users can seamlessly project their feelings onto the virtual representation, fostering a deep connection and emotional resonance. This interactivity enhances the immersive experience, enabling users to feel more authentic and present within the virtual environment. I aim to create an avatar that acts as an extension of the user's emotions and gestures. Witnessing their emotions seamlessly translated onto the avatar's face engenders a sense of authenticity, bridging the gap between the virtual self and the real-world self. This sense of connectedness promotes a heightened sense of presence and embodiment, elevating the avatar from a mere digital creation to a powerful conduit for personal expression and communication within the dynamic realm of virtual reality.

Besides presence, other factors can improve the avatar's general look and feel. Firstly, the likability factor. Crafting an appealing and relatable avatar is vital, as it facilitates positive connections and engagement. Striving for an avatar with qualities capable of resonating with a wide range of individuals can enhance its overall appeal and accessibility, fostering a sense of affinity among diverse users. Additionally, uniqueness is a crucial factor to consider. Creating an avatar with distinctive features, traits, and visual aesthetics can set it apart, capturing attention and leaving a lasting impression. A unique avatar has the potential to evoke intrigue and curiosity, thereby fostering deeper engagement and exploration within the virtual realm. To achieve this effect, I will harness the power of fictional anthropomorphic

creatures to achieve this, as they have almost limitless creative potential for visual looks and character. This deliberate departure from human characteristics enables an exploration of the distinction between the individual and the avatar. This approach allows for a more diverse and expansive exploration of identities within virtual reality.

A clear vision can guide the avatar creation process by defining these objectives, enabling intentional design choices that align with the desired goals. These considerations pave the way for the development of avatars that captivate, inspire, and foster meaningful connections within the virtual realm. However, planning extends beyond the avatar itself, encompassing the project as a whole. Key questions arise, such as the necessary applications for creating the avatar, which include *Blender*⁴³, *Substance Painter*⁴⁴, *Photoshop*⁴⁵, *Unity Editor*⁴⁶, and *VRChat* itself. The avatar's scope is carefully considered, balancing originality with feasibility, as overly ambitious plans can impede progress. The goal is to introduce something groundbreaking and unique in both character design and concept, pushing the boundaries of what has been seen before in the virtual world.

In the following stages, I will delve into the multifaceted aspects of avatar creation, from the general ideas of conceptualization and planning to the more specific topics of *3D modeling*, *texture mapping*, and beyond. Each step will be carefully dissected, providing insights and practical guidance for aspiring creators while following the creation of a specific avatar.

Concepting

The concepting phase will delve into the critical phase of creating the avatar's identity. This stage is dedicated to generating visual ideas, seeking inspiration, and crafting a compelling vision for the avatar's visual and conceptual elements. To accomplish this, it is important to start by taking inspiration from multiple existing characters and avatars, examining various sources to identify what resonates with users, and discover patterns that align with popular preferences.

⁴³ Blender Foundation (1994). *Home of the blender project - free and open 3D creation software*. <https://www.blender.org/>

⁴⁴ Adobe Inc. (2014). *3D painting software - Adobe Substance 3D Painter*. <https://www.adobe.com/products/substance3d-painter.html>

⁴⁵ Adobe Inc. (1990). *Official Adobe Photoshop - Photo & Design Software*. <https://www.adobe.com/products/photoshop.html>

⁴⁶ Unity Technologies (2005). *Unity real-time development platform: 3D, 2D, VR & AR engine*. <https://unity.com/>

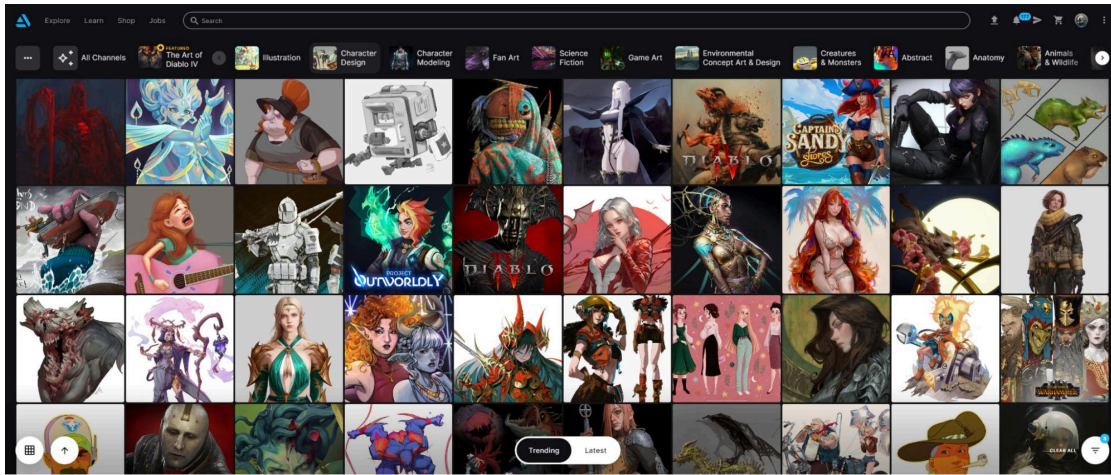


Figure 12a. Epic Games Inc. (2014). *Artstation - explore*. ArtStation. <https://www.artstation.com/>

The process begins by immersing myself in a world of creativity, drawing inspiration from various artistic and cultural sources. I gain a broader perspective on design possibilities and aesthetics by exploring different art forms. *Artstation* is a wellspring of inspiration, particularly from its *Trending* tab (Figure 12a). This platform stands as a beacon of high standards, showcasing the exceptional works of experienced artists and veterans as it serves as an online art portfolio website. The trending section offers valuable insights into popular design trends and ideas that resonate with the creative community and offers a wide selection of art pieces with a broad appeal. This enables me to craft avatars that not only meet high artistic standards but also align with the preferences of the art community outside of virtual reality.

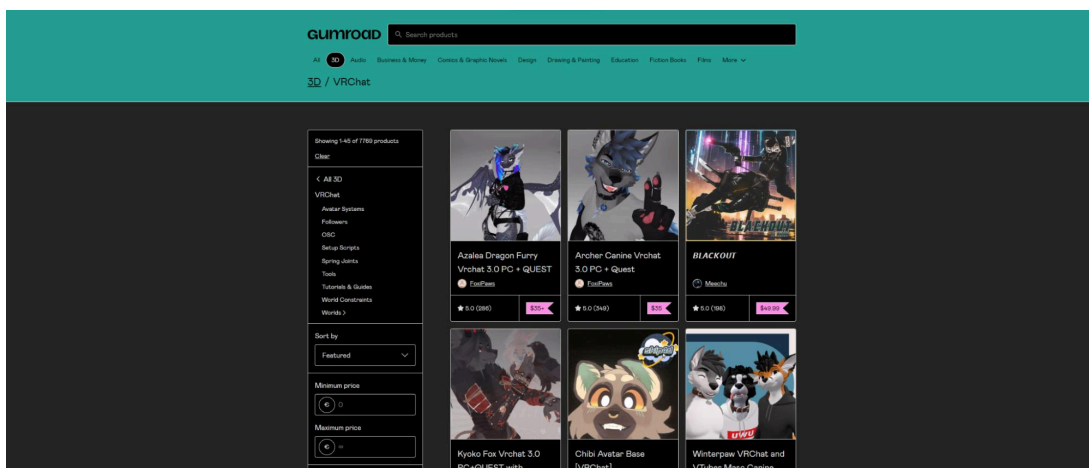


Figure 12b. Gumroad, Inc. (2011). *Gumroad discover: Assets, books, Courses & More*. Gumroad. <https://discover.gumroad.com/>

Moreover, in the pursuit of crafting an avatar tailored for *VRChat*, I thoroughly explore existing popular avatars available on platforms like *Gumroad* (Figure 12b). This research allows me to analyze the characteristics, features, and visual styles that captivate the community and contribute to the avatars' popularity and success within the virtual reality game. By identifying recurring patterns and elements that resonate with users, I gain valuable insights into what captures their imaginations and what they seek in their virtual representations. This understanding forms the basis for most of my previous characters, which have been refined as avatars intending to be sold. Initially created for a comic called *Lifespan*,⁴⁷ these characters involved some experimentation in their design. However, they were subsequently adapted to prioritize meeting user preferences, fulfilling expectations, and enhancing immersion in virtual reality. This approach has effectively produced commercially viable avatars that appeal to a broader audience, exemplified by characters like my *Zuri Sloth* and *Belai Bunny* (as previously showcased in p.28 and p.35 respectively).

For this project, I am opting for a more experimental approach to crafting the avatar. In other words, I will not necessarily adhere to trends set by platforms like *Gumroad*. The aim is not to conform to popular or mainstream characteristics, but to explore different forms of expression and create an avatar that enhances immersion. On a related note, one of my previous avatars, *Parasite Robot*⁴⁸, exemplified this experimental approach. It featured an overgrown parasite inhabiting a dead robot's shell, a relatively unique concept. However, since its design prioritized being expressionless and enigmatic, lacking facial expressivity, it does not fit the criteria for this project. Consequently, I have decided to develop a new avatar tailored for this project, emphasizing expressiveness and originality.

⁴⁷ Santos, D. (2021). *Lifespan*. <https://lifespancomic.wordpress.com/>

⁴⁸ Santos, D. (2021). *Parasite Robot* (VRChat Avatar Base). Gumroad. <https://gh0stt.gumroad.com/l/uwltH>




Figure 12c. Santos, D. (2023). Rough silhouettes. Retrieved August 1, 2023.

To establish the fundamental structure of the avatar, I initiate the process by drawing a series of silhouettes. This approach is highly effective, as the avatar starkly contrasts the background, enabling a clear and distinct definition of its shape. Initially concentrating solely on the character's silhouette ensures that the avatar stands out from others and possesses a bold appearance. Through careful refinement and exploration of various silhouettes, I arrive at a visually compelling and distinctive shape that lays the groundwork for the avatar's subsequent development. The emphasis on silhouette ensures that the avatar possesses a strong and recognizable outline, forming the backbone of its visual identity and setting the stage for further refinement and artistic exploration. For this avatar, my vision included a daring combination of floating limbs, organic material, and machinery. On my second attempt, I achieved a strong and visually compelling silhouette, propelling me into the next phase with enthusiasm and confidence. The distinct blend of elements in the avatar's design promised a unique and captivating virtual presence, driving me forward into the sketching phase.

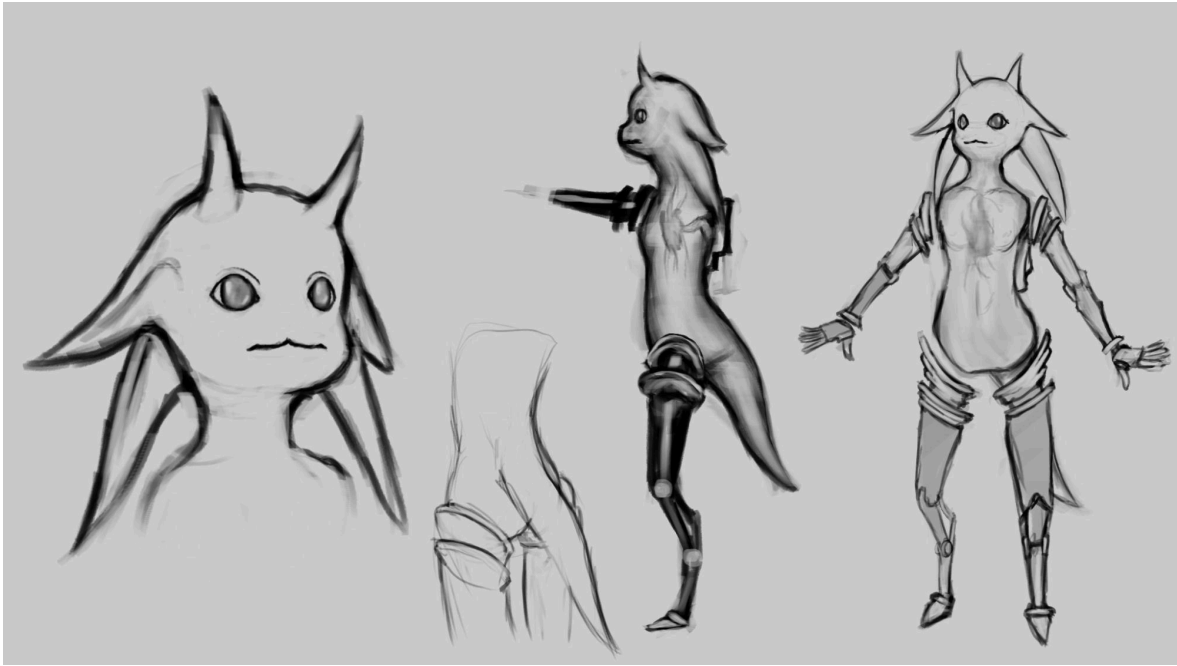


Figure 12d. Santos, D. (2023). Original Avatar Concepts. Retrieved August 1, 2023.

After gathering enough references that appeal to my sense of aesthetics and choosing the silhouette most pleasing and exciting to my eyes, I created various initial sketches based on it, engaging in ideation and iteration while considering all the references, generating multiple ideas, and experimenting with various visual styles, personalities, and themes. These sketches serve as a foundation to translate abstract concepts into visual form, capturing the essence of the avatar-to-be. Although these initial sketches lack intricate detail, they are crucial in refining and shaping the avatar's ultimate form during the sculpting and modeling stages. As the avatar takes shape, it transforms, evolving into a more intricate and refined representation of the original vision, culminating in a virtual entity that encapsulates the essence of the initial ideas while exuding a wealth of intricacies and depth.

In crafting this avatar, I deliberately chose to harness the creative potential of non-human creatures, which enabled experimenting with diverse body shapes, sizes, forms, and colors. The fusion of organic and mechanical elements within this humanoid creature creates a striking contrast, enhancing both aspects and contributing to the avatar's unique identity. As the initial sketch evolved and refined, a compelling narrative emerged, influencing the avatar's shape and form. This symbiotic relationship between visual design and storytelling led to a continuous and iterative process, where each element informed and enriched the other, creating a cycle where visual influences story, which influences visual.

