



Griefbots, Deadbots, Postmortem Avatars: on Responsible Applications of Generative AI in the Digital Afterlife Industry

Tomasz Hollanek¹ · Katarzyna Nowaczyk-Basińska¹

Received: 17 November 2023 / Accepted: 8 April 2024
© The Author(s) 2024

Abstract

To analyze potential negative consequences of adopting generative AI solutions in the digital afterlife industry (DAI), in this paper we present three speculative design scenarios for AI-enabled simulation of the deceased. We highlight the perspectives of the *data donor*, *data recipient*, and *service interactant* – terms we employ to denote those whose data is used to create ‘deadbots,’ those in possession of the donor’s data after their death, and those who are meant to interact with the end product. We draw on the scenarios to map out several key ethical concerns posed by ‘re-creation services’ and to put forward recommendations on the ethical development of AI systems in this specific area of application. The recommendations, targeted at providers of AI-enabled re-creation services, include suggestions for developing sensitive procedures for retiring deadbots, ensuring meaningful transparency, restricting access to such services to adult users only, and adhering to the principle of mutual consent of both data donors and service interactants. While we suggest practical solutions to the socio-ethical challenges posed by the emergence of re-creation services, we also emphasize the importance of ongoing interdisciplinary research at the intersection of the ethics of AI and the ethics of the DAI.

Keywords Griefbots · Deadbots · Chatbots · Generative AI · Digital afterlife · AI ethics · Thanatosensitivity · Postmortem privacy

✉ Tomasz Hollanek
th536@cam.ac.uk

Katarzyna Nowaczyk-Basińska
kn395@cam.ac.uk

¹ Leverhulme Centre for the Future of Intelligence, University of Cambridge, 16 Mill Lane, Cambridge CB2 1SB, UK

1 Introduction

For a small payment, the online platform Project December (PD) grants users access to a ‘deep AI running on one of the world’s most sophisticated super-computers’ and allows them to participate in simulated ‘text-based conversation[s] with anyone’ – including ‘someone who is no longer alive’ (2023). The platform’s earlier version came under public scrutiny when stories about a man who used the PD website to interact with his deceased fiancée’s avatar started to circulate the web in 2021 (Fagon, 2021), and when OpenAI, whose GPT-3 model initially powered the platform, reportedly terminated PD’s access to its API, citing Project December’s failure to abide by its safety guidelines (Robitzski, 2021). While OpenAI’s usage policies do not outright prohibit the use of its large language models (LLMs) in the production of the so-called ‘deadbots,’ the guidelines indeed specify that any conversational AI system ‘simulat[ing] another person’ – with the exception of ‘historical public figures’ – is required to ‘either have that person’s explicit consent or be clearly labeled as “simulated” or “parody”’ (OpenAI, 2023). Failing to follow OpenAI’s safety team’s instructions, Project December was forced to temporarily suspend its operations – but only to evolve into the platform that it is now, built upon its own ‘patent-pending technology’ and continuing to offer users the opportunity to ‘simulate the dead’ (Project December, 2023).

The story of Project December’s evolution, interweaved with that of OpenAI’s usage policy development, points to how the rapid progress in the broadly construed field of ‘generative’ AI – with advancements in natural language processing in particular – relates to the accelerated expansion of what we refer to, following Öhman and Floridi (2018), as the digital afterlife industry (DAI). While the DAI comprises new data management services in charge of ‘digital remains’ on behalf of the deceased and digital memorial services targeting the bereaved, our interest lies specifically in AI-powered simulations of the dead, akin to those offered by Project December, concerning both the deceased and the bereaved. Drawing on Öhman and Floridi’s categorization (2017), we adopt the term ‘re-creation service’ to denote an AI service specializing in postmortem simulation of the dead. Additionally, we use the term ‘deadbot’ to refer to an AI-enabled digital representation of a deceased individual created by a re-creation service.¹

Responding to the ongoing, unrestricted ‘democratization’ of ‘immortalization’ technologies, through this paper, we aim to bridge the persistent gap between the fields of AI ethics and the ethics of the DAI and map out the social and ethical challenges posed by the unregulated use of AI in the digital afterlife industry. Within our study, we identify three primary stakeholder groups: *data donors*, *data recipients*, and *service interactants*. The term *data donors* refers to individuals whose data contributes to the creation of deadbots; *data recipients* are in possession of the kind of data that can be used to create a data donor’s deadbot; *service interactants*, in turn,

¹ It is also important to highlight that the literature employs a range of sub-terms for ‘deadbots,’ including ‘thanabots,’ ‘postmortem avatars,’ ‘griefbots,’ ‘ghostbots,’ and ‘mind clones,’ which, as of now, are used largely interchangeably without a clear differentiation or specification.

are those meant to engage with the resulting deadbot. Most academic work analyzing the ethical and legal implications of simulating the deceased revolves around the perspective of the departed (e.g. Buben, 2015; Öhman & Floridi, 2017; Harbinja in: Savin-Baden and Mason-Robbie, 2020; Stokes, 2021), with less attention given to the perspective of the bereaved (e.g. Krueger & Osler, 2022; Lindemann, 2022). However, as of now, the complex relationships between the mentioned stakeholder groups – *data donors*, *data recipients*, and *service interactants* – remain unaddressed. The advent of re-creation services has introduced a particularly intricate situation in which the person whose data is used to inform the design of a given interactive product (the data donor) is not its intended end user (the service interactant). This complexity necessitates that, to determine what constitutes responsible deployment of AI in the DAI, we consider the interconnected interests, rights, and needs of different groups of stakeholders that partake in re-creation projects.

Bearing the fundamentally relational nature of re-creation services in mind, we draw on speculative design as a method for considering the socio-ethical dimensions of technology development and a means for eliciting alternative design values, principles, or methods that should be prioritized to allow for socially desirable outcomes of technological development. We present three speculative design and business scenarios focusing on different uses of re-creation services to then formulate a set of recommendations for providers of such services. These recommendations draw on already existing frameworks for responsible AI development but focus specifically on the use of generative AI in the digital afterlife industry – an area of AI application that remains understudied by AI ethics and human–computer interaction scholarship. The exercise of mapping the ethical challenges posed by re-creation services and conceiving of potential solutions through speculative design is intended to lay the groundwork for future interventions in technology design standards and policy development that, as we demonstrate, are needed to mitigate the risks posed by the use of AI in the digital afterlife industry.

2 The Intersection of the Ethics of AI and the Ethics of the DAI

In the last two decades, academia and industry have witnessed a surge in initiatives aimed at tackling challenges pertaining to death and mortality within product and interaction design. The development of ‘thanatosensitivity’ as a new design paradigm was one of the early responses to these challenges within human–computer interaction. Massimi & Charise (2009), who coined the term, argued that prevalent interaction design practices had failed to account for death as the key element of the human experience; thanatosensitivity, or the attention to the matters of death in interaction design, serves to identify potential design problems and delineate areas for improvement. Building upon the thanatosensitivity framework, design and research teams have developed new design approaches, such as the ‘lifespan-oriented approach’ (Massimi et al., 2011), and concrete large-scale solutions, such as Facebook’s Legacy Contact feature (Brubaker et al., 2014) and Google’s Inactive Account Manager, as well as smaller projects like ReFind (Wallace et al., 2020).

While these new solutions and functionalities aim to acknowledge the inevitable mortality of technology users, we can also observe a growing number of technology design efforts with the opposite goal: instead of acknowledging death, they aim to transcend it. The explorations of technology-enabled ‘immortalization,’ akin to Project December, encompass the development of memorialization and art projects (e.g. James Vlahos’s *Dadbot*; Hanson Robotics’ *BINA48*; Marlynn Wei’s *Elixir*), as well as the introduction of new functionalities to existing products (e.g. Amazon’s Alexa speaking with the voice of a deceased relative; see: Allyn, 2022), and the establishment of start-ups (e.g. *HereAfter*). These examples signal the development of a wider trend in technology design, whose sheer scale is attested by the term ‘digital afterlife industry,’ which underscores the growing significance of ‘immortality’ as a market segment. Indeed, the story of Microsoft’s recently secured patent for software that could ‘resurrect’ the dead as chatbots points to the fact that the question of technology-enabled ‘immortality’ has already appeared on the radar of tech giants (Smith, 2021). At the same time, thanks to the rapid advancements in generative AI, the option to simulate the deceased has become more widely available. Unlike in the past, when setting up re-creation services demanded specialized skills and a substantial budget, today, almost anyone can bring a deceased loved one ‘back to life,’ as evidenced by numerous instances in China (Loh, 2023) and the United States (Percy, 2023).

Despite the rapid growth of this sector within the DAI, the matter of socio-ethical risks posed by re-creation services has been largely overlooked within the broader field of AI ethics.² This oversight within AI ethics scholarship has also resulted in a persistent void in AI-related policy and design standards work; to the best of our knowledge, the already mentioned OpenAI’s usage policy is the only document of its kind that acknowledges, albeit indirectly, that the use of AI in the simulation of deceased individuals is an area of application that necessitates additional precautions. Re-creation services raise ethical concerns that neither the thanatosensitivity framework – focused on the mortality of users, rather than their postmortem activity – nor the already available guidelines for responsible AI development can help re-creation service providers resolve comprehensively. Meta-analyses of available responsible AI guidelines (Jobin et al., 2019; Attard-Frost et al., 2023; Wong et al., 2023) demonstrate that these guidelines may be useful for considering technical aspects of responsible AI production, such as data bias, but fail to guide developers through addressing more complicated socio-ethical challenges. This is in part due to the guidelines’ ‘high-level’ nature. The recommendations that we put forward in this

² To be clear, we are aware of the work in philosophy, anthropology, and sociology dedicated to re-creation services (including: Kasket, 2019; Lagerkvist, 2022; Stokes, 2021; Sumiala, 2021), but we are referring here to the absence of dedicated scholarly papers on the topic featured at the key AI ethics conferences, specifically: the ACM Conference on Fairness, Accountability, and Transparency (FAcT) and the AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society (AIES). Similarly, none of the journals that constitute popular outlets for AI ethics research – *Minds and Machines*, *Philosophy & Technology*, *AI & Society*, and *AI and Ethics* – seem to have published articles that directly delve into the subject of re-creation services and the need for their regulation (as of the submission date and to the best of our knowledge).

article for the providers of re-creation services are meant to help them consider and navigate through the complex socio-ethical issues that are specific to this particularly delicate area of AI application.

A few previous interventions have already highlighted the special nature of re-creation services that distinguishes them from other types of AI systems and gestured towards the need for additional guardrails for integrating AI into the DAI. For instance, in their article on the ethical framework for the DAI, Öhman and Floridi (2017, 2018) suggest several measures for protecting the dignity for those who are 're-mediated' through deadbots – focusing on the perspective of data donors. In a more recent paper, Lindemann (2022), who analyzes the technology's influence on the grieving process, suggests that deadbots should be regulated as medical devices to protect the end users' wellbeing – focusing on the perspective of service interactants. Despite these early contributions that move beyond the examination of risks posed by re-creation services to proposing concrete guardrails for their development, a comprehensive framework for the ethical production of deadbots that highlights the rights and (sometimes conflicting) needs of data donors, data recipients, and service interactants *in tandem* remains absent and this is precisely what we hope to develop through this article. While we build on these earlier recommendations for the ethical development of deadbots, we modify them and put forward additional ones – to fully account for the intricate, deeply relational nature of re-creation services that we highlight through our design fictions.

3 Methodology and Scope

In this article, we draw on *design fiction* to distill several key ethical concerns posed by re-creation services and to put forward recommendations on the ethical development of AI systems in this specific area of AI application. As defined by Bruce Sterling, design fiction is a practice that aims towards 'a suspension of disbelief about change achieved through the use of diegetic prototypes' (in: Bosch, 2012). It falls under the broader category of speculative design, or the kind of design practice whose products are not meant to be widely adopted or sold, but which prompt audiences to pose questions about possible futures and their relationship to the present, including the socio-economic and political realities that make only some of these futures – and, therefore, only some objects of design – appear realizable or desirable (Dunne & Raby, 2013). Design fiction draws on the narrative property of design – the fact objects themselves can tell stories and that broadly understood stories often rely on 'diegetic prototypes' to make the worlds they represent appear plausible (Bleecker, 2009) – and has been applied to future policy scoping work (Imagination Lancaster, 2023) or human–computer interaction research (Sturdee et al., 2016), as well as in eliciting and challenging ethico-political assumptions behind dominant design practices, to then make recommendations on alternative, socially desirable practices (Bardzell & Bardzell, 2013).

In what follows, we showcase three such 'diegetic prototypes' of re-creation services – *MaNana* (Fig. 1), *Paren't* (Fig. 2), and *Stay* (Fig. 3) – and three scenarios presenting their imagined use cases and potential users. We created

the prototypes paying attention to catchy names and taglines (summarized in Tables 1, 2, and 3) to ensure they appeared plausible. Before we delve into the scenarios, we must stress that the fictional products represent several types of deadbots that are, as of now, *technologically possible* and *legally realizable*. Our scenarios are speculative, but the negative social impact of re-creation services is not just a potential issue that we *might* have to grapple with at some point in the future. On the contrary, Project December and other products and companies mentioned in Part 2 illustrate that the use of AI in the digital afterlife industry already constitutes a legal and ethical challenge today.

To expound the logic behind our work on imagining the prototypes and constructing the accompanying user-focused stories, we must first elaborate on the key perspectives that we underscore in the scenarios: of those whose ‘digital remains’ are utilized in the process of deadbot creation; of those who have access to the kinds of data that can be used to produce a deadbot; and of the living users of re-creation services meant to interact with deadbots. We refer to these three types of stakeholders in the DAI as *data donors*, *data recipients*, and *service interactants*.

The term *data donor* alludes to previous work on the ethics of posthumous medical data donation (Krutzinna & Floridi, 2019; Harbinja, 2019); in our framing, the *donor* is the *source* of data – extending beyond medical records to include other forms of data such as emails or text messages – that can be used to produce a deadbot. The term refers to those who provide a re-creation service with their personal data directly and willingly with the intention of creating their own deadbot; to individuals who do not provide their data directly to any re-creation service, but who consent to the use of their personal information in this context by a third party, such as a relative or friend; as well as those individuals whose data is provided to a re-creation service by a third party without the donor’s explicit and meaningful consent.

The *data recipient* constitutes the ‘third party’ mentioned above. While the term *data recipient* has been used in different contexts to refer to a broader set of actors (e.g. the European Union’s regulatory framework for data protection), for the purposes of this study it signifies, more specifically, those individuals who are in possession of the kinds of data that can be used by a re-creation service to create a donor’s deadbot. The data we have in mind is generated during interactions between donors and recipients – for instance, when exchanging text messages or emails – hence the recipients have immediate access to the data after the donor’s demise; further considerations of the legal status of other forms of posthumous personal data are beyond the scope of this article.

Service interactants are the intended users of re-creation services, meant to interact with a deadbot after the donor’s death. In some cases, service interactants are also data recipients – when those in possession of a donor’s data supply it to a re-creation service to produce a deadbot they would like to interact with. In other cases, service interactants are not synonymous with data recipients – when it is the donor who creates their own deadbot and designates a service interactant not involved in the process of deadbot production, or when a data recipient creates a deadbot with someone other than themselves in mind as the intended interactant. We distinguish between these different roles and positions among the key stakeholder groups within

the digital afterlife industry to underscore the fundamentally relational nature of re-creation services. We refrain from using the term *end user*, as both *data donors* and *data recipients* can employ a re-creation service to ‘immortalize’ themselves or their loved ones, while the term *service interactant* refers specifically to those who are supposed to *interact* with a deadbot.

Appreciating the complexity of the relationships between different stakeholders and their roles in re-creation projects constitutes the necessary first step in analyzing the socio-ethical dimensions of deploying AI in the DAI. We conceived and visualized three re-creation service ‘prototypes’ to foreground these intricacies. The prototypes represent different modes of deadbot production, different goals of technological immortalization, different types of engagement they facilitate, and different re-creation service revenue models.

Each user-focused scenario is followed by an analysis of the ethical dimensions of the re-creation services’ impact on different stakeholder groups. In our discussion of MaNana, we focus on the impact of re-creation services on data donors and the role that data recipients play in determining whether this impact is negative; in the analysis of Paren’t, we foreground the influence on service interactants; and in the discussion of Stay, we delve into the impact on the relationships between donors and interactants, as well as between different interactants. Each of the imagined re-creation services affects all of the mentioned stakeholder groups and the relationships between these groups. However, we split up our analysis of individual products and scenarios this way to ensure that our recommendations for the providers of re-creation services clearly tie to the analyses of the impact of deploying AI in the DAI on specific stakeholders. We present our recommendations this way to ensure clarity, but to have a positive effect on re-creation services, they must be followed concurrently.

Finally, we should note that, while our recommendations point to concrete solutions, each recommendation should also be read as highlighting the need for further research, including user studies, in this particular area of AI application that remains, as we have noted, understudied by AI ethics and HCI scholarship.

4 Impact of Re-creation Services on Data Donors

4.1 Design Fiction I: *MaNana*, Bianca and Laura

Let us explore a hypothetical scenario featuring Bianca, a thirty-five-year-old woman who decides to use a speculative – yet plausible – re-creation service called MaNana (outlined in Table 1). MaNana enables users to construct deadbots of their deceased grandmothers (with alternative versions of the service enabling the ‘resurrection’ of grandfathers or similar significant figures in an individual’s life) to provide companionship and entertainment, rather than help with processing grief.

Bianca lost her grandmother, Laura, when she was twenty-eight. Bianca and Laura were close and – after Bianca left her home country to take up a new job abroad – they would often call, text, or send voice messