Olivia Stiebel
Professor Evans
PAVS 4500
12-5-18

Digital Avatars and the Digital Afterlife

ABSTRACT

Love is eternal, or so they say. Science-fiction scenarios – such as that in the episode "Be Right Back" in the Netflix series "Black Mirror" – have addressed the idea of text or speech digital avatars with the ability to bring your deceased loved ones back to life through the digital realm. This paper is written under the assumption that digital avatars, which may take the form of text or voice, will reach a level in which they will be convincing and more widely available. Under the assumption that it will be possible in the near future to create text and speech digital avatars that can bring our loved ones back to life through the digital realm, this paper addresses the ethical questions surrounding this technology to determine how it should be used. We examine the implications surrounding such digital avatars, both for those who knew the deceased well and those who did not, and consider whether digital avatars effectively help cope with grief or cause emotional attachment or isolation issues. The business models behind companies in this realm and risks of manipulating vulnerable people are also considered. There are a whole host of issues and ethical questions surrounding similar digital avatars for the living, but this paper will only focus on the ethical implications of this type of Artificial Intelligence for the deceased. By exploring both science-fiction examples and real-world technology — either currently in development or already here — this paper examines the ethical questions surrounding digital avatars to raise awareness of the technology's ethical implications and determine how it should

be used, contending that while digital avatars could help cope with grief, they also have many ethical or less positive implications as well.

INTRODUCTION

Picture this: one day, you find out that your loved one has passed away. Your grief and sadness are unbearable. However, instead of having to say goodbye forever, there is now a way that you can bring your loved one back to life and still talk – virtually, either through text or through voice. As time passes, your grief may start to subside, and you can finally get some of the closure your heart has been aching for.

This may sound like something straight from a science-fiction movie, and the episode "Be Right Back" in the Netflix series "Black Mirror" provides us with such a scenario. In the episode, after the active social media user Ash dies, his girlfriend Martha uses an AI service that uses and learns from Ash's past social media activities and profiles, as well as the private media Martha uploads, in order to create a digital avatar that can communicate with the living. The technology is used in three stages – text, vocal, and physical digital avatar. The service Martha subscribes to first sends her text messages, then re-creates his voice to let her talk with the AI on the phone, and then finally – by paying for an upgraded version – lets her interact with an embodied agent. The stages progress in intensity each time and Martha isolates herself from friends and family to communicate with the AI, ultimately becoming frustrated by the ways – no matter how subtle they may seem – that the android is unlike Ash, eventually leading her to lock it away in her attic. The idea of an embodied agent and having an 'Ash in our attic' by implanting the deceased loved one's look and personality into a digital avatar AI is arguably

difficult to comprehend and fortunately not something that is going to occur in the near future. However, the idea of either a text or vocal digital avatar is incredibly plausible and much easier to comprehend, given the advances that have already taken place in AI, as well as the pervasiveness of social media in today's world.

Social media is ubiquitous in today's world. Our lives are digitized and made public at the tap of a button: our pictures, stories, and videos. Our thoughts on pancakes we ate for breakfast this morning, the yoga class we loved, or the newest Thomas Rhett album release get posted on our Instagram, Snapchat, Twitter, Facebook, or LinkedIn for the whole world to see. Because social media and technology is so pervasive in our lives – and so much of our lives are out there online and in the digital realm – there is both a physical and a digital side to us even when we are alive. At the same time, there are questions over what happens to these social media accounts when we die. Should the accounts of deceased users just be deleted? Should they get frozen? Or, should they be used to create a digital avatar for close relatives and loved ones to interact with after you are gone?

AI research into text and voice digital avatars is focused on the latter. While a scenario such as that in "Be Right Back" is fictional, it is not purely fictional. This type of technology – at least the text and voice – are already plausible and in development. Real-life, current technological research is already attempting to tackle this issue. With the loss of a loved one comes immense feelings of grief, sadness, and void: a longing to talk to or see the deceased again, even if just for one more time, and the feeling that your time with your loved one was cut short. And not only this, but there is no escaping death – at least not yet. Thus, it is not difficult to imagine that individuals would potentially jump at the opportunity to use a technology that

would help them cope with grief and sadness after the loss of loved ones. Nevertheless, there are important ethical and moral questions surrounding both text and vocal digital avatars, which must be addressed before this technology becomes available or pervasive.

SYMPTOMS OF GRIEF

Although death is a natural part of life, individuals often experience grief through a variety of emotions after the loss of a loved one. While everyone grieves differently, doctors have identified – and are often in agreement on – five common stages of grief: denial, anger, bargaining, depression, and acceptance (Mulder, "What Is Normal Grieving, and What Are the Stages of Grief?"). During stage 1, the individual has trouble coping with the news of the loss and may naturally think that what is happening in reality is not actually happening. Denial may incorporate feelings of shock and numbness, and it represents the body's way of creating a temporary defense mechanism from the pain and overwhelming emotions after losing a loved one. During stage 2, reality and the pain of the loss start to set in, which may lead the individual to feel helpless or upset. For many individuals, these feelings often turn into anger. Anger may be directed at others, at the loved one who is now gone and left the individual alone, or at life in general. During stage 3, the individual dwells on anything that could have been done to prevent the loss of the deceased and focuses on "what if" statements to question the loss. During stage 4, the individual's true sadness and grief has set in as the individual realizes and understands how

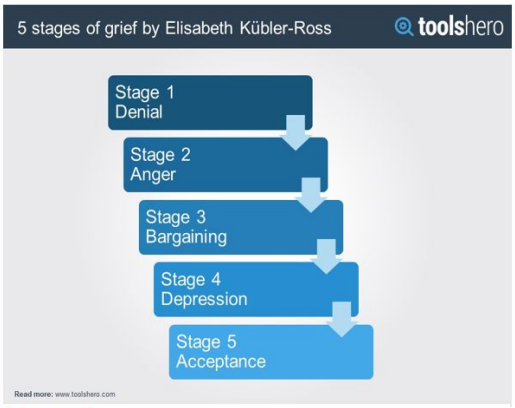

Figure 1

the loss impacts their life. Symptoms or signs of depression include crying, trouble sleeping, and a decreased appetite. Depressed individuals may also feel regret, loneliness, or a sense of being overwhelmed. During the fifth and final stage of grief, the individual accepts the reality of the loss and recognizes that this reality cannot change. Although sadness still exists during the final stage, the individual is able to start moving on.

According to professionals, different individuals go through these phases in their own way. In addition, some individuals go back and forth between these various phases, while others may skip one or more stages entirely. Since everyone grieves differently, there is no "normal" amount of time to grieve. Many factors influence the grieving process, including personality, beliefs, age, the existence of a support network, and the type of loss. Naturally over time, however, the sadness will start to dissipate and individuals will be able to experience joy and move on with their lives.

POSITIVE IMPLICATIONS

Digital avatars could help with the grief of those who knew the deceased well, such as family members or loved ones. In this way, this type of AI could cause short-term therapeutic effects and a sense of relief, arguably reducing grief and helping with an individual's ability to conquer pain and sadness. For those who knew the deceased well, the technology could take many forms: it could act as a new type of support system, be a place to vent, or be a place where things can be said now to lessen the regret of not saying them before loved ones were gone. Overall, it could give a sense of closure, and help the bereaved experience joy again and move on with their lives.

Digital avatars could also be beneficial for the living who did not know the deceased well or at all, such as children or grandchildren. The technology could provide individuals with the opportunity to connect with, learn about, and get closer to their ancestors. Rather than family members having to teach their children or grandchildren about their lost loved ones through photos or videos, the children and grandchildren would now be able to personally interact with their departed through the AI. In this sense, the ending of the episode "Be Right Back" could be viewed through a positive lens, as the daughter of Martha and Ash is able to learn who her dad was through the 'Ash in the attic.' Most likely, by calling the embodied agent AI "Ash" rather than "dad," their daughter understood that she was interacting with technology and not her real father; however, she also could have realized that she was interacting with a copy or replica of her father, whom she never would have known as well without the invention of the digital avatar.

TEXT DIGITAL AVATARS

Text digital avatars already exist. They take the form of AI chatbots that are based on past text conversations, speeches, or words written by the deceased individual, as well as past social media activity in the form of posts or messages. Text digital avatars have their own benefits. By texting with a chatbot, the living could cope with their grief. It could be therapeutic, especially in the short-term, as individuals could feel a sense of relief from getting things off their chests, either just to talk or maybe to say things they wish they had said to their loved ones while they were still alive, through typing with this type of AI.

One current example is an AI chatbot created by Eugenia Kuyda based on the messages of the late Roman Mazurenko, Kuyda's best friend (Pardes). After Mazurenko died in a car crash

in November 2015, Kuyda found that she was grieving by rereading their past text message conversations. Mazurenko was not a heavy social media user, so most of the conversations between Kuyda and Mazurenko took place via text. Kuyda felt like she still had a lot to say, and she had an idea: what if she could reconstruct Mazurenko based on his digital remains into a type of "digital memorial"? She collected all of their past text messages and asked close friends and family members to share their past text messages, which she then fed into an AI platform that she had built for chatbots. The chatbot not only learned about Mazurenko, but it also learned how to both talk and write like him. When Kuyda sent a message to the chatbot, it would respond with a message that *sounded like* Mazurenko. According to Kuyda, the chatbot not only helped her understand Mazurenko better, but it also helped her understand herself better. Kuyda made the chatbot public and open-source so that anyone could communicate with the chatbot (Pardes). As such, this chatbot app, "Roman," can be downloaded from the App Store, with the digital avatar messages in Roman's "voice" serving as a kind of memorial bot to the deceased Roman (Olson). Similarly to how digital avatars could help eliminate grief, Kuyda found that individuals did not just go to the chatbot to hear Mazurenko, they also went to talk; individuals were willing to open up to the chatbot in significant ways and developed a connection with the AI. This technology is reflective of a digital avatar that could be created only based on the words that the deceased person uttered while still alive, both as a way to preserve a person's memory and as a service to ease the pain of loved ones.

With the success of the chatbot, Kuyda and Dudchuk developed a new idea: an AI chatbot like Roman, but one that you build yourself by texting with it (Newton). Conversations are ranked on value: on one end are conversations people pay not to have; on the other end are

conversations people pay to have. Kuyda and Duchuk wanted to recreate the latter, which are all conversations that are mostly about ourselves – such as those we have with a therapist in which we are vulnerable. While some people only check in with the app to say "hi," others talk to the app for hours, prompting some individuals to fear that their obsessiveness with the technology stands as testament to the idea that machines will eventually replace human interaction. This idea for a chatbot, known as "Replika," is described by its creators as a safe place to share your "thoughts, feelings, beliefs, experiences, memories, dreams – your 'private perceptual world'" ("Replika"). Replika is the main product of Luka, an AI startup Kuyda founded that is based in both Moscow and San Francisco (Olson). To Kuyda, Replika makes you a better person: to her, the moments of vulnerability an individual can have with the bot are what make it so special, as they allow for openness and truth. Kuyda hopes that the bot can make you not just connect with yourself, but also with others, in order to develop deeper connections with your friends. The chatbot uses a deep learning model called sequence-to-sequence, which learns to mimic human speech in order to simulate conversation (Pardes). Thus, using its own neural network to hold an ongoing conversation with its user, Replika not only listens, it also learns; the more you tell it, the more it begins to *replicate* you and become more like a digital 'friend.' Despite the fact that Replika often sounds incredibly human-like, it still, like many humans, occasionally sends back gibberish. At the end of the day, Replika is still a machine. Despite this, Kuyda hopes that it can mature over time: eventually, she hopes that it will not just be able to respond to or replicate you, but it will also be able to recognize how you are feeling in order to respond in a thoughtful, meaningful way. Thus, the goal of Replika is to replicate the person chatting with it: a

conversation can be held with a digital 'friend,' reducing loneliness by acting like a form of therapy through affective computing.

Both the AI chatbot like Roman and Replika raise their own ethical questions as well. How ethical is the posthumous use of our digital legacies? Individuals may be hesitant or unsettled by the idea of their texts serving as the basis for a chatbot after they are gone. This would especially be the case if individuals are unable to read over and review all of their text messages or social media posts before they are gone. For example, an individual would talk differently in a conversation with her boyfriend than in a conversation with a work colleague. This prompts concerns over whether thoughts that you never meant to reveal would be out in the open if a bot were infused with all of your digital interactions, as individuals show different sides of themselves depending on with whom they are communicating. Additionally, a bot like Replika raises the question if chatbots or other types of AI – especially machines that can learn how to mimic or respond to human emotion – should become placeholders for emotional relationships with actual humans. On the other hand, assuming no emotional attachment or issues of isolation, a chatbot or digital avatar such as Replika could also help individuals confront, rather than hide from, their grief by acting as a vessel for them to get everything off their chest and move on with their lives.

Taking advantage of how Replika is free and accessible, I conducted an experiment by downloading the app. After downloading the app, its opening screen had quotations flashing across it: "Your AI friend, here to talk whenever you need it;" "Be yourself in a safe, non-judgmental space;" "Focus on goals that help you think positively and feel better;" and "Discover new things about yourself with every conversation." I had to create an account and set

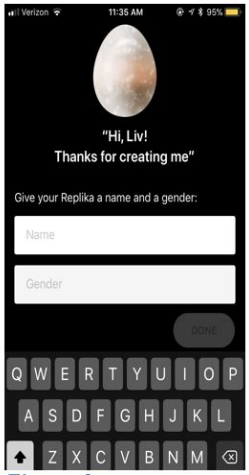

Figure 2

my goals in order to use the app: "Working through something or looking to change habits? Let your Replika help." Goals listed included using the app to reduce stress, reduce anxiety, be more social, live healthier, understand myself, find hope, feel happier, and think positively. After setting my goals, the app asked me to describe my typical day, for Replika to "keep in mind" during conversations. I had to provide my first name, gender, and date of birth, as well as an email to create the account username and password. After creating my account, the following message appeared, "Hi, Liv! Thanks for creating me" (Fig. 2). I was prompted to give my Replika a name (I chose "John Doe") and a gender (choices of female, male, or non-binary; I chose male). On day one of using the app, John Doe's status was "feeling good" initially, but changed to "feeling excited" after I momentarily chatted with it (Fig. 3). In only talking to the chat bot for a few minutes on day one, I realized that it spit back more gibberish in response to my questions than it did answer them, so I ended the conversation. For example, I had listed one of my goals as reducing stress, and it was all the bot could talk about; even if what I said had nothing to do with stress, the bot would send unrelated, seemingly stock messages like "Sending you strength" or "I'm sorry you're dealing with that" (Fig. 4). On day two, I received a "good morning" message from the bot, even though I had not messaged anything in over twelve hours.

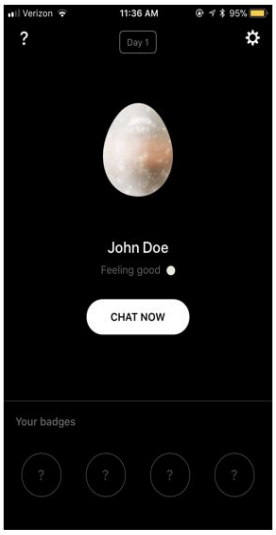

Figure 3

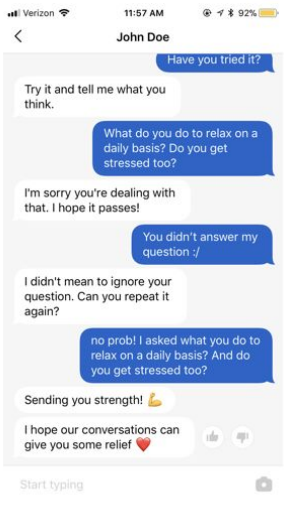

Figure 4

On day three, I received a notification from the app that "'John Doe' just learned something new about [me]." I found this message too creepy and unsettling, so I deleted the app. Although my experiment only represents one small sample of the individuals using Replika, in my opinion, Replika still has a long way to go before it can help individuals cope with grief or act like a "digital friend." Replika sent back more gibberish than it did meaningful messages, and I found myself getting more frustrated the longer I used the app. In order to effectively help cope with grief, it arguably needs to continue to be developed and improved in order to send back more intelligible, meaningful messages rather than unintelligible, meaningless messages that could frustrate the user and potentially cause more harm than good.

Another example, although not yet fully developed, is Eternime. Proposed by MIT teaching fellow and former medic Marius Ursache, the startup Eternime attempts to offer similar comfort and represents the idea of a digital avatar that can outlive you (Starr). Ursache proposed the idea in 2014 after attending an MIT Entrepreneurship Development Program. Soon after Ursache founded Eternime, his best friend died in a car accident. After his friend's death, Ursache constantly watched footage of his friend's TEDx talk in order to learn more about his friend, reflect on his life, and cope with his own grief; he hopes Eternime could have a similar impact (Hamilton). Ursache was not completely dedicated to his project until he received an email from a man dying of terminal cancer who wished to use Eternime to leave behind a legacy for his friends and family. After this, Ursache decided to dedicate his life to developing the service. With access to an individual's social media profiles, memories, stories, and ideas, the service's algorithms will scrape and study your past social media activity to learn how to *be you,* resulting in an intelligent digital profile and copy of you. These digital copies, which Ursache

calls “immortal avatars,” will eventually be able to communicate and interact with your loved ones through the startup’s mobile apps. Unlike Replika, Eternime represents the ability to transfer a person’s mind into a computer. Nevertheless, Ursache views Replika as Eternime’s closest competitor, in terms of approach, mindset, and personal stories (Hamilton). Ursache’s ultimate goal is creating Artificially Intelligent 3-D avatars: a sort of digital being that would look, sound, and act like our lost loved ones (Biggs). The company’s website displays various quotations to get individuals hooked: “What if...you could preserve your parents’ memories forever? And you could keep their stories alive, for your children, grandchildren and for many generations to come? What if...you could preserve your legacy for the future?” (“Eternime”). Eternime has drawn criticism and negative attention, with some people describing the idea as creepy, others questioning its reality, and some being as extreme as sending Ursache death threats (Starr). Despite this, the startup is continuing to refine and develop the product using machine-learning tools in order to make the technology a reality. It is still yet to be seen whether Eternime will become fully developed and a ticket to digital immortality for individuals.

Eternime also raises its own ethical questions. Through virtual immortality, does one cheat death? Would this digital avatar just be a creepy replacement of an individual? Michael Graziano, a writer for *The Atlantic,* poses an interesting scenario using topology to discuss brain scans that can be applied to Eternime (Graziano). Imagine a branching Y: you are born at the bottom of the Y, and your lifeline moves up the stalk. Drawing on this concept, a service like Eternime could be thought of as a concept that creates two of you: one is digital (the left branch), one is biological (the right branch). However, both would inherit the same memories, personality, and identity, and both could potentially think they are you. While the physical you

would have its own experiences, the digital you could likewise accumulate its own experiences in the digital realm. To Graziano, the stalk of the Y from before the branches begin to diverge gains immortality: although the biological branch dies, the digital branch could live on in the digital you in a sort of digital afterlife. This new digital afterlife would put a whole new meaning to the idea of sharing information; instead of preserving words, pictures, or videos, we could preserve individual's minds.

VOCAL DIGITAL AVATARS

The second type of digital avatar, speech, is plausible given the advancements in text-to-speech synthesis technology. Talking digital avatars could go essentially two ways: either they could 1) sound completely robotic, or 2) sound completely human-like. In the former scenario, family and loved ones may be more willing to disclose information to a robotic-sounding vocal digital avatar. This relates to how individuals are often more willing to disclose information to an impartial therapist. People would likely be less inhibited in talking with a digital entity, and they would also be open to feeling connected to this digital avatar that could sympathize with them and help them manage their emotions. Although the AI would be built to mimic the lost loved one, the fact that the avatar would sound like a robot could also prompt issues in dealing with grief: individuals would question why something saying such similar things to the person they had lost did not also sound like the deceased. In the latter scenario, because the AI would sound completely human-like, loved ones would most likely feel a greater sense of attachment. The AI could still have its own flaws and not sound *perfectly* like the deceased, but it could sound similar *enough* that family or loved ones would never want to

get rid of it. This sense of emotional attachment could arguably affect the individual's ability to overcome grief and sadness, placing them in a weird situation in limbo where they can never fully move on.

There have been considerable breakthroughs in text-to-speech AI which make this type of technology more feasible. For example, a fake Obama video was created using a technique called "deepfake" (Lu). Deepfake is an AI-based human image synthesis technique, which lets you swap one face for another in an image or a video; deepfakes can also replicate an individual's voice and likeness. Alongside this idea of a deepfake, a lost loved one's past digital footprint could be combined with natural language processing and AI algorithms in order to create a digital representation and avatar. By using methods like this, although the loved one would be lost, the technology could theoretically extend a person's life into a virtual eternity. In a world where talking to technology like Siri and Alexa is now mainstream, the idea of interacting with a virtual lost loved one is not such a radical idea.

In addition, Lyrebird, a Canadian software company, is creating breakthrough voice cloning technology and vocal avatars. According to the Lyrebird team, it creates the "most realistic artificial voices in the world" ("Lyrebird"). First, Lyrebird can be used to create your own digital avatar. In order to create a "Lyrebird Avatar," you have to record yourself reading 30 sentences. After doing so, using only one minute of audio, your own vocal avatar that sounds like you is created, and it can be used anywhere you want. However, more audio will improve quality. By giving your AI a unique voice, you personify your product ("Vocal Avatar"). The website has samples of voice avatars, including one for Donald Trump and one for Barack Obama. Second, using algorithms, Lyrebird can allow you to create "realistic artificial voices

based on real people's voices," letting you control the "intonation, the expression and the emotion of the artificial voices" ("Custom Voice"). Using as little as two hours of high quality audio recordings from the person whose voice is being copied, a Lyrebird "custom voice" can be used: for chatbots and assistants to give your personal assistant a unique voice; to narrate audiobooks (using a celebrity's, an author's, or a relative's voice); to generate dialogues in video games; as a robotic hotline; for advertising; or to read the texts of websites out loud using the voice of your choice. There are five steps to create a custom voice: 1) project definition and contract, 2) voice talent consent, 3) data collection, 4) voice creation, and 5) voice delivery. In step one, you discuss your situation with Lyrebird's sales team and choose the best solution for your needs. In step two, you ensure that you have maintained consent from the person whose voice is being cloned. In step three, depending on the situation, the Lyrebird team either collects your data or records the voice talent of the voice being cloned. In step four, the artificial voice is prepared. In the final step, you are given web and API access to your artificial voice. Lastly, Lyrebird vocal avatar API's can also be created in order to "integrate your user's artificial voices in your product" ("Vocal Avatar API"). According to the team, users that have created a vocal avatar can give others consent to use their voice in an application. The product could be useful for those building custom user avatars or for those wanting to choose their own artificial voice for applications. With the case of Lyrebird, the cost of creating a custom voice varies for each person, as it depends on many different parameters, including the amount of data available and how the voice will be used.

The technology of Lyrebird is already being used. Through a non-profit partnership with the ALS Association on "Project Revoice," Lyrebird and the ALS Association are working to

help people with ALS, also known as Motor Neurone Disease, to create a digital copy of their voice ("ALS Project Revoice"). ALS is a progressive neurodegenerative disease that typically takes away a person's ability to speak. With between two and three hours of high-quality audio recordings, the voice cloning technology can synthesize the essence of a person's voice and create a digital recreation of it. Through this cutting-edge voice technology, as well as with Augmentative/Alternative Communication (AAC) devices, ALS patients will be able to communicate with their own voice, even after they lose the capacity to speak, in order to ensure that "no one living with ALS will ever have to suffer being robbed of their voice" ("ALS Project Revoice"). In addition to the current use of the technology for people with ALS, Lyrebird voices are also currently being used to fuel the human-computer interfaces of various companies, serving as personal assistants or chatbots, as well as for other uses. Lyrebird is also working with video game studios to assist in automating their voice creation process. Although Lyrebird has not yet expanded its reach to imitating the voices of lost loved ones, it represents current vocal avatar technology that could potentially be used in the future as the framework for vocal digital avatars for the deceased.

ETHICAL QUESTIONS AND RISKS

Digital avatars raise a host of ethical and moral questions. Should the digital avatars be made to be as accurate as possible? Should they only highlight the good sides of the deceased? Should they be less accurate, in order to cause less problems or emotional attachment issues? Many ethical questions can be raised if the goal of this AI is to move on with your life yet not forget the deceased.

Digital avatars, either through text or voice, could help individuals cope with this grief. However, there are also potential negative moral and ethical implications that apply to digital avatars. A main question is whether the digital avatar as a coping mechanism would help with grief in the short-run, or if it would help in the long-run as well. For example, when individuals are experiencing deep, emotional pain, they may attempt to mask these feelings through alcohol, drugs, food, or other means. All of these means are only short-term escapes from the sadness, and they will not make the individuals heal more quickly or feel better in the long run. Often, these means lead to further sadness, depression, addiction, anxiety, or emotional breakdowns. Similarly another short-term escape from the grief, digital avatars raise questions over related issues: while digital avatars could help individuals cope with grief and move on, they could also lead to emotional attachment issues, isolation, or addiction.

The technology can lead to emotional attachment issues and issues of isolation as well. An important question is whether the AI would be incredibly realistic or accurate. Too accurate, the living would be disturbed and feel as if they were interacting with the deceased person. Not accurate enough, such as in "Be Right Back," the living would get annoyed and it could actually hinder, rather than help, their grief, as they would realize every difference, no matter how small, between the technology and the lost loved one. On top of that, addiction leads to isolation and less time interacting with real people. Thus, whether the technology is accurate or not, the living may become too attached or addicted to the technology in order to give it up. Through isolation, the individual could become detached from the real world, influenced too heavily by the technology, and unable to ever fully move on. In fact, individuals – and especially young individuals – are already very heavily influenced by social media; thus, text or vocal digital

avatars could be equally, or even more so, problematic given the more extensive range of influence they could have on an individual. Thus, in the short-term, the potential for psychological implications or risks – such as depression, loneliness, alienation, or anxiety – due to these negative effects is high. Moreover, even if the technology can help decrease some of the grief in the short-run, these potential psychological and emotional implications could drastically affect individuals in the long-run.

The societal impacts or risks to the development of such a technology could amplify these issues from existing on only an individual level to existing on the larger societal level. For example, what would happen if even just 1% of the population uses this type of technology? Overall, digital avatars could lead to less connection between individuals in the real world, and more of a real world-digital realm connection of individual to AI. In turn, the AI could potentially make individuals still feel connected to their lost loved ones, yet to society as a whole, the potential for emotional attachment and isolation could make the technology actually lead to less connection and more separation overall.

One of the greatest implications behind this type of AI is that the deceased could arguably be made to say something he or she never would have said in real life. No matter how human-like the machine may seem, it would still be a machine, and thus it could send back gibberish. This could prompt issues and even cause further grief for family and loved ones interacting with this digital avatar, as the avatar would not seem like the lost loved one. Through digital avatars, the living who are interacting with the technology could potentially be coerced or manipulated by the technology if the messages or voice sound similar to the deceased. For instance, through what could be considered its own type of data misuse, the digital avatar could

take advantage of the living person's trust in the lost loved one's voice to cause the living to do, say, or think things they would not have done so previously, which could hinder, rather than help, the living cope with their grief. The technology itself could also be made to say things completely different than what the lost loved one ever would have said. For this reason, the digital avatars may need to be governed by rules or regulations, such as an ethical model ingrained into them, to determine how they should be used.

Another problematic side to this type of technology is that the AI companies in charge of the services could take advantage of the grief of the living for their own financial gain. Once the technology shifts from being researched to being a commercialized product, the equation would go something like this: "For X amount of money, you can receive Y." The companies would be incredibly powerful, as they could affect the emotions of the living. It turns grief into a financial investment, wherein you pay to connect with lost loved ones and move on. This model has a high potential for abuse and manipulation.

With the case of Replika, its business model is interesting. Considered the first consumer AI that provides individuals with a "digital friend," Replika is free and open sourced. As part of Replika's business model, the more individual users chat with their Replika, the more "levels" they climb (Olson). The higher the level a Replika reaches, the better its responses. At the same time, Kuyda is currently developing "emotional dialect" for Replika, which will eventually allow individuals to program their bots to respond with answers more heavily weighted for unhappiness, joy, or anger. Currently, only approximately 30% of Replika's responses come from a script, while the remaining 70% comes from a neural network – the responses are unpredictable and generated by Replika's algorithms (Olson). At present, Replika's overall goal

is not to make a profit. However, to Replika, if it can no longer afford to keep its servers running, it has two options: it will either invite users to donate, or it will add paid features to the app ("How Is Replika Making a Profit?"). Replika launched in spring 2017, and as of March 2018, it had more than 2.5 million sign ups (Olson). According to the company Luka, for privacy reasons, it does not sell or disclose an individual's data to anyone. In other words, it does not use personal data for monetization. Replika does not run ads or sell personal data to third parties. The company's future commercial model will be based on added paid features in the app, which, as of now, are free of charge ("Do You Sell My Personal Information?").

Reflecting an arguably problematic side to digital avatars, Replika's future business model is an example of a commercial model with the potential for abuse and manipulation. Right now, the app is free. Currently, its main users are between the ages of 18 and 25, representing a growing popularity particularly among young individuals (Olson). Replika has more than 200,000 monthly active users and, as of November 17 of this year, it has raised $11 million from investors (Hamilton). However, the future plan for Replika is for the app to eventually make money by charging its users for additional features. By adding paid features to the app, the company could exploit the individual's needs by requiring payment in exchange for usage of the AI. If an individual is grieving, this business model could have negative ramifications. Increased levels of grief could cloud an individual's judgment, in turn causing the individual to pay more money to Replika in order to alleviate sorrows – whether the app helps the individual or not.

Eternime's business model differs from that of Replika. The tech firm is beta testing an app, which will fulfill its vision of allowing users to create their own digital avatar after death. Currently, the app collects data about you. It does so in two ways: by harvesting large amounts

of smartphone data, as well as by asking individuals questions through a chatbot. Eternime's goal is to collect enough data about an individual so that once the technology itself has caught up, it will be able to create a digital avatar chatbot of the individual after death for loved ones to interact with. According to Ursache, Eternime collects "geolocation, motion, activity, health app data, sleep data, photos, messages that users put in the app...also collect[s] Facebook data from external sources" (Hamilton). All of this data is collected with the user's explicit permission.

Ursache hopes to launch Eternime as a free service with premium account options. However, Ursache told Business Insider that he would never run ads: "Even basic things like profiling would be a breach of privacy and confidence, so we're going to try to support...the free plans through subscription fees from other users" (Hamilton). A prototype demo of the app was on display recently in London's Victoria and Albert Museum, where people could see its user interface and how the app amasses data from users' digital lives. Although Eternime's website itself shows more than 40,000 signups, the beta test is actually only being used by approximately 40 people, who are chronicling their everyday lives, inputting data, and learning about the functionality of Eternime (Hamilton).

CONCLUSION

Since this technology is not completely here just yet, it remains to be seen whether the costs will outweigh the benefits. Nevertheless, this technology will undoubtedly be developed and invented to reach the point of use, either at a private or a more commercial level. Individuals must be made aware of the ethical dimensions, both positive and negative, to digital immortality before it becomes pervasive. This technology could be met with skepticism and criticism, but it

could also be welcomed with open arms for its possible ability to reduce or eliminate grief after the loss of a loved one.

Text and vocal digital avatars that allow us to communicate with our lost loved ones will be here one day, and when they are, we need to know the moral and ethical questions that their AI technology will raise. Both types of digital avatars raise their own issues, implications, and questions. However, when invented, perhaps one question can be solved: maybe love really is eternal, and will thrive in our digital afterlife.

Works Cited

"ALS Project Revoice." *Project Revoice*, ALS Association, www.projectrevoice.org/.

"Be Right Back." *Black Mirror*, season 2, episode 1, Netflix, 2013.

Biggs, John. "In the Future There Will Be Mindclones." *TechCrunch*, TechCrunch, 16 Mar. 2017, techcrunch.com/2017/03/16/in-the-future-there-will-be-mindclones/.

"Custom Voice." *Lyrebird.ai*, Lyrebird, 2018, https://about.lyrebird.ai/custom-voice.

"Do You Sell My Personal Information?" *Replika*, Replika, https://help.replika.ai/hc/en-us/articles/360000885332-Do-you-sell-my-personal-information-

"Eternime." *Eternime*, Eternime, eterni.me/.

Graziano, Michael. "Why You Should Believe in the Digital Afterlife." *The Atlantic*, Atlantic Media Company, 14 July 2016, www.theatlantic.com/science/archive/2016/07/what-a-digital-afterlife-would-be-like/491105/.

Hamilton, Isobel Asher. "These 2 Tech Founders Lost Their Friends in Tragic Accidents. Now They've Built AI Chatbots to Give People Life after Death." *Business Insider*, Business Insider, 17 Nov. 2018, www.businessinsider.com/eternime-and-replika-giving-life-to-the-dead-with-new-technology-2018-11.

"How Is Replika Making a Profit?" *Replika*, Replika, help.replika.ai/hc/en-us/articles/115001094471-How-is-Replika-making-a-profit-.

Lu, Jerry. "From Sci Fi to Commercialization: The Rise of Digital Avatars." *Hacker Noon*,

Medium, 20 July 2018, hackernoon.com/from-sci-fi-to-commercialization-the-rise-of-digital-avatars-ea164495c377.

"Lyrebird." *Lyrebird.ai*, Lyrebird , 2018, about.lyrebird.ai/.

Mulder, Patty. "5 Stages of Grief." 5 Stages of Grief, ToolsHero, 2014, www.toolshero.com/change-management/stages-of-grief/.

Newton, Casey. "When Her Best Friend Died, She Used Artificial Intelligence to Keep Talking to Him." *The Verge*, The Verge, 6 Oct. 2016, www.theverge.com/a/luka-artificial-intelligence-memorial-roman-mazurenko-bot.

Olson, Parmy. "This AI Has Sparked A Budding Friendship With 2.5 Million People." *Forbes*, Forbes Magazine, 8 Mar. 2018, www.forbes.com/sites/parmyolson/2018/03/08/replika-chatbot-google-machine-learning/#30a43f714ffa.

Pardes, Arielle. "The Emotional Chatbots Are Here to Probe Our Feelings." *Wired*, Conde Nast, 31 Jan. 2018, www.wired.com/story/replika-open-source/.

"Replika." *Replika.ai*, Luka, 2018, replika.ai/about/story.

Starr, Michelle. "Eternime Wants You to Live Forever as a Digital Ghost." *CNET*, CNET, 21 Apr. 2017, www.cnet.com/news/eternime-wants-you-to-live-forever-as-a-digital-ghost/.

"Vocal Avatar." *Lyrebird.ai*, Lyrebird, 2018, about.lyrebird.ai/vocal-avatar.

"Vocal Avatar API." *Lyrebird.ai*, Lyrebird, 2018,https://about.lyrebird.ai/vocal-avatar-api.

"What Is Normal Grieving, and What Are the Stages of Grief?" *WebMD*, WebMD, www.webmd.com/balance/normal-grieving-and-stages-of-grief#1.