As the creative vision crystallized, a fictional story emerged behind this character: an experiment creating an amalgamation of a slime-based creature imbued with artificial enhancements through mechanical limbs. This intriguing concept led to the creation of cutting-edge technology to mold and morph the creature, resulting in a seamless fusion of organic and mechanical elements. The mechanical limbs were crafted to exude a futuristic and slightly alien appearance with their simple but effective sharp and clean lines, evoking a sense of advanced civilization and technological prowess. Incorporating mechanical limbs adds to the avatar's visual appeal and serves a practical purpose in the 3D medium. Major stretching areas around the elbows and glutes, which tend to suffer from texture stretching during joint rotations, are alleviated using these floating limbs. This technique allows for seamless movement of the arms and legs without causing unwanted texture distortion around these joints. As a result, the avatar maintains its realistic appearance and smooth joint movement, enhancing the overall immersive experience for the user. This clever solution adds to the avatar's unique aesthetic and enhances its functionality.

The avatar's face was designed with face-tracking capabilities in mind. Recognizing the significance of the eyes and mouth as the most expressive features of the human face, I purposefully focused on enhancing and enlarging these key elements. By accentuating these expressive areas, the avatar becomes highly responsive to face-tracking technology, mirroring the user's real emotions and expressions. The eyes were carefully detailed to convey various emotions, from joy and excitement to sorrow and contemplation. Similarly, the mouth was considered similar to a human's, allowing it to transition through various shapes smoothly and naturally, creating an avatar that is truly alive and engaging. As such, the avatar is poised to serve as an authentic extension of the user, accurately reflecting their emotions and reactions. The intricate interplay of these exaggerated facial features enhances the avatar's capacity for emotional expression, forging a closer connection between the user and their virtual representation. By closely aligning the avatar's facial design with the capabilities of face-tracking technology, I aimed to create a powerful and immersive experience, enabling users' emotions to come to life within the virtual environment.

Still, not every part of the character's anatomy coincides with the human anatomy. To reinforce the concept of the avatar's organic components resembling a gel-like substance, I intentionally designed it to possess a unique and distinct transparent appearance. This distinctive characteristic was achieved by infusing intricate detailing on the inner structure of the avatar, providing a glimpse into its intriguing composition. One striking example of this design choice is the incorporation of a vividly rendered beating heart, which resides within

the transparent confines of the avatar's torso. Similarly, the tail, the horns, and the double set of ears are also an example of the disparity between the avatar's anatomy and the human form, with their central objective being to give the avatar more character and enhance its identity.

While there may be concerns about the potential impact of these extra limbs on the user's sense of presence when embodying the avatar, the positioning of these distinct features, which do not obstruct the face and attract limited attention, alleviates these worries. More significantly, an intriguing possibility emerges: users might perceive these supplementary attributes as integral to their physical body. This concept finds support in the previously mentioned "*Transcending the Self in Immersive Virtual Reality*" study and the "*Presence and Phantom Touch*" chapter (pp.32-35), which reveal the human brain's remarkable capacity to assimilate novel bodily elements. This psychological phenomenon suggests that when these elements align with the brain's embodiment tendencies, a fusion occurs, effectively integrating the avatar's extra features into the user's virtual self-concept. This intricate interplay between physical sensations and cognitive processes challenges the notion of rigid boundaries between real and virtual body elements. This cognitive adaptability indicates that the brain can readily incorporate the avatar's unique characteristics, forging a near-seamless unity between the user's real and virtual embodiments. The strategically positioned and carefully designed non-human limbs may transcend their initial strangeness to become integral components of a virtual identity that bridges perception and cognition.

After capturing the avatar's dynamic presence and identity, as the final step of the concepting stage, I will create additional sketches offering different specific perspectives. This includes a front-on view and a side view, allowing for a comprehensive understanding of the avatar's physical attributes and proportions. These views serve as valuable references in the subsequent stages of the avatar creation process, particularly in developing the 3D model. By translating the avatar's three-dimensional form into a two-dimensional isometric representation, I lay the groundwork for the subsequent stages of realizing the avatar in a virtual environment. These illustrations showcase the avatar's visual aesthetic but also aid in accurately translating its physical attributes and features into a fully proportioned and fully realized 3D model.

Structuring

Building upon the groundwork established during the concepting phase, the structuring stage is a pivotal continuation of the avatar creation process. Transitioning from two-dimensional sketches to the third dimension, this stage marks the moment when the avatar truly starts to materialize within the digital environment. At this stage, I harness the power of 3D modeling software, with the chosen platform being *Blender* - an open-source application known for its accessibility and robust capabilities. In *Blender*, I create planes and assign them the images containing the proportion sketches finished in the concepting stage. By aligning them to the frontal and side views, I set up an environment where I can easily switch to those views as I morph the initial shape of the avatar. This strategic alignment facilitates a dynamic environment that facilitates effortless shifts between these key perspectives as I shape and mold the avatar's foundational contours. *Blender*, in this aspect, emerges as an invaluable ally in this creative journey, its intuitive interface and extensive toolkit aiding in the visualization and construction of the avatar.

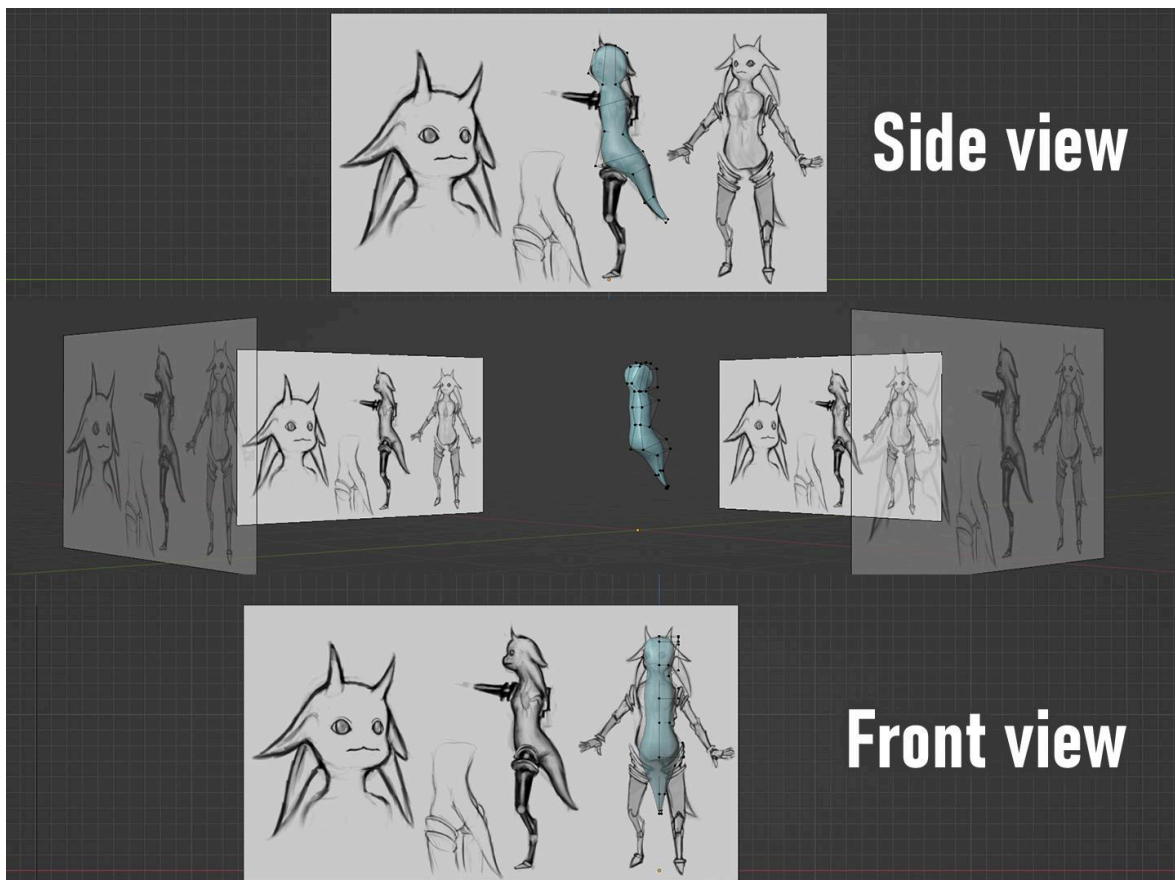


Figure 13a. Santos, D. (2023). Various views of the structuring phase. Retrieved August 18, 2023.

Starting with a basic cube as my foundational template, I employ a tool known as the *subdivision modifier*. This modifier serves as a pivotal technique in 3D modeling, allowing for the transformation of a simple shape into a more intricate and detailed structure. This is achieved by dividing the faces of the initial cube, generating additional vertices and edges that add complexity to the form. This also smooths sharp edges, forming a very rudimentary body form. Afterward, I meticulously cut and extrude faces from the cube using this modifier. This process involves dividing, adding, and extending specific parts of the cube's surface outward or inward, gradually molding the shape to align with my intended design. This intricate dance of extrusion and manipulation is an intricate art form wherein I gradually align the amorphous geometry with the avatar's conceptual blueprints. Each extruded face is carefully moved, manipulated, and aligned to coincide precisely with one of the predefined views, and then the same is done with the other view to align it on both axis. The alignment ensures that the evolving 3D structure faithfully reflects the contours and proportions outlined in my initial sketches.

As this phase unfolds, the cube undergoes an expressive transformation, shedding its initial box-like appearance to adopt a more intricate and purposeful form. The cube's once-geometric edges and planes gradually morph into an intricate shape marked by jagged corners. These corners guide the *subdivision modifier* and the targeted extrusions in their unique arrangement. In the wake of these manipulations, the cube evolves from its original shape, gradually assuming a form resembling the intended avatar's proportions and outline. While its appearance during this mid-process stage may be unconventional, characterized by abundant sharp points and facets, this transient metamorphosis is a vital intermediary step as it shows the edited cube before the subdivision modifier is applied. This evolving shape's smoothed and subdivided form holds the essence of the avatar's ultimate silhouette and physical characteristics. This is where the avatar's true identity begins to take shape, emerging from the cocoon of a seemingly discordant arrangement of vertices and edges.

The *mirror modifier*, another key tool employed in this process, is pivotal in shaping the avatar's symmetry, operating seamlessly along the X-axis and generating a mirrored mesh replication. This mirroring mechanism captures the left and right directions, ensuring adjustments to one side of the model are automatically mirrored onto the other. This modifier proves immensely beneficial in maintaining the avatar's bilateral symmetry, an essential aspect of its design. Given that the avatar's left and right sides are intended to be exact replicas of each other, akin to the bilateral structure found in humans' external appearance, the *mirror modifier* serves as a critical tool. This functionality prevents the creation of a

jarring asymmetry that could disrupt the overall harmony of the avatar's form. This symmetrical consistency aligns seamlessly with users' expectations and enhances the avatar's visual coherence, ultimately contributing to a more appealing and natural embodiment.




Figure 13b. Santos, D. (2023). The final structure of the organics. Retrieved August 18, 2023.

The model assumes a foundational role during this structuring stage, acting as a preliminary form that will serve as the canvas for the subsequent sculpting phase. Given its intermediary nature, it emphasizes establishing a simplified representation rather than delving into intricate details. The primary objective is to make a rudimentary base to serve as the foundation for the ensuing sculpting process. It is important to refrain from dedicating excessive time to refining the model at this juncture, as the subsequent sculpting phase will build upon and refine this rough base and is much more adequate to this procedure. This approach ensures a streamlined progression, allowing for a more focused and detailed model transformation in later stages.

Sculpting

The sculpting phase marks a transformative juncture where the avatar's visual intricacies begin to flourish in remarkable detail. Building upon the foundational structure forged in the preceding structuring stage, this phase employs the advanced capabilities of sculpting tools to detail the model. This process operates like digital clay molding, similar to

sculpting in real life, allowing for the careful addition and removal of virtual material. Through this intricate dance of creation, every nook and cranny is refined, every curve sculpted, and every feature meticulously honed. Subtle imperfections emerge, mirroring the organic aspect of real-life forms and endowing the avatar with a tactile sense of authenticity.

Expanding on the groundwork laid in the previous phase, the first step involves applying the modifiers to solidify their changes within the mesh. This prepares the model for sculpting, ensuring it builds upon the well-defined structure established during the initial phase. Integrating these modifications directly into the mesh, the sculpting phase begins with a strong foundation and works off the final model created by the structuring stage. After applying the modifiers, the subsequent progression involves activating symmetry in the X-axis. This ensures that any modifications introduced during this sculpting phase are mirrored across both sides of the avatar. This compensates for the absence of the *mirror modifier*, which was previously responsible for bilateral mirroring. With symmetry engaged, sculpting symmetrically works properly, and consistent alterations are upheld on both sides of the avatar, preserving its proportional balance and visual consistency, just as the concept drawings of the character showcase.



Figure 14a. Santos, D. (2023). Sculpted organics. Retrieved August 21, 2023.

Commencing the actual sculpting process entails grasping its fundamental methodology. Analogous to sculpting with physical clay, a systematic approach is adopted. It starts by altering the overarching shape from a macroscopic, comprehensive perspective and subsequently delves into incremental refinements. This method involves progressively zooming in on the model, thoroughly augmenting its intricacies in the final stages. This sequential progression ensures that the initial work establishes a coherent and expansive foundation upon which intricate elements are carefully added and integrated. I implement this approach by adding virtual clay atop the existing material. This involves extending it outward and shaping convex forms as needed. Conversely, Blender possesses a wealth of sculpting tools that allow for many changes to the model, like pinching, smoothing, and subtracting material to create cavities and concave shapes, thus instilling depth and dimensionality into the evolving model. Similarly to drawing, these additions and subtractions can be drawn onto the form, creating contrasting lines that change the depth between surfaces. This adds small details like textures, wrinkles, and patterns towards the end of the phase to make it look more natural and lifelike. This step makes the avatar more appealing and closer to looking like a detailed, realistic character.

You might have noticed that the sculpture lacks features like eyes, limbs, and a translucent appearance from its concept art. However, these omissions are intentional and serve specific purposes within the design process. Eyes are absent because they will be created separately from the main mesh later. This decision prevents them from casting static shadows during the texture-baking process, which would occur if integrated into the current sculpture. Since eyes can rotate within their sockets, having those shadows fixed in place would result in an unnatural visual effect when the eyes move. The limbs are intentionally excluded because they will be developed using a distinct technique called *hard-surface modeling*, which differs from the current sculpting approach. This will be explored more comprehensively in the subsequent chapter, ensuring a focused and effective development of this specific aspect of the avatar. The transparent appearance featured in the concept art will be introduced during the *texture creation* process. Deliberately avoiding its application during sculpting allows for a clearer view of how light and shadow interact with the model's surface during its formation. If the transparency effect were applied during sculpting, it could potentially complicate the process by obstructing the visibility of these vital aspects of the model's development.



Figure 14b. Santos, D. (2023). Sculpted face. Retrieved August 22, 2023.

The innate tendency to direct attention toward faces during social engagements is central to human interaction, as explored in the “*Importance of the face*” (pp.12-16) and “*VR Face-Tracking*” (pp.19-22) chapters. Therefore, a significant emphasis should be placed on shaping the avatar's face within the sculpting phase. This concentrated effort is grounded in the understanding that observers instinctively fixate on facial expressions during virtual interactions, underscoring the importance of a compelling facial representation. Detailing the avatar's face involves more than just a visually appealing facade. It necessitates a comprehension of the web of muscular interactions underlying human expressions. By adhering to this understanding, the avatar's design is structured to resonate with the biomechanical realities of the human visage. This entails ensuring the avatar's facial muscles are congruent with their anatomical counterparts, even if the overall shape deviates creatively. This adherence to established muscle groups underpins the user's ability to interpret the avatar's expressions as authentic, fostering an intuitive connection between the avatar's emotional cues and the user's interpretation. The dynamic interplay of these facial muscle groups enables the avatar to convey a spectrum of emotions, transcending the limitations of

mere visual design. This sculpting approach acts as an interpretive bridge between human expressions and virtual representation, creating a fluid channel for communication and empathy in the virtual reality environment. In essence, the sculpting process is not confined to aesthetics alone. Still, it extends to psychological resonance, where authenticity and relatability converge to create a truly engaging and expressive virtual persona.

However, as the sculpting process requires getting closer to the minute, it becomes apparent that the model needs more polygons and vertices to move around and generate the level of quality necessary to perceive those details. This requirement is met by increasing the model's polygon and vertex count, which allows for capturing the nuances and intricacies that give it a more realistic and refined appearance. To achieve this, two primary methods can be employed, each with its unique approach. The first method, *Voxel Remeshing*⁴⁹, is a tried-and-true technique that has stood the test of time. It operates through a relatively straightforward process - when activated, it analyzes the model's curvature and layout. Subsequently, it strategically places points on a three-dimensional grid, aligning them with the contours of the model's surface. By reducing the dimensions of this grid, the resolution is effectively increased, resulting in a greater number of polygons. This method's advantage lies in its stability and predictability, making it a preferred choice for generating more polygons and enhancing detail. On the other hand, the more recent technique, known as *Dynamic Topology*⁵⁰, offers a real-time approach to the polygon-count enhancement process. It dynamically adjusts the resolution as virtual clay is added or removed during sculpting. This dynamic adjustment is influenced by the proximity of the perspective to the model, ensuring that detail remains optimal and appropriate for the viewer's proximity and works with the philosophy of being hidden in the background without the sculptor noticing it. However, a notable drawback of *Dynamic Topology* is its real-time operation, which can reduce computer performance. The model's complexity and the view's closeness can impact the system's responsiveness.

Given these considerations, while *Dynamic Topology* offers a dynamic and context-sensitive approach to polygon generation, it is essential to factor in potential performance implications. As a result, *Voxel Remeshing* often emerges as the preferred choice, offering a stable and reliable method for enhancing model detail without

⁴⁹ An official video by the Blender Foundation explaining the usefulness of the Voxel Remeshing algorithm and how it works. Dobarro, P. (2019, October 26). The new sculpting workflow in Blender - Pablo Dobarro. YouTube. <https://www.youtube.com/watch?v=lxkyA4Xslzs>

⁵⁰ An article explaining what Dynamic Topology is and how it works. Garofalo, E. (2022, February 17). Dyntopo in Blender: Dynamic Topology for beginners. MUO. <https://www.makeuseof.com/dyntopo-for-beginners/>

compromising overall system performance. The strategic selection of these methods is crucial in striking a balance between detail enhancement and computational efficiency during the sculpting phase.




Figure 14c. Santos, D. (2023). Topology after applying a voxel remesh algorithm. Retrieved August 22, 2023.

The sculpting process is a powerful tool for creating organic-looking models imbued with authenticity. This technique generates imperfections, subtle indentations, and forms that deviate from rigid, perfect lines. This deliberate departure from flawless geometry adds realism and character to the organic part of the model, reflecting the natural variations found in the real world. Nevertheless, transitioning to creating the robotic components within the model, a noteworthy revelation emerges: sculpting, while adept at organically imperfect forms, proves inadequate for the perfect precision and flawless contours demanded by mechanical elements, which is one of the reasons why the concept art of the avatar contains mechanical elements - to exemplify and demonstrate the creation of both organic sculpting and hard-surface modeling, making it the perfect avatar for demonstrating a wide array of aspects to the overall creation process. The intricacies of the robotic parts necessitate an alternative approach that caters to their specific design requirements. In response to this demand for exactitude, the hard-surface modeling technique emerges as the ideal solution, enabling the crafting of these mechanical components with utmost precision and ensuring that their shapes meet the exacting standards demanded by their mechanical nature.

Hard-Surface Modeling

Like the structuring process, the hard-surface modeling technique commences with basic primitive shapes bolstered by the subdivision modifier. However, it distinctly diverges as its detailing intricacies are predominantly self-contained. This technique's strength lies in its proficiency in crafting precise geometries from elementary forms, obviating the necessity to transition to the sculpting phase.




Figure 15a. Santos, D. (2023). The results of the subdivision modifier. Retrieved August 23, 2023.

During the creation of this avatar, the hard-surface modeling technique will be employed to shape the legs, arms, and all mechanical components requiring extreme geometrical precision. As illustrated in the visual representation (*Figure 15a*), these flawless geometrical structures demand the utilization of the subdivision modifier and its associated functionalities. The diagram presents distinct perspectives of the model, each offering unique insights. On the left, the unmodified base form is depicted. The right side showcases the model with the subdivision modifier applied, revealing refined details and smoother contours. Meanwhile, the center previews the subdivision modifier's impact, showing both the original mesh (transparent) and the subdivided result (opaque), allowing adjustments before the final application. This visual representation highlights the diverse viewpoints these different modes of presentation offer in understanding the model's intricacies.

Integrating seamlessly into this workflow, *Blender's edge crease*⁵¹ function is pivotal in conceiving the avatar's mechanical components. Designed to synergize almost exclusively with the subdivision modifier, this tool, as its name suggests, imparts creased edges to the subdivided mesh, mirroring the precision of real-world geometric structures. Its strategic application ensures that the avatar's mechanical elements attain an authentic appearance, fitting seamlessly within the broader design.




Figure 15b. Santos, D. (2023). The effects of edge creasing a subdivided mesh using a value of 0.5. Retrieved August 24, 2023.

Within *Blender*, the edge crease function operates by instructing the subdivision modifier on the distribution of subdivisions across the initial faces. When a creased edge bridges two faces, the subdivisions gravitate towards the original edge, resulting in heightened definition (*Figure 15b*). Here, the creased edges, visualized in a deep red hue, starkly contrast with the remaining edges depicted in black. Consequently, any subdivided edge influenced by a creased edge will retain a more pronounced sharpness, contrasting with the relatively smoother appearance of unaffected edges.

However, the functionality of *edge crease* goes beyond its basic use. Unlike a simple on-off mechanism, edge creasing operates on a gradient scale, ranging from 0, which means no creasing, to 1, indicating maximum creasing. This nuanced approach allows precise control over how much creasing is applied to each edge. This control also extends to vertices,

⁵¹ A video demonstrating how the edge crease functions in addition to the subdivision modifier. PIXXO 3D. (2020, April 27). *Blender: Edge creases and how to use them*. YouTube. <https://www.youtube.com/watch?v=rhJTaD3BzWo>

known as *vertex creasing*. This advanced feature applies the same functionalities of *edge crease* to vertices, enabling even finer adjustments contributing to the detailed refinement of the model's overall shape. By adjusting this creasing factor, I can achieve a range of sharpness or smoothness that suits the model's specific intricacies and design requirements.



Figure 15c. Santos, D. (2023). The final hard-surface models, along with the organic sculpt of the body. Retrieved August 25, 2023.

The *hard-surface modeling* process creates robotic limbs with flawlessly shaped parts characterized by precise mathematical curves and straight lines that are challenging to achieve with sculpting. It is important to note that objects intersecting in the final avatar must be physically separated in the digital model. This is crucial for the following *texture-baking* and painting steps, ensuring a smooth integration of these precisely shaped components into the avatar's overall design. This separation applies to the leg joints, the arm joints, and the eyes, ensuring they will not cast a sharp ambient occlusion shadow during the texture baking.




Figure 15d. Santos, D. (2023). High-poly backpack-heart model. Retrieved August 30, 2023.

Nonetheless, I deliberately deferred the creation of the backpack, reserving it as the concluding segment of the avatar concept. This strategic choice exemplifies a harmonious fusion of the sculpting and hard-surface modeling techniques, effectively yielding a distinctive piece that embodies both semi-organic and semi-robotic characteristics. The effect achieved through this approach highlights the potential for seamless integration of these diverse methods, the central heart being the organic material coherently joined with the precise form of the pumps, which then join into the robotic backpack.

This piece began with applying hard-surface modeling techniques to create the backpack itself. Following this, designing the pumps adopted a similar methodology to *hard-surface modeling*, using mathematics to achieve perfect shapes. Notably, the difference was the introduction of the *bezier curve*⁵² tool, which marked a significant advancement. This tool facilitated the precise generation of bezier curves, employing interpolation points and smoothing handles to construct intricate three-dimensional mathematical curves, forming the basis for the radius of a tube. By extending a perfect circle along this curve, the outer edges of the tube emerged. The finalization of this process involved sealing the tube's ends with a face, ensuring structural integrity. After that, I started sculpting the heart using the techniques described in the sculpting section. This means starting from the general view and honing in

⁵² A video explaining the different types of curves in Blender, including the bezier curve. King, R. (2021, May 31). How to use curves in Blender (tutorial). YouTube. <https://www.youtube.com/watch?v=Ve9h7-E8EuM>

on the details: pushing, moving, adding, and removing clay. However, since this is a smaller model than the torso, I could use *dynamic topology* to automatically make the vertex density high or low in real-time without affecting the performance. In the end, I had three pieces that overlapped and did not connect, but I needed to connect them to look seamless.

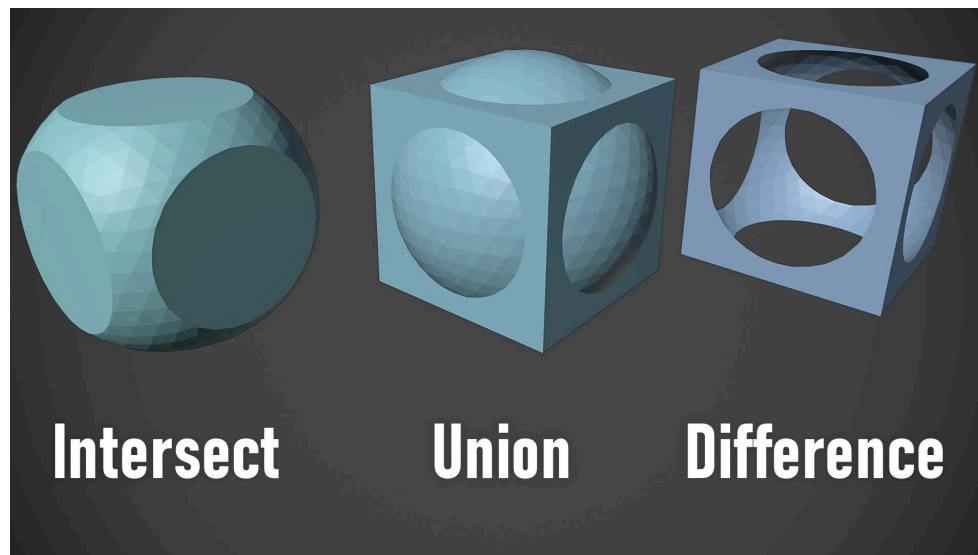


Figure 15e. Santos, D. (2023). Various effects of the boolean modifier. Retrieved August 30, 2023.

This is where the *Boolean modifier*⁵³ truly shines in its functionality. The *Boolean modifier* operates by combining two objects in three distinct ways. The first method, *Intersect*, generates a new object formed solely from the overlapping space of both input objects, effectively removing the volume exclusive to each object. The second technique, known as *Union*, fuses the two objects together, resulting in a unified object encompassing the combined volumes of both components. Lastly, the third operation, labeled *Difference*, centers on the original object to which the modifier was initially applied (such as the cube depicted in *Figure 15e*) and subtracts the volume corresponding to the target object (in this case, the sphere). This operation produces a final object that appears as if the first shape has been hollowed out by the second. Notably, the *Difference* operation's outcome is influenced by the order in which the objects were assigned to the modifier, unlike *Intersect* and *Union*.

Nevertheless, the ideal operation for amalgamating the backpack, pipes, and heart is undoubtedly the *Union* operation. I initially merged the backpack with the pipes to facilitate this process, seamlessly combining their forms. Subsequently, this unified backpack-pipes form was merged with the heart, resulting in an intricate composite unit comprising these

⁵³ A video demonstration the different types of boolean operations within Blender. Jayanam. (2018, January 22). Blender boolean modifier tutorial. YouTube. https://www.youtube.com/watch?v=RDKt_INAHss

three distinct components with interconnected topological integrity. With the application of the respective modifiers, the subsequent steps encompassed refining the connection points through sculpting, ensuring a cohesive and smooth integration. To enhance the seamlessness of these connections, dynamic topology was engaged, allowing for the organic blending of these regions and culminating in a cohesive, unified whole.

It is important to note that the sculpting and hard-surface phases mark a crucial stage in the avatar creation process, wherein the model evolves into its most refined and detailed form. So much so that, as the detailing process reaches its culmination, the final model, with its numerous nuanced subtleties and intricate elements, can potentially transcend the limits of real-time computer rendering. This level of intricacy ensures that the avatar possesses a wealth of subtle nuances and intricate features. However, it requires additional optimization and adaptation for real-time rendering within virtual reality environments and translating those details into a simplified form. With this in mind, I move on to the next phase of the avatar creation process.

Retopology

With the culmination of the detailing phase for the avatar, the focus now shifts towards a pivotal aspect of the creative process: optimization. The intricate sculpture is a testament to digital artistry's potential; however, a meticulous optimization process is essential to seamlessly integrate this intricate creation into real-time environments such as game engines. This phase can involve diverse approaches, from automated techniques to well-established industry standards, extending to precisely tailored methods for specific use cases. Despite their diversity, these strategies share a common objective: to maximize the optimization level while upholding visual quality. A shared characteristic among these optimization processes lies in reducing triangles or polygons. This involves refining the highly detailed model, which boasts an extensive count of triangles and polygons, and dramatically reducing this count while preserving the model's overall form and outline. In many instances, especially in the context of video-game characters, the final count is less than 1% of the original sculpt's tally, achieved through a careful and intricate transformation that ensures the model's essence remains intact.

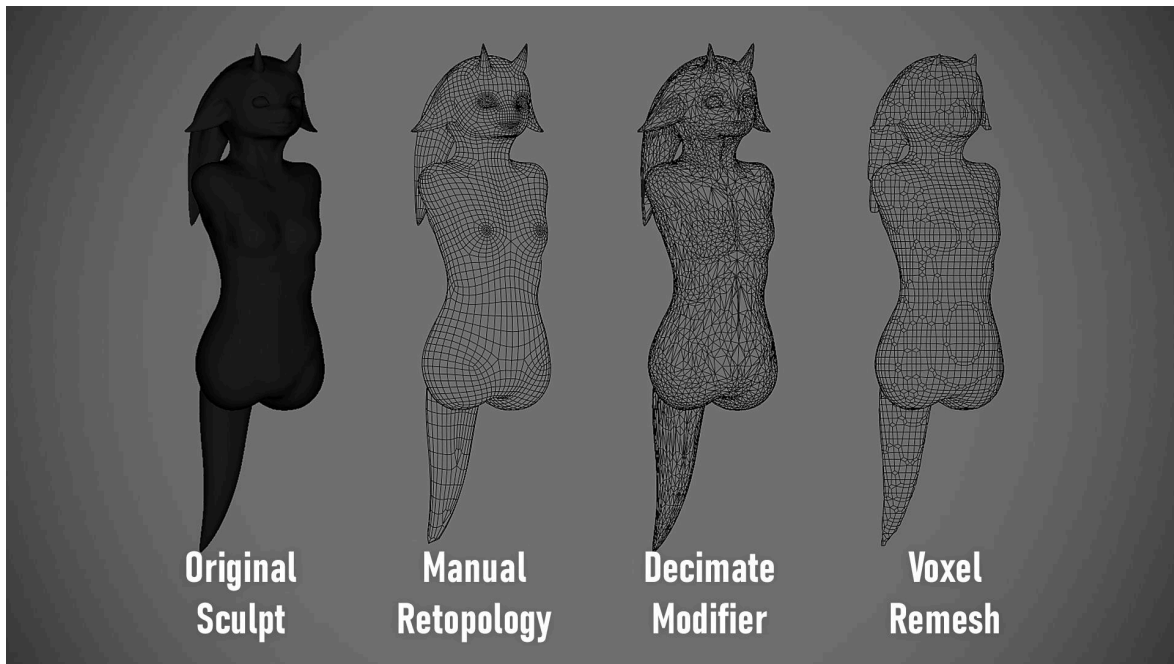


Figure 16a. Santos, D. (2023). Various methods of retopology. Retrieved August 30, 2023.

The fastest approach to reducing polygon counts involves utilizing automated tools that harness computer processing power through a gamut of operations. These operations encompass a spectrum of tasks, from identifying specific angles to performing un-subdivision. Among these methods, one of the earliest algorithms employed is the same technique outlined in the “*Sculpting*” chapter (pp.51-56), the *Voxel Remesh*. However, in this context, Voxel Remesh serves a different purpose. It transforms a highly detailed sculpt into a low-poly model using a larger grid size instead. However, this operation bears a notable limitation: it tends to sacrifice perceived details. In other words, while it effectively reduces the mesh's complexity, it often fails to preserve the density of intricate details, diminishing the model's nuanced features.

Another automated technique involves the application of the *Decimate Modifier*⁵⁴, a more intricate algorithm than *Voxel Remeshing*. This method seeks to streamline polygon counts by leveraging a more comprehensive range of information inherent to the model, encompassing detail density and angle magnitude. Nevertheless, the *Decimate Modifier* has its drawbacks. One significant concern pertains to the potential creation of a topology that lacks organization. This drawback holds particular implications for models that undergo morphing and movement via a skeletal rig, a characteristic common to character models. In

⁵⁴ A video demonstrating the Decimate modifier, how it works, and some of its advantages and disadvantages. MK Graphics. (2021, August 13). Decimate modifier | blender tutorial. YouTube. https://www.youtube.com/watch?v=JzOC_xtFV0E

these instances, a haphazard topology could undermine the model's ability to maintain its intended form during dynamic transformations. However, these automated methods are tailored to their distinct functions, primarily highlighting the rapidity with which they execute. This efficiency is particularly noteworthy when contrasted with the manual approach. Furthermore, the decimation process proves valuable in specific contexts. For instance, it excels in swiftly generating low-poly versions that do not morph, a favorable use-case for static props like environmental elements or background components. Such items, devoid of the necessity for morphing abilities, benefit from the streamlined results of the decimation process.




Figure 16b. Santos, D. (2023). The high-poly sculpt and the low-poly retopology. Retrieved September 2, 2023.

In this particular project involving the avatar, I opted for the manual Retopology approach. This method involves the precise placement of each polygon, aligning them with the contours of the sculpt. The rationale behind this choice primarily rests on its manual nature, which inherently affords greater control to the artist. This elevated level of control empowers the artist to make informed decisions regarding the precise positioning of polygons, tailoring them to suit the specific utilization requirements of the 3D model.

In my retopology process, I employ a range of modifiers to streamline the workflow and enhance efficiency. The process begins with placing an initial polygon on the sculpt's

chest area. To achieve this, I activate *Face Snapping*⁵⁵, a critical feature that ensures precise alignment of the polygon's vertices with the surface of the chest. Enabling this feature facilitates the seamless extraction and placement of subsequent polygons, contributing significantly to the overall workflow efficiency. Following this initial step, I proceed by enabling the *Mirror modifier*. As previously elucidated in the “*Hard-Surface Modeling*” chapter (pp.57-62), this functionality ensures that any modification made during the retopology process is mirrored accurately on the opposite side. The *Mirror modifier* proves especially advantageous in this context, as it expedites the current workflow and lays the foundation for smoother progress in subsequent phases, notably in the “*Rigging and Weight-Painting*” (pp.80-83) and “*Iteration*” (pp.91-97) chapters. Maintaining perfect mirroring throughout the retopology process is instrumental in preventing potential complications and inconsistencies down the line. Subsequently, I activate the *Subdivision modifier*, a tool previously introduced in the “*Structuring*” chapter (pp.49-51). Its role in the retopology process aligns with its function there, primarily smoothing the mesh and enhancing the overall organization of the topology. This automated operation proves exceptionally time-efficient, streamlining the workflow significantly. With the *Subdivision modifier* in play, I only need to manually extrude a single polygon, which then undergoes automatic subdivision into four, simplifying the process considerably compared to the manual extrusion of four separate polygons. After these preliminary adjustments, a pivotal step in this context involves the utilization of the *Shrinkwrap Modifier*⁵⁶. By designating the sculpt as the target mesh and applying this modifier to the retopologized plane, a precise and neat wrapping of the plane around the sculpt is achieved, akin to the process of vacuum-forming plastic in the physical world. This particular step is very important in the retopology workflow, serving as its foundational element. Essentially, the *Shrinkwrap Modifier* is a ubiquitous tool employed in most retopology workflows across various 3D modeling programs, owing to its straightforward mathematical principles and the indispensable role it plays.

⁵⁵ An official Blender article exposing the different kinds of snapping that are possible. Blender Foundation. (2023, October 30). Snapping. Snapping - Blender Manual. <https://docs.blender.org/manual/en/latest/editors/3dview/controls/snapping.html>

⁵⁶ An official article explaining the different operations within the Shrinkwrap Modifier, as well as a brief explanation of the modifier itself. Blender Foundation. (2023, October 30). Shrinkwrap modifier. Shrinkwrap Modifier - Blender Manual. <https://docs.blender.org/manual/en/latest/modeling/modifiers/deform/shrinkwrap.html>




Figure 16c. Santos, D. (2023). Effects of the modifiers utilized in the retopology process. Retrieved September 3, 2023.

A pivotal consideration in this process is the precise order of each modifier within the stack. When a model possesses a series of modifiers, as is the case with the retopologized model, the arrangement of these modifiers in the stack significantly influences the outcomes of each. As such, it is imperative to recognize that modifiers are applied sequentially from the top to the bottom of the stack. In the specific context of retopology, the *Mirror modifier* must occupy the top position in the stack. Failing to adhere to this order, mainly by applying the *Mirror modifier* after the *Subdivision modifier*, would lead to undesired results. Specifically, the mesh would undergo subdivision and smoothing on only one side of the model before mirroring rather than ensuring a symmetrical mirroring of the mesh prior to subdivision. Such an incorrect sequencing would introduce artifacts along the symmetry plane at the model's center. Hence, the correct stacking order is pivotal to achieving the desired retopology results. Furthermore, it is essential to emphasize that in this particular context, the *Shrinkwrap modifier* must be positioned after the modifier stack. This deliberate order is necessary because the *Shrinkwrap modifier* securely adheres each vertex and sub-vertex to the sculpt's surface. This adherence must occur after the model has undergone subdivision and smoothing actions performed by the *Subdivision modifier* earlier in the stack. Should the *Shrinkwrap modifier* precede the *Subdivision modifier*, its influence would be confined solely to the original, non-subdivided vertices. Ineffectively, it would not extend its impact to the newly created subdivided vertices. Consequently, such misalignment of modifier order would result in uneven smoothing across various areas of the mesh. The subdivided vertices could drift

inside the sculpt, causing an undesirable outcome. Therefore, in this context, the correct sequence would be the mirror being the first, the subdivision the second, and the shrinkwrap being the final, ensuring it acts upon the mirrored, subdivided, and smoothed mesh as intended.

After understanding the tools, I will now explain the essential principles governing polygon placement and the orientation of topology. While the specifics of polygon placement can be highly subjective and vary significantly between different models, specific fundamental guidelines underscore good topology practices. These principles offer valuable insights into effectively structuring and organizing a 3D model's topology for optimal results. The initial general rule to underscore is the importance of contextual planning for the model's intended use. For this project's specific avatar, designed to be worn and experienced from a first-person perspective within a social environment, consideration is given to the expected functionalities and visual requirements. The avatar must exhibit the ability to articulate its joints without undesirable deformations. Furthermore, specific areas demand heightened attention to detail. Notably, the face plays a pivotal role in mirroring and conveying the user's expressions, necessitating a greater allocation of polygons to facilitate nuanced expressions. Likewise, the hands, as they are positioned in proximity to the user's perspective and serve as a conduit for interacting with the virtual environment, demand a comparable focus on detail. These strategically allocated areas of interest heighten the entire model's overall fidelity and perceptual quality.

Furthermore, it is advisable to prioritize using polygons over triangles within the topology. This choice stems from several key advantages. Foremost among these is the inherent stability in smooth shading that polygons offer compared to triangles. Smooth shading across polygons exhibits greater predictability and consistency in rendering than the sometimes unpredictable behavior seen in transitioning from triangles to polygons. Another significant advantage is the versatility of polygons when it comes to subdivisions. Unlike triangles, which only subdivide into polygons, polygons can undergo subdivisions and unsubdivisions more seamlessly. This capacity is particularly beneficial for refining the model's level of detail. For instance, four polygons can readily merge to form a larger polygon, facilitating adjustments in the model's structure or simplification when necessary. In contrast, triangles lack this ease of manipulation, making polygons the preferred choice for achieving adaptable, high-quality topology.

Another fundamental guideline in effective polygon placement involves the creation of a structured grid pattern. This entails configuring the mesh to approximate a grid,

minimizing distortions and irregularities. This deliberate organization serves several crucial purposes. Firstly, it promotes smooth shading transitions between neighboring polygons, minimizing unsightly artifacts and discrepancies in the visual presentation. Additionally, it contributes to the avatar's overall aesthetics by enhancing the clarity of its contours, resulting in a refined, less jagged silhouette. To successfully implement this guideline, particular emphasis is placed on ensuring that connections from edge to edge produce a clean, gently curved path that gracefully traces the mesh's surface, devoid of zig-zag or erratic patterns.




Figure 16d. Santos, D. (2023). Difference between grid interiors and grid corners in the retopology process. Retrieved September 5, 2023.

To maintain optimal topology, it is crucial to adhere to the underlying contours of intricate details across the sculpt, such as the subtle curves of eyelids and lips, the complex geometry of areas like skin folds, or the junction between horns and the head. In some instances, rigidly adhering to a uniform grid pattern may prove challenging while following the patterns of these details. Thus, it becomes imperative to recognize these scenarios and adjust connection points accordingly. The most prevalent and effective strategy to address such situations involves segregating grids and carefully applying corners where they meet. In the illustrative example (*Figure 16d*), the interior of a grid, highlighted in blue, necessitates the seamless conformity of its curvature to that of the underlying sphere. To achieve this transition and interconnection with an adjacent grid, the strategic application of a corner, highlighted in green, becomes pivotal. Employing this technique results in a consistently exemplary topology that accommodates the majority of use cases, ensuring the retention of intricate details while preserving the overall efficiency of the model.

One final vital guideline pertains to maintaining an awareness of the overall polygon count or, more commonly expressed, the final tally of triangles (also known as *tris*). In a polygon-based model like the avatar being created, this triangle figure is roughly twice the number of polygons. The rationale behind this doubling effect is rooted in the nature of polygons, where each polygon invariably divides into two triangles. Using *Overwatch*⁵⁷ and its characters as a reference point, each character is crafted with a triangle count ranging from 30,000 to 50,000⁵⁸. However, it is crucial to recognize that this specific numerical benchmark may not be an appropriate yardstick for our avatar creation endeavor. The rationale for this deviation lies in the unique contexts that define each model. *Overwatch* characters typically occupy a distinct space within the gaming landscape. They are often positioned further away from the camera's perspective, necessitating a different approach in terms of polygon allocation. Conversely, virtual reality avatars are intended to be directly embodied by users, placing them closer to the viewer's perspective. Furthermore, as a competitive shooter released in 2016, *Overwatch* prioritized broad accessibility across a spectrum of computer and console hardware, requiring a stringent commitment to performance optimization. In contrast, virtual reality avatars, designed for casual social interactions, benefit from the enhanced computational abilities of more modern computers compared to their 2016 counterparts. This disparity results in a more lenient approach to triangle count. Another key distinguishing feature between the two contexts is the presence of face-tracking capabilities in the avatar being created. This advanced functionality demands a heightened polygon density around the facial region to enable nuanced expressions and interactions. In contrast, *Overwatch* characters do not require as many intricacies in facial animation.

⁵⁷ Blizzard. (2016). *Overwatch*. <https://overwatch.blizzard.com/>

⁵⁸ Agilethief. (2016, May). Technical study: *Overwatch* [image heavy]. polycount. <https://polycount.com/discussion/170394/technical-study-overwatch-image-heavy>

Avatar Quality	Excellent	Good	Medium	Poor
Polygons	32,000	70,000	70,000	70,000
Bounds Size ¹	2.5m x 2.5m x 2.5m	4m x 4m x 4m	5m x 6m x 5m	5m x 6m x 5m
Texture Memory	40 MB	75 MB	110 MB	150 MB
Skinned Meshes	1	2	8	16
Meshes	4	8	16	24
Material Slots	4	8	16	32

Figure 16e. VRChat Inc. (2023, July 19). VRChat’s performance ranking system. Performance ranks: Vrchat creation. Retrieved September 6, 2023. <https://creators.vrchat.com/avatars/avatar-performance-ranking-system>

Therefore, to determine a suitable target triangle count for our avatar, we turn to *VRChat's* guidelines, a reputable source for creating avatars within its virtual reality environments. According to their recommendations, an avatar's maximum advised triangle count (incorrectly labeled as *Polygons*) stands at 75,000 tris. Considering all these factors, our avatar will aim for a count ranging between 50,000 and 75,000 tris, equivalent to approximately 25,000 to 37,500 polygons. However, it is crucial to acknowledge the influence of modifiers, which alter the polygon count once applied. For instance, the *Mirror modifier* effectively doubles the polygons from one side to the other, while the *Subdivision modifier* quadruples each manually placed polygon. Consequently, when we focus on the manually positioned polygons before being affected by the modifiers, our target range falls between 3,125 and 4,687. This careful calibration ensures that our avatar strikes a harmonious balance between visual fidelity and optimal performance in virtual reality.




Figure 16f. Santos, D. (2023). Difference between flat shading and smooth shading. Retrieved September 7, 2023.

Upon completing the manual retopology process, several final steps must be executed to prepare the avatar for its subsequent phases. Firstly, due to the inherent characteristics of the *Shrinkwrap modifier*, not all vertices may be perfectly aligned symmetrically. The "*Snap to Symmetry*" option becomes instrumental in rectifying this potential misalignment. This seemingly minor operation is vital to avoiding significant complications in the future when working with symmetry, a prevalent aspect of the 3D modeling process. Following this pivotal step, transforming the shading type from flat to smooth is crucial. This alteration significantly impacts the avatar's visual presentation and allows the future *texture-baking* process to proceed smoothly. The smooth shader engenders a graceful gradient between adjacent polygons, facilitating the smooth transition of light across the surface. Conversely, the flat shader imparts a distinctly sharp edge to each polygon, resulting in a more pronounced and angular appearance. Additionally, verifying the topology's alignment and integrity at this juncture is imperative. Addressing any discrepancies or inaccuracies at this stage is essential, as attempting to rectify such issues after commencing the *texture creation* process can be exceptionally challenging and time-consuming.

UV Editing

The next phase after retopology, *UV editing* is a crucial step in the 3D modeling process, as it prepares the model for applying textures. To grasp this concept, it is essential to

comprehend how 2D images are wrapped around 3D models. Akin to taking a paper cube and carefully unfolding it until it lies flat. In this unwrapped, 2D state, the surface of the 3D model is laid out in a manner that resembles a flat, 2D canvas. On this canvas, the 2D image, which contains the texture, can be precisely placed, accounting for every nook and cranny of the model's surface. Once the image is applied, the model is "re-folded" back into its original three-dimensional shape, resulting in the intricate and realistic application of textures across its surface. This process, named *UV Unwrapping*⁵⁹, ensures that textures seamlessly conform to the contours and details of the 3D model, enhancing its visual fidelity and realism.




Figure 17a. Santos, D. (2023). 2D texture being applied to a 3D cube. Retrieved September 8, 2023.

Like the analogy of unfolding a paper cube to create a flat template, *UV editing* involves making precise incisions along the model's edges to facilitate its transformation into a 2D canvas. These strategic cuts, known as *UV seams*, are strategically positioned to ensure that no areas exhibit visible texture stretching or distortion when the model is unwrapped onto a two-dimensional surface. *UV seams* serve as virtual "fold lines" that guide the unwrapping process, allowing the 3D model's surface to be laid out to minimize texture deformities and ensure a seamless and natural texture application. In the *UV Editing* phase, creating seams is the bulk of the work, making it the central and most substantial component of our current endeavor.

⁵⁹ An official Blender article which contains various aspects of UV editing. Blender Foundation. (2023, October 30). UV tools. UV Tools - Blender Manual. <https://docs.blender.org/manual/en/latest/modeling/meshes/editing/uv.html>

Much like retopology, various methods are available for automating the creation of *UV seams*. However, given the relatively swift pace of this phase, I opt for the manual placement of seams. This choice is driven by the fact that the speed gained from automation does not outweigh the imperfections it might introduce. Before diving into the manual seam placement process, it is essential to consolidate all objects that will share the same material. This consolidation organizes the *UV unwrapping* process and ensures consistency in texture application. Additionally, it is crucial to set up a placeholder texture, which serves a dual purpose: aiding in assessing the texture density of each island, providing a visual reference for how textures will be applied, and helping to quickly identify any texture stretching instances that may occur during the unwrapping.




Figure 17b. Santos, D. (2023). Avatar's body and its seams (in red). Retrieved September 11, 2023.

When assigning seams (indicated by the red edges in *Figure 17b*) I follow a set of principles that guide this part of *UV editing*. These principles streamline the unwrapping process while upholding the quality and coherence of the final textured model. They encompass strategic seam placement, minimized island count, and uniform texture density. Strategic seam placement involves thoughtfully positioning seams in areas with reduced visibility. This often entails situating seams beneath obstructing elements like arms, between limbs, or within recessed contours. Minimized island count focuses on streamlining the UV layout's complexity by keeping the number of islands to a minimum and reducing seam visibility. Uniform texture density is about maintaining a consistent texture density across the

model. This is achieved by judiciously placing seams around regions with geometric intricacies, such as twists, curves, and bulges. Ensuring uniform texture density prevents perceptible variations in texture quality and removes stretching. By diligently adhering to these principles during the seam assignment process, I strive to craft an optimized UV layout that facilitates impeccable texture mapping while mitigating any visible irregularities or anomalies in the final rendering.

After assigning seams, the next step is to *Unwrap*. This automatic operation involves separating the mesh into different islands, and then unfolding those islands into a flat plane, spreading them across a 2D square, the *UV map*. With this operation, the visible texture on the model gets applied. It usually takes a back-and-forth between assigning seams and unwrapping to optimize the *UV seams'* positions. Afterward, the final step is to optimize each island's position, scale, and rotation on the *UV map*. Manually scaling each island is integral to ensuring the desired texture density across the *UV map*. However, when it comes to optimizing the placement, rotation, and utilization of available UV space, relying on the computational power of specialized software proves to be more efficient. *Blender's Unwrap* operation provides an automated solution for UV packing and can deliver acceptable results for creating a model with good texture quality. Howbeit, an optional but even more practical approach, which I personally employ, involves using a paid add-on known as *UVPackMaster*³⁶⁰. This sophisticated tool harnesses heuristic functions and intelligent trial-and-error algorithms. It actively seeks optimal island placement within the *UV map*, maximizing the occupied area. This advanced approach enhances the perceived texture resolution of the model, all without necessitating an increase in the image's actual resolution. The result is a UV layout that optimally utilizes the available texture space, ensuring high-quality textures on the final model and a model ready for the next phase.

Texture Creation

Following the preparations involving the model's topology, an optimized UV layout, and properly assigned materials, the subsequent phase revolves around creating the essential textures and giving the model color. During this transition, a fundamental shift in software takes place, moving from *Blender* to *Substance Painter*. Consequently, steps must be taken, beginning with exporting models from *Blender* and their subsequent import into *Substance Painter*. The initial action in this phase involves creating a new file within *Substance Painter*

⁶⁰ glukoz. (2022, January 6). UVPackMaster for Blender. UVPackmaster 3. <https://uvpackmaster.com/>

and selecting the exported file of the retopologized model, most commonly in the *FBX file format*. Simultaneously, it is crucial to configure the texture quality settings to a high resolution, which I opt for a resolution of 4096, setting the stage for creating textures with exceptional detail and clarity.




Figure 18a. Santos, D. (2023). Effects of baked textures on the retopologized model. Retrieved September 12, 2023.

Before adding color and lifelike details to the model, transferring all the sculpted details to the optimized model is essential. This step significantly enhances the model's visual intricacy without sacrificing its runtime performance. This process involves creating a new FBX model file exclusively for the high-poly sculpts. This high-poly model is precisely positioned at the exact spatial coordinates as the optimized low-poly version, allowing *Substance Painter's* advanced engine to seamlessly extract data from the high-poly sculpt and apply it to the optimized model, akin to an artist's brushstroke. This process generates several essential images to create lifelike materials and realistic textures. The first is the *Normal* map, a colored image that guides how light interacts with the model by indicating the direction in which light should bounce off the surface, enhancing the perception of three-dimensionality. Next is the *Ambient Occlusion* map, presented in grayscale. This map represents soft shadows formed in recesses and crevices, contributing to the model's depth and realism. The *Curvature* map, another grayscale image, showcases the curvature of the high-poly sculpt projected onto the low-poly model, denoting lightness to convex shapes and darkness to concave shapes.

The *Thickness* map, a greyscale map, communicates the thickness or depth below the surface, aiding in creating convincing material properties. Finally, the *Position* map, a colored utility image, assists in positioning specific features in three-dimensional space. Each point on the model's surface is assigned a unique gradient color, facilitating precise placement and orientation of elements.




Figure 18b. Santos, D. (2023). Various baked maps generated by Substance Painter. Retrieved September 13, 2023.

After transferring the sculpt's intricate details to the model, the subsequent stage involves creating materials that leverage these details. Creating realistic textures for a model extends beyond simple surface coloration. In real life, objects possess material properties that significantly influence their interaction with light. To replicate this natural behavior and imbue the model with a lifelike appearance under various lighting conditions, a principle known as *Physically-Based Rendering*⁶¹ (PBR) is employed. This approach seeks to faithfully recreate the behavior of real-world materials, enhancing the model's visual fidelity.

Physically-based rendering relies on various technical parameters linked to diverse material characteristics. For instance, the *smoothness* parameter, which spans from 0 to 1, indicates a material's 'shininess.' A value of zero corresponds to a rough material akin to rubber, while a value of one signifies an exceptionally smooth material reminiscent of

⁶¹ A comprehensive website explaining Physically-based rendering. Pharr, M., Jakob, W., & Humphreys, G. (2018, October 15). Physically based rendering: From theory to implementation. *Physically Based Rendering: From Theory to Implementation*. <https://www.pbr-book.org/>

porcelain. The continuous nature of this parameter, expressed as a *float*⁶² value, enables the selection of any point within the 0 to 1 range, accommodating an array of material choices. Moreover, intricate details can gradually manifest on the model's surface by translating these chosen values into pixels within a 2D image. For instance, a smooth surface characterized by a high smoothness value can coexist with a layer depicting dust and scratches, corresponding to a low smoothness value. This dynamic interplay facilitates the creation of nuanced and realistic material intricacies.

Similarly to a game engine, *Substance Painter* employs a rendering engine capable of dynamically translating these greyscale PBR values into a range of material parameters visible in real-time. However, it is essential to consider that not all of these parameters find direct application in *Unity*, the game engine powering *VRChat*. Therefore, the focus should gravitate toward parameters that align with *Unity*'s requirements. One such crucial parameter is the *Base Color*, which essentially provides the model with its primary color. *Opacity* also plays an integral role, residing in the alpha channel of the *Base Color* texture and regulating transparency and translucency. Additionally, the *Normal Map* shares similarities with the previously mentioned map of the same name, but bears added intricacies from materials. *Ambient Occlusion*, a distinct grayscale texture, ushers in shadowing effects, darkening the primary color and reducing reflective and specular properties. The *Emissive* texture is worth noting, denoting the degree of self-glow and contributing to the model's radiance, particularly in dimly lit environments. *Metallic*, a grayscale texture, signifies whether a material exhibits metallic attributes and will be converted into the Red Channel within the MetallicSmoothness map. Lastly, *Smoothness*, as previously mentioned, another grayscale texture, conveys shininess and will be converted to the Green channel in the MetallicSmoothness map.

⁶² A website explaining the concept of parameters and, more specifically, the float parameter. The explanation is given within the python programming language, but the concept applies to computer graphics as well. Ravikiran, A. S. (2023, May 18). Understanding float in python [with examples]. Simplilearn.com. <https://www.simplilearn.com/tutorials/python-tutorial/float-in-python>



Figure 18c. Santos, D. (2023). Different textures from a previous model, exported from Substance Painter. Retrieved September 18, 2023.

This collection of textures is utilized within *Poiyomi's Toon Shader*⁶³, an open-source shader designed for *Unity*. This shader serves as the computational powerhouse behind the scenes, converting standard PBR textures into a comprehensive framework that enhances the object's realism. It orchestrates light reflections from the environment with remarkable finesse, resulting in many realistic lighting effects. The decision to employ this shader for the avatar stems from several compelling reasons. Firstly, its open-source nature encourages a collaborative and adaptable development community. Secondly, its code is optimized for efficient performance within VR, ensuring the avatar functions smoothly in *Unity*. Most importantly, its user-friendly interface and extensive array of customizable options enable the creation of an extensive range of visual styles, making it the ideal choice for achieving the desired aesthetics.

Armed with an understanding of how *Substance Painter* translates sculpt data into various maps and the intricacies of *Physically-Based Rendering* materials, the next pivotal step unfolds: crafting materials tailored for the model. This endeavor breathes life into the form, infusing it with texture, color, and, in a meaningful sense, a vibrant semblance of existence. The artistry in *Substance Painter* parallels painting color onto a tangible sculpture, albeit with a distinct divergence: here, the artist grapples with materials rather than mere paint. The process commences by delving into *Substance Painter's* extensive array of

⁶³ Poiyomi. (2023, January 15). Poiyomi shaders. Poiyomi Shaders RSS. <https://www.poiyomi.com/>

materials and applying a chosen one onto the model, draping it in its unique texture. The crux of material authenticity in *Substance Painter* lies in the adept utilization of the previously created baked *mesh maps*. Take, for instance, crafting a material designed to accumulate dirt in crevices and recesses. Leveraging the *Curvature* mesh map, the software pinpoints concave curves on the model and automatically applies a dirt texture to these specific areas, simulating the organic accrual of grime. To illustrate further, using human skin as an example, it is a canvas for nuanced hues that bestow lifelike qualities. Imagine bestowing a subtle rosy tint to delicate protrusions like the nose and ears, an artful touch that contributes to realism. This effect can be masterfully achieved by employing the *Thickness* mesh map as a guiding mask, permitting the precise application of a pinkish hue within its contours. The outcome transcends mere flat coloring, infusing the model with a heightened sense of authenticity. Harnessing a multitude of such intricacies, artists can craft materials within *Substance Painter* that authentically mirror their real-world counterparts or conjure entirely novel substances that adhere to the same visual principles as their tangible counterparts.



Figure 18d. Santos, D. (2023). Multiple views of the final Avatar materials within Substance Painter. Retrieved September 18, 2023.

Nevertheless, akin to real-life scenarios, an avatar typically comprises multiple materials in different zones, and this is where the utility of masks in *Substance Painter* becomes evident. These masks are instrumental in designating precise regions where a particular material becomes visible. Like the masks automatically generated within materials

via various mesh maps, artists can leverage this ability to manually create masks by drawing them directly onto the model. This process can be aptly characterized as 'painting a material onto the model,' as artists witness the material materializing in real-time as their brush strokes traverse the model's surface. As such, during this phase, the avatar creation process most assimilates to traditional painting.

Once the artist has achieved satisfaction with the materials they have carefully created and their precise application to specific areas of the model using masks, the subsequent phase revolves around preparing these materials for export. *Substance Painter* provides a range of templates tailored for exporting textures in various formats to suit different purposes. However, it is worth noting that the specific template required for *Poiyomi's Toon shader* is not pre-installed in the software. This essential template can readily be located on the official website⁶⁴, ensuring compatibility with the shader's unique requirements. Upon importing this template, the artist can export their carefully designed materials into a collection of diverse textures, each serving distinct functions in the rendering pipeline of 3D models. Once these textures have been successfully exported, the next destination is *Blender*, where they will aid in the subsequent stages of the avatar creation process. However, it is paramount to remember that these textures were exported using *Poiyomi's Toon Shader's* specific export preset. Consequently, one must exercise diligence in ensuring that each texture's RGB channels are correctly assigned to their respective PBR parameters. This assignment process guarantees that the materials behave as intended when incorporated into the shader system, resulting in a visually compelling and correctly lit 3D model. In *Blender*, this process is currently done using the *Material Node* system with the *Principled BSDF*⁶⁵ node, a graphical interface that displays each parameter of the PBR material.

Rigging and Weight-Painting

Following the application of textures onto the model within *Blender*, with each texture correctly assigned to its corresponding material, the creative path advances to the pivotal rigging phase. Rigging is the art of imbuing the model with the ability to move dynamically and convincingly, replicating natural and lifelike motion. This intricate process hinges on developing a mathematical skeleton (an *Armature*), a digital framework composed of

⁶⁴ Poiyomi. (2023b, January 15). Substance Painter: Poiyomi Shaders. Poiyomi Shaders RSS. <https://www.poiyomi.com/general/substance-painter>

⁶⁵ Blender Foundation. (2023, September 19). Principled BSDF. Principled BSDF - Blender Manual. https://docs.blender.org/manual/en/latest/render/shader_nodes/shader/principled.html

infinitesimal points in space intricately interconnected. These virtual bones influence the model's vertices, orchestrating their movement in unison as the character is animated. It is essential to adhere to a structured approach to ensure that the rig faithfully emulates the character's anatomical structure. Imagine an ethereal skeleton residing within the character and visualize the placement of bones at each pivotal joint where movement occurs. This visualization technique is a valuable guide, facilitating the strategic positioning of bones to align with the character's underlying anatomy. In essence, the rigging process gives life to the character, enabling it to articulate and express itself realistically in response to animations and interactions.

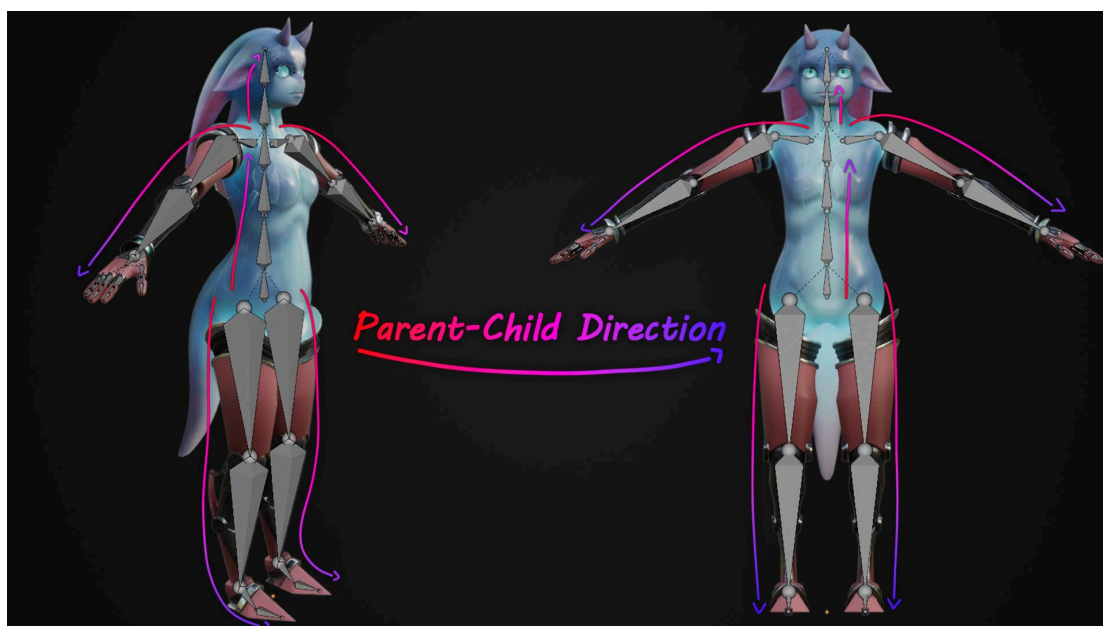


Figure 19a. Santos, D. (2023). The proper parent-child bone direction. Retrieved September 19, 2023.

Each bone within the rig operates as either a parent or child in relation to other bones. This hierarchical arrangement, where bones exhibit parent-child relationships, is fundamental to the dynamic articulation of the character. When a bone assumes the role of a child, it becomes intrinsically tied to the translation and rotation of its parent bone. This connection means that the child bone moves and rotates in harmony with its parent, seamlessly mirroring its movements. Conversely, when a bone serves as a parent, it exercises influence over the position and rotation of its respective child bones. This hierarchical structure generates a cascading effect, where each parent bone imparts its motion and orientation to its offspring. In most cases, bones within the rig possess both parent and child relationships, resulting in a complex web of interdependencies. Each bone is influenced by its parent and, in turn, exerts

control over the position and rotation of its own children bones. However, it is vital to establish these parent-child connections in a specific way during the Armature creation process. *Unity*, the game engine utilized for *VRChat*, imposes specific requirements on the bone order within the hierarchy. The initial bone, known as the hip, serves as the root from which the hierarchy branches into three primary divisions: the spine, the right leg, and the left leg. Each of these divisions corresponds to its respective chain of connected bones, ensuring a structured, standardized arrangement that aligns with *Unity's* expectations. This organization lays the foundation for seamless character animation and representation of the user's body movements in the virtual environment. Once the bones have been thoughtfully positioned within the avatar's body, the subsequent step involves *weight-painting* the model.




Figure 19b. Santos, D. (2023). The shoulder bone's zone of influence on the model. Retrieved September 19, 2023.

The *Weight-Painting* phase ensures the model's vertices respond correctly to the rig's movements. During this stage, specific areas of influence, usually called *Vertex Groups*⁶⁶, are assigned to individual bones. These zones then effectively dictate how the vertices move in response to the bone's motion. Moreover, each zone, akin to a mask, can exhibit a gradient of influence (*Figure 19b*) and the shoulder bone's *Vertex Group*. Typically, this gradient is visually represented by a spectrum of colors transitioning from indigo to blue, cyan, green,

⁶⁶ An official Blender video discussing the principles behind Vertex Groups. Blender Foundation. (2019, August 24). Vertex groups - blender 2.80 fundamentals. YouTube. <https://www.youtube.com/watch?v=dKZrzG5r13g>

yellow, and finally to red, signifying the degree of influence exerted. The artist can, in turn, paint this influence across the model's surface.

However, undertaking the laborious process of manually weight-painting each individual vertex group from the ground up can be exceptionally time-intensive. Fortunately, many 3D modeling applications come equipped with an automated feature that streamlines this task by assigning vertex groups based on the expected positions of bone articulations. This predictive approach hinges on the logical placement of articulation points, a generally straightforward concept. In Blender, for instance, this can be executed by selecting the model and its corresponding armature, finally utilizing the "*Armature Deform with Automatic Weights*" function. This initial step serves as the foundational scaffold for weight-painting. Following this automated allocation, the artist can transition into *Pose Mode*, wherein the bones are manipulated to observe how they influence the model. However, this automation is not an end-all solution. It merely lays the groundwork, and a manual adjustment process becomes necessary to achieve realism in the model's movements. During this phase, each vertex group undergoes deliberate fine-tuning. The objective is to ensure that the model undergoes deformation that closely mimics real-world movements when subjected to dynamic test poses.

Following the manual weight-painting of the fundamental humanoid bones, the artist emerges with a model poised for dynamic movement, responding faithfully to animations that dictate the articulations of these predefined humanoid bones. The subsequent phase entails transferring the model from *Blender* to *Unity*, which demands careful attention to specific export settings. These settings play a pivotal role in standardizing and harmonizing the values between *Blender* and *Unity*, ensuring that the model and its orientation maintain uniformity when viewed in both applications. The most crucial adjustment pertains to the '*Apply Scalings*' option, which necessitates a change from '*All Local*' to '*FBX All*.' This alteration is instrumental in avoiding scale discrepancies between *Blender* and *Unity*, thus facilitating a seamless transition and alignment of the model across both platforms.

Unity Setup

Upon completing the initial phase of avatar development, the subsequent task is to seamlessly integrate it into the game engine, *Unity*, and, by extension, the *VRChat* platform. This entails a series of essential steps: importing the previously created textures produced in *Substance Painter*, bringing in the FBX model from *Blender*, and creating a *Unity* project that

neatly incorporates all these vital components. Nonetheless, specific preparatory measures must be undertaken before the avatar is ready for testing and eventual use within *VRChat*. These preliminary tasks encompass the installation of the appropriate *Unity* version, *Poiyomi's Toon Shader*, and *VRChat's Unity SDK*. Given the rapidly evolving nature of both the game and engine, an in-depth exploration of the specific installation procedures for these components is beyond the scope of this dissertation. Instead, I will redirect to the relevant instructions,⁶⁷ recognizing the dynamic and frequently changing landscape of *Unity's* and *VRChat's* SDK development.

Once the *Unity* project has been set up and all the necessary files are imported and in place, the process of adjusting each asset's settings to align with its appearance in other programs begins. This procedure ensures that the avatar maintains its visual fidelity across different platforms. Starting with the textures, the initial step involves matching each texture's resolution to its native setting, which, for this project, should be set to 4096 pixels. However, it is essential to note that as part of the optimization process, these resolutions will be manually reduced later to minimize the processing power required. Subsequently, depending on its purpose, each specific texture must be configured to use either the *sRGB* or *linear* color space. The *sRGB* color space should be applied to color textures, such as *BaseColor* and *Emission* maps, which require *gamma-correction* for accurate color representation.⁶⁸ In contrast, the linear color space is the appropriate choice for data textures, like the *MetallicSmoothness* and *AmbientOcclusion* maps, as it ensures proper handling of their information. This alignment of settings ensures that the textures' visual integrity is maintained throughout their transition from program to program. Moreover, the *Normal* map should be set to use the *Normal* map standard. Afterward, a similar process must be undertaken with the model file, albeit with different settings. The model file necessitates the configuration of its rig to '*Humanoid*.' This designation informs *Unity* that the model incorporates an armature with standard bones that replicate the structure of a human skeleton, enabling the utilization of humanoid animations and *VRChat's Inverse Kinematics*,⁶⁹ a system designed to predict a chain of bones' positions based on the location of two end-points. This stage also demands scrutiny to ensure all bones are correctly assigned to their corresponding counterparts within the *Unity* humanoid rig system.

⁶⁷ VRChat Inc. (2023, June 12). Setting up the SDK: Vrchat creation. VRChat Creation RSS. <https://creators.vrchat.com/sdk>

⁶⁸ Vries, J. de. (2015). Gamma correction. LearnOpenGL. <https://learnopengl.com/Advanced-Lighting/Gamma-Correction>

⁶⁹ The Mathworks, Inc. (2020). What Is Inverse Kinematics? - MATLAB & Simulink. <https://www.mathworks.com/discovery/inverse-kinematics.html>




Figure 20a. Santos, D. (2023). The Avatar's appearance in the Unity editor using Poiyomi's Toon Shader. Retrieved September 20, 2023.

With the assets prepared, it is time to construct a Scene that will encompass the avatar and its components, ultimately slated for *VRChat* integration. Commencing by importing the model file into the Scene hierarchy, the subsequent step entails the creation of materials and their application to the avatar. This phase imbues the model with color, texture, and the desired appearance akin to that created in *Substance Painter*. This step initiates with creating a material, which is subsequently configured to employ the previously mentioned *Poiyomi's Toon Shader*. Each texture is then assigned to its corresponding slot within the material. For instance, the *BaseColor* texture finds its place in the Main Color slot, while the *Normal* map takes its designated position, and so forth. At this juncture, the avatar for this project comprises two distinct materials previously planned in the “*UV Editing*” chapter (pp.71-74), one tailored for the organic core and another for the mechanical components.

After finalizing the avatar's appearance within the game engine, the next phase involves incorporating *VRChat's Avatar Descriptor* component into the avatar and customizing its attributes. These adjustments encompass configuring the viewpoint to align seamlessly with the avatar's perspective (typically situated between the eyes), defining the maximum rotation values for the eye bones, implementing colliders for different body parts, and more. Ultimately, the artist can finally upload a test version of the avatar. This marks the

final preparatory step before the artist can don their initial version of the avatar within the game.

Following this, a simplified iteration phase commences. During this step, the artist assumes control of their creation in-game and documents any imperfections or issues that manifest. These issues can encompass various aspects and are best addressed through iterative refinement. For instance, if the avatar exhibits incorrect proportions when tested with full-body tracking, the rig's proportions must be adjusted in *Blender*. Accurate armature proportions can be achieved through reference materials, such as an existing model with ideal humanoid proportions or real-life images of the human body matching the avatar's pose. Subsequently, the revised file is imported back into the Unity project, effectively replacing the original one. Furthermore, any discrepancies stemming from inadequate weight-painting become apparent when embodying and manipulating the avatar. The conclusion of this stage signifies the completion of a fundamental and conventional avatar. Nevertheless, considering this project's specialized avatar is intended to incorporate additional features to enhance expressiveness and immersion, the journey is far from its conclusion.

Face-Tracking

The initial specialized phase of this project involves implementing face-tracking support for the avatar. As previously mentioned, the primary objective of integrating face-tracking into an avatar is to enhance the precision of emotions conveyed through the user's facial expressions. Consequently, this significantly augments the avatar's expressiveness, as it leverages the face, arguably the most important tool employed by humans during social interactions, to bridge information from the real world into the virtual one.

However, another system must be employed due to the limitations of current face-tracking technologies. While impressive in their ability to capture facial expressions, these technologies cannot yet track the intricate movements of the mouth well enough as they articulate words during speech. To address this, a distinct system comes into play, one that operates in tandem with a microphone input. This system focuses exclusively on animating the lower face of the avatar, primarily responding to the user's spoken words. This alternative system relies on a sophisticated algorithm designed to identify and interpret specific syllables, known as '*visemes*.' These *visemes* correspond to phonemes and are the building

blocks for replicating speech-related mouth movements in the avatar. When the algorithm detects these *visemes* in the user's speech, it transmits this information to an animator system. This animator system, in turn, activates the appropriate *blendshapes* that deform the avatar's mouth region over time. Through this careful blending, the avatar's mouth mimics the user's speech patterns, creating a natural and synchronized mouth movement that enhances the overall realism of interactions. The initial selection of *visemes* was drawn from a comprehensive paper by *Visage Technologies*, the *Viseme MPEG-4 Standard*.⁷⁰ However, recognizing the complexity of this task, *Oculus* (a company later acquired and merged into *Meta*) took the initiative to curate a set of 15 key *visemes*. *Oculus* then developed a specialized algorithm, known as *Oculus Lipsync*⁷¹, capable of reliably detecting and interpreting these specific phonemes.



Figure 21a. Santos, D. (2023). Example visemes on the Avatar. Retrieved September 22, 2023.

VRChat subsequently integrated this system as a default feature for avatars, simplifying the process for artists. To implement this system successfully, artists must create the corresponding 15 *blendshapes* for their avatars, each carefully named to correspond to the specific *viseme* it represents. These *visemes* should ideally mirror human mouth movements associated with common speech sounds, ensuring that users can readily recognize and relate to the avatar's expressions. However, based on my personal experience, it is advisable to

⁷⁰ Visage Technologies. (2012, August). MPEG-4 Face and Body Animation (MPEG-4 FBA). <https://visagetechnologies.com/uploads/2012/08/MPEG-4FBAOverview.pdf>

⁷¹ Meta Platforms, Inc. (2018). Viseme reference: Unity: Oculus developers. Developer Center. <https://developer.oculus.com/documentation/unity/audio-ovrlipsync-viseme-reference/>

exercise caution with the intensity of each *viseme*. Overly exaggerated *visemes* can significantly distort the avatar's facial appearance compared to the user's, potentially disrupting the sense of immersion, mainly when used in conjunction with face-tracking. After finalizing the *visemes*, the subsequent step involves testing them in the in-game environment and fine-tuning them if needed. However, given their relatively straightforward creation, the artist may streamline this process and proceed directly to crafting the face-tracking *blendshapes*. After completing both the lip sync *visemes* and the face-tracking *blendshapes*, comprehensive testing ensures seamless integration of these crucial features into the avatar's expressions and speech replication.

Before delving into creating the face-tracking *blendshapes*, selecting a face-tracking standard that aligns with the project's objectives is essential. Several face-tracking standards cater to distinct purposes, each with its unique features. For this project, the chosen standard is a slightly modified version of *ARKit* by *Apple*, a decision motivated by several compelling reasons. First and foremost, *ARKit* enjoys widespread popularity within VR applications. This popularity translates into broad device support, making it accessible to diverse users. Additionally, *ARKit* exhibits significant overlap in *blendshapes* with other face-tracking standards. This overlapping characteristic is pivotal, as it facilitates the development of a face-tracking-supported avatar compatible with various face-tracking devices. Furthermore, from a personal standpoint, *ARKit* balances complexity and manageability, offering a robust framework for accurately reproducing facial movements while maintaining a reasonable number of *blendshapes*. This equilibrium ensures the avatar can faithfully replicate nuanced expressions without an excessively intricate *blendshape* system.



Figure 21b. Santos, D. (2023). Multiple ARKit blendshapes on the Avatar. Retrieved September 23, 2023.

When creating the face-tracking *blendshapes*, the artist should initiate the process by closely examining the chosen standard. It is imperative to gain a comprehensive understanding of which blend shapes correspond to specific muscle groups responsible for facial expressions. Carrying out this process can be achieved through two primary methods. The first involves referencing various image-based websites that provide detailed facial anatomy and muscle movement references.⁷² The second, which I personally prefer, entails studying a pre-existing model that has already been equipped with the necessary *blendshapes*. This hands-on approach allows for a more tangible and practical understanding of how these *blendshapes* interact and influence the model's facial expressions. Once this understanding is in place, the task at hand is to represent these muscle movements on the avatar adeptly. *Blendshapes* function additively, meaning multiple *blendshapes* can simultaneously influence the model's appearance (*Figure 21b*). This inherent feature grants the ability to exhibit a wide range of facial expressions by enabling different *blendshapes* concurrently. Moreover, the flexibility of *blendshapes* extends to their activation percentages. This nuanced control allows for smooth transitions between *blendshapes* or their gradual melding at various values, thereby facilitating the creation of entirely novel expressions and allowing for accurate face-tracking.

Following the modification of the avatar with the newly incorporated *blendshapes*, the subsequent phase involves the reintegration of the avatar into *Unity*. This process entails

⁷² Arkit blendshapes. ARKit Blendshapes. (2021). <https://arkit-face-blendshapes.com/>

replacing the previous model file with the updated version and subsequently integrating it into the *VRChat* environment. Notably, *VRChat* currently lacks native support for face-tracking, prompting a proactive part of the *VRChat* community, collectively known as *VRCFaceTracking*⁷³ (VRCFT), to start an initiative to implement face-tracking capabilities within the platform. This implementation harnesses the power of *VRChat* implementation of *OpenSoundControl*⁷⁴ (OSC), a standardized protocol adapted to facilitate communication between various software components seamlessly. The resourceful VRCFT community has developed a dedicated program that adeptly parses data from diverse VR headsets. This program then standardizes this data and channels it to *VRChat* through the OSC framework. Intriguingly, avatars can interpret this OSC data to influence many parameters within the avatar itself dynamically. This ingenious approach effectively enables avatars to incorporate highly responsive face-tracking functionality compatible with a broad spectrum of face-tracking VR headsets, due to the overlapping *blendshapes* being standardized through the external program.

Nonetheless, implementing the intricate logic required for face-tracking is a task that demands a nuanced understanding of the data transmitted by the VRCFT program. To enable this functionality, the artist must navigate this data and use it to construct an *Animator Controller*⁷⁵, a *Unity* script that operates based on predefined conditions to trigger animations on the avatar. Although it is possible to build such an *Animator Controller* from scratch, there are readily available templates tailored to address this specific challenge. Given the dynamic nature of this face-tracking implementation, it is essential to note that the precise steps to implement it may evolve over time. At the time of writing, one such template I intend to use for this project is part of *Adjerry's VRCFaceTracking-Templates*.⁷⁶ These templates come complete with detailed instructions. Following these instructions (at the time of writing, involving actions like importing a program such as *VRCFury*⁷⁷ and attaching the *ARKit* face-tracking package to the avatar within the *Unity* hierarchy), the avatar then becomes primed for in-game testing. When executed correctly, the avatar should mirror the user's facial movements, imbuing it with a heightened sense of realism and expressiveness.

⁷³ James, B. (2021). Docs: Vrcfacetracking. VRCFaceTracking RSS. <https://docs.vrcft.io/>

⁷⁴ Wright, M., & Freed, A. (2021, August 13). OpenSoundControl.org. OSC index. <https://opensoundcontrol.stanford.edu/index.html>

⁷⁵ Unity Technologies. (2023a, September 22). Animator controller. <https://docs.unity3d.com/Manual/class-AnimatorController.html>

⁷⁶ Adjerry91. (2021). Adjerry91/vrcfacetracking-templates: VRC face tracking templates. GitHub. <https://github.com/Adjerry91/VRCFaceTracking-Templates>

⁷⁷ VRCFury Devs. (2023). Home. VRCFury. <https://vrcfury.com/>

Iteration

The final stage in this avatar creation process is iteration. In this critical phase, the artist rigorously evaluates the avatar's performance in-game, seeking out any missing features and flaws that need addressing. Despite its appearance as a potentially swift final step, this phase is frequently quite the opposite, primarily due to the complex and multifaceted nature of building an avatar from the ground up. Effective testing must encompass various scenarios and situations to ensure the avatar functions seamlessly. This extensive testing process can be time-consuming as it involves putting the avatar through its paces in numerous contexts. However, it is essential to strike a balance and avoid falling into the trap of perfectionism. While attention to detail is crucial, excessive tinkering can lead to diminishing returns, consuming valuable time without yielding substantial improvements.



Figure 22a. Santos, D. (2023). A new optional helmet created for the avatar. Retrieved September 25, 2023.

This stage of the process is highly subjective and relies on the artist's values regarding the ideal avatar. For this project, I prioritized a distinct appearance, experimental elements, and, notably, a swift and efficient means of conveying expressiveness to enrich social interactions. Consequently, the primary emphasis during the design process was creating an avatar embodying these values. However, a crucial lesson I have learned from selling avatars is the importance of valuing others' opinions and feedback. In the unique context of *VRChat*, this involves temporarily loaning the avatar to different users and encouraging them to share their impressions and suggestions. Including these diverse perspectives does not necessarily

diminish the value of prior research conducted before avatar creation; it can enrich it by incorporating additional opinions and external viewpoints. While external influences can potentially detract from the project's original message and intent, the crucial aspect is to be selective when considering feedback. That is, to prioritize feedback that aligns with the original vision and stays within the project's goals. This feedback can even originate from a broad spectrum of users, from those new to VR to those who have spent considerable time in virtual environments. The broad spectrum of testers helps identify issues and preferences that might have been overlooked. In the case of the avatar developed for this project, initial feedback revealed that some users found its original facial features slightly unsettling or 'uncanny.' This valuable insight prompted a thoughtful retexturing and slight resculpting of the facial elements. An optional helmet piece was also introduced, complete with integrated face-tracking support. This addition aimed to cater to a broader range of user tastes and perspectives and allowed for more customization in expressivity.

In addition to the feedback from friends, personally and through testing, there was a recurring sentiment that the avatar felt 'too stiff' despite its dynamic and fluid appearance. This perception was primarily attributed to certain elements, such as the ears and tail, being physically unresponsive to movement. In response to this feedback, several measures were taken to enhance the avatar's realism and fluidity. The first step involved rigging the ears and tail with additional bones. These supplementary bones were introduced to enable more precise control over these specific elements. To then enable lifelike movement, *VRCPHysBone* components⁷⁸ were applied to these bones. These *VRChat Unity* components, akin to the *Avatar Descriptor*, endow real-time physics calculations to the designated bones. Combined with proper weight-painting on the mesh, this integration allowed for natural secondary movement in response to the avatar's motions.

⁷⁸ VRChat Inc. (2023, May 19). Physbones: Vrchat Creation. VRChat Creation RSS. <https://creators.vrchat.com/avatars/avatar-dynamics/physbones>



Figure 22b. Santos, D. (2023). Morph Heartbeat Animation. YouTube. Retrieved September 25, 2023, from <https://youtu.be/o26kNEotEdk>

Moreover, reflecting on the feedback regarding the body's stiffness sparked an idea. I conceived a heartbeat animation designed to mimic the rhythmic pulsing of an artificial heart. This addition accentuates the main body's gel-like substance and demonstrates its artificial life support. This additional animation layer introduced a distinctive identity. It heightened the believability of the avatar, reinforcing the impression that it was 'alive' and adhered to its internal logic (as demonstrated in *Figure22b*). The process commenced with the creation of two *blendshapes* in *Blender*. These *blendshapes* were designed to represent the heart's behavior: one *blendshape* reflected the expansion of one side of the heart and the contraction of the other, and the other *blendshape* mirrored the effect. Since the organic body is a separate mesh from the robotics and the heart, a complementary set of *blendshapes* was developed to influence the protrusion of the chest to ensure a seamless visual transition between the heart's pulsations and the avatar's external appearance. Upon importing the model, now enhanced with these *blendshapes*, into *Unity*, the next stage was the construction of the heartbeat animation. This animation was carefully calibrated to smoothly activate and deactivate the *blendshapes* in sequence, ensuring that the pulsating motion closely mimicked the rhythm of a genuine heartbeat. The heartbeat animation should be integrated into an *Animator Controller*, which should then be placed in the FX section of the *Avatar Descriptor*. In the case of this project's avatar, there was a preexisting FX controller in the *Avatar Descriptor*. Therefore, a new layer was created within it to house the heartbeat animation. When examined in-game, the result was a remarkably lifelike heartbeat that conferred

unprecedented dynamism and vitality to the avatar. This subtle yet impactful addition further solidified the avatar's illusion of life.



Figure 22c. Santos, D. (2023). Promotional material showcasing different body shapes available through blendshapes. VRA-05 “Morph” (VRChat Avatar Base). Gumroad. Retrieved September 25, 2023, from <https://gh0stt.gumroad.com/l/morph>

Furthermore, a final significant enhancement introduced to the avatar was the ability for users to modify the body's shape directly within the virtual world. This feature aimed to cater to users with diverse body types, enhancing the sense of presence and alignment with their virtual representation. This system's development paralleled the heartbeat system's creation in many ways, as it relied on custom *blendshapes* for its functionality. However, a crucial distinction lies in the system's dynamic nature, allowing users to alter their avatar's body type while in-game. The process involved the creation of custom *blendshapes* that represented various body types, ranging from slender to muscular or curvy, for instance. These *blendshapes* are then dynamically adjusted based on a menu accessible to the user within the virtual environment. In essence, this system empowers users to reshape their avatars according to their preferences during gameplay. An *Animator Controller* was employed to realize this functionality and different parameters and animations associated with each custom *blendshape* were established. Afterward, a *Direct Blend Tree*⁷⁹ enabled custom animations driven by these parameters. Subsequently, the parameters needed to be

⁷⁹ Unity Technologies. (2023, September 22). Blend trees. <https://docs.unity3d.com/Manual/class-BlendTree.html>

linked within the *VRCEXpressionParameters* Script, a tailored *VRChat* Script that synchronized *Animator Controller* parameters among different *VRChat* users. The final step in this process was the creation of a customized *VRCEXpressionMenu*,⁸⁰ yet another specialized script. This script detected the custom parameters from the Expression Parameters list and added radial menus to each parameter. These radial menus allow users to gradually modify these parameters directly within the virtual environment, allowing them to transform their avatar's body appearance at will and increasing immersion after doing so.

Finally, after various adjustments to the textures in *Substance Painter* and subtle refinements in proportions, the intricate process of creating this avatar draws to a conclusion, aligning with my exacting standards. The resulting avatar, *VRA-05 "Morph"*, carries a chosen name inspired by science fiction, evoking the essence of a fictional experiment and paralleling the more experimental nature of the creation of the avatar itself. This nomenclature harmonizes seamlessly with the avatar's underlying theme and conceptual identity. The moniker "*Morph*" itself is the primary name, derived from one of my cherished characters with the same name in *Treasure Planet*⁸¹, an animated masterpiece that has left an indelible mark on my childhood. However, the acronym '*VRA-05*' holds a hidden layer of significance that is not immediately apparent. In truth, it encapsulates a piece of my avatar-making history. '*VRA*' stands for Virtual Reality Avatar, while the number '*05*' signifies that this avatar marks the fifth avatar I have built and shared with others. It symbolizes the progression and evolution of my skills and the vibrant journey of creativity.

⁸⁰ Expression menu and controls: Vrchat creation. VRChat Creation RSS. (2023, June 20). <https://creators.vrchat.com/avatars/expression-menu-and-controls>

⁸¹ Musker, J., & Clements, R. (2002). *Treasure planet*. Walt Disney Pictures, Walt Disney Feature Animation.



Figure 22d. Santos, D. (2023). Final promotional video showcasing the avatar. VRA-05 “Morph” (VRChat Avatar Base). Youtube. Retrieved October 27, 2023, from https://youtu.be/2SMZrE_07FA

I showcase the avatar (*Figure 22d*) being embodied by several people and their experience with it. In it, the result of this extensive creation process can be finally seen. In crafting VRA-05 “Morph”, I ventured deep into the intricate art of avatar creation, traversing through planning, concepting, structuring, sculpting, retopologizing, UV unwrapping, sculpting, texturing, rigging, and face-tracking. Each step brought unique challenges and rewards, resulting in an avatar that captures the essence of expression and embodies the culmination of artistic exploration. As this chapter of the creative journey concludes, Morph serves as an example of the fusion of artistry, technology, and storytelling and the boundless potential of digital artistry within the immersive medium of virtual reality.

Closing Thoughts

Throughout the research and avatar creation process, my perspective on the potential of avatars and their role in human-computer interaction has undergone a meaningful transformation. Delving into the intricacies of various avatars and their creation has expanded my knowledge and significantly altered my perception of avatars' expressive and influential abilities. One of the most impactful revelations of this journey has been realizing the power inherent in face-tracking and full body-tracking technologies. These tools have illuminated the potential for avatars to become powerful conduits for human expression and embodiment in virtual environments. Witnessing how avatars can mirror users' nuanced movements and expressions has deepened my appreciation for how technology can bridge the gap between the physical and the digital. Furthermore, this research has instilled in me a heightened awareness of the user's experience and original self-expression modes within virtual spaces. Understanding the intricacies of how individuals interact with avatars has been very informative and has equipped me with valuable insights into enhancing my character creation and 3D modeling skills. In essence, this journey has been a transformative one, not only in terms of academic knowledge but also in personal growth. As I move forward, I carry with me a newfound appreciation for avatars' creative and expressive potential and a deeper understanding of their significance in the evolving landscape of virtual reality and social experiences within.

Findings, statements and implications

The culmination of this project has exposed several key findings. First and foremost, it has introduced me to a more efficient approach to avatar creation. The process was distilled to its core essentials, significantly reducing the time required. A stark example of this efficiency is the comparison between my previous avatar, *Zuri Sloth*, which took around one and a half months to complete, and the creation of *Morph* in just three weeks. This substantial time reduction emphasizes the success of simplifying and optimizing the avatar design process. Furthermore, meticulous planning emerged as a vital lesson. Strategic thinking and structured planning at every stage of the avatar's development proved essential in achieving effectiveness and alignment with the avatar's intended objectives. This transition to a more organized approach saved time and contributed significantly to the avatar's overall success. A substantial focus of this project was delving into various face-tracking standards. This

exploration deepened my understanding of these standards and equipped me with the knowledge needed to create convincing blend shapes that adhere to each unique standard. The outcome was an avatar that boasts the facial expressiveness aimed in the scope of this research, effectively bridging the digital and physical gap. This project expedited the avatar creation process, underscored the critical role of planning, and deepened my understanding of technology standards. The experience gained through this journey has significantly enhanced my expertise in avatar creation, propelling me into the realm of more effective and expressive virtual characters.

The core purpose of this project was multifaceted, driven by the ambition to achieve two primary objectives. Firstly, it aimed to facilitate advancing my skills in the field of avatar creation. This journey was not merely a creative endeavor but a dedicated effort to elevate my expertise to new heights. The significance of this mission extended beyond personal growth; it encompassed the aspiration to contribute to the broader discourse on avatars and their role in virtual spaces. The second pillar of this undertaking was equally significant: sharing my acquired knowledge into a form that would be accessible and comprehensible to a diverse audience. This commitment to clarity and accessibility stemmed from recognizing that avatars are integral components of digital spaces, shaping and reflecting the identities of their users. One must navigate a diverse landscape of appearances, behaviors, and expressive potential to truly comprehend the intricate relationship between users and their avatars. This project, therefore, explored how users harness avatars to express themselves in virtual realms. It delved deep into the myriad ways in which these digital representations extend and amplify the identities of their users. In doing so, it sought not only to unravel the complexities of user identity and social dynamics within virtual environments but also to process these insights into a form that would resonate with a wider audience.

This project has revealed the close association between avatars and identity, demonstrating how simple features like body shape toggles and hue sliders enable users to imbue their avatars with a distinct sense of self. It is common practice for users to select their preferred color or body type, reflecting the deep-rooted connection between avatars and personal identity. Furthermore, this journey has provided insights into the diverse ways individuals respond to the experience of embodying expressive avatars. While reactions may vary, there is a consistent thread of wonder and awe when individuals first engage with immersive technologies like face-tracking and body-tracking. This initial fascination underscores the transformative potential of avatars, offering users a portal to explore and express facets of themselves that extend beyond the constraints of the physical world.

Limitations

While this dissertation offers valuable insights into the creative process of crafting avatars, emphasizing expressiveness, it is important to acknowledge certain limitations. Firstly, the primary focus of this dissertation is on the technical and artistic aspects of avatar creation, with a secondary consideration of the psychological aspects related to prolonged embodiment. While it introduces these psychological elements, the exploration is not as comprehensive as possible. A more in-depth analysis of the psychological and sociological implications of avatar embodiment warrants further research and examination outside of my area of expertise. Secondly, this project is tailored towards avatars created with user expressiveness in mind. It may not fully accommodate avatars designed for other specific purposes. However, it is worth noting that overlapping methodologies and principles in avatar creation can be adapted and applied to diverse avatar objectives, allowing for extracting relevant insights and techniques.

This dissertation serves as a stepping stone toward a deeper understanding of avatar embodiment and its multifaceted dimensions. While it primarily explores the creative process and user expressiveness, there is ample room for future research to delve into the intricate interplay between avatars and psychology and tailor avatar creation to various objectives and applications.

What was, what is, and what will be?

The concept of avatars, once synonymous with anonymity and the freedom to explore an alternative persona, has evolved into a far more intricate phenomenon than initially conceived. In the past, avatars were perceived as users' extensions, offering a hidden, liberated identity shielded from the constraints of real-life struggles, akin to the concept portrayed in "Snow Crash." They were seen as the embodiment of an anonymous online persona, allowing individuals to express themselves more freely. However, the present-day understanding of avatars transcends mere anonymity. Extensive studies and real-world examples, such as the immersive social platform *VRChat*, have illuminated the complexity of avatar-user interactions. Avatars, in addition to affording a heightened sense of freedom, expressively influence users' behaviors and perceptions when they feel deeply engaged. This influence can transform personality, body perception, perceived social roles, and even gender identities. Contemporary virtual reality technologies, exemplified by platforms like *VRChat*, have garnered a dedicated and steadily growing user base. With VR technology becoming

more advanced, accessible, and comfortable, it is gradually approaching a point where it can convincingly simulate alternate realities. Much like the transformative role of smartphones in modern life, virtual reality is poised to assume a central position, potentially replacing many everyday functions and activities.

However, unlike the simplification of experiences associated with smartphones, virtual reality promises to augment and enhance sensory experiences. Today's impact on users' sense of self due to virtual reality environments and avatars will likely intensify as technology revolutionizes the human sensory experience. In many ways, it is shaping up to be the new reality for a significant portion of the population, fundamentally altering how individuals interact, perceive themselves, and engage with the world around them. As this technological journey unfolds, the influence and implications of avatars in shaping our virtual and, by extension, our real lives will naturally change.

Bibliography

Alexdottir, S., & Yang, X. (2022). Phantom touch phenomenon as a manifestation of the visual-auditory-tactile synaesthesia and its impact on the users in virtual reality. 2022 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct) pp. 727-732. <https://doi.org/10.1109/ismar-adjunct57072.2022.00218>

Barfield, W., & Furness, T. A. (1995). *Virtual Environments and advanced interface design*. Oxford University Press.

Belting, H., & Hansen, T. S. (2017). *Face and Mask: A double history*. Princeton University Press.

Braun, N., Berisha, A., Anders, D., Kannen, K., Lux, S., & Philipson, A. (2020). Experimental inducibility of supernumerary phantom limbs: A series of virtual reality experiments. *Frontiers in Virtual Reality*, 1. <https://doi.org/10.3389/frvir.2020.00012>

Edwards, G. J., Taylor, C. J., & Cootes, T. F. (1998). Interpreting face images using active appearance models. *Proceedings Third IEEE International Conference on Automatic Face and Gesture Recognition*. <https://doi.org/10.1109/afgr.1998.670965>

Ekman, P. (1972). *Universals and cultural differences in facial expressions of emotion*. University of Nebraska Press.

Gerhard, M., Moore, D., & Hobbs, D. (2004). Embodiment and copresence in collaborative interfaces. *International Journal of Human-Computer Studies*, 61(4), 453–480. <https://doi.org/10.1016/j.ijhcs.2003.12.014>

Lebowitz, J., & Klug, C. (2017). *Interactive storytelling for video games: A player-centered approach to creating memorable characters and stories*. Focal Press.

Mori, M., MacDorman, K., & Kageki, N. (2012). The uncanny valley [from the field]. *IEEE Robotics & Automation Magazine*. <https://doi.org/10.1109/mra.2012.2192811>

Nilsson, N. J. (2009). *The quest for artificial intelligence*. Cambridge University Press.

Parrinder, E. G. (1982). *Avatar and incarnation*. Oxford University Press.

Pilacinski, A., Metzler, M., & Klaes, C. (2023). Phantom Touch Illusion, an unexpected phenomenological effect of tactile gating in the absence of tactile stimulation. *Scientific Reports*, 13(1). <https://doi.org/10.1038/s41598-023-42683-0>

Poor Zamany Nejatkermany, M., Modirian, E., Soroush, M., Masoumi, M., & Hosseini, M. (2016). Phantom Limb Sensation (PLS) and Phantom Limb Pain (PLP) among Young Landmine Amputees. *Iranian journal of child neurology*, 10(3), 42-7.

Ravikumar, S. (2017). *Performance-driven facial animation with Blendshapes*. University of Bath.

Seegmiller, D. (2003). *Digital character design and painting*. Charles River Media.

Skarbez, R., Smith, M., & Whitton, M. C. (2021). Revisiting milgram and Kishino's reality-virtuality continuum. *Frontiers in Virtual Reality*, 2. <https://doi.org/10.3389/frvir.2021.647997>

Slater, M., & Sanchez-Vives, M. V. (2014). Transcending the self in immersive virtual reality. *Computer*, 47(7), 24–30. <https://doi.org/10.1109/mc.2014.198>

Stephenson, N. (2022). *Snow crash*. DEL REY. Initially published in 1992.

Thomson, T. J. & Greenwood, K. (2020). Profile pictures across platforms. *Handbook of Visual Communication*, 349–363. <https://doi.org/10.4324/9780429491115-32>.

Trappl, R., Payr, S. & Petta, P. (Eds.). (2003). *Creating emotional relationships with virtual characters. Emotions in Humans and Artifacts*. <https://doi.org/10.7551/mitpress/2705.003.0013>

Viola, P., & Jones, M. (2001). Rapid object detection using a boosted cascade of Simple features. Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001. <https://doi.org/10.1109/cvpr.2001.990517>

Williams, L. (1990). Performance-driven facial animation. Proceedings of the 17th Annual Conference on Computer Graphics and Interactive Techniques. <https://doi.org/10.1145/97879.97906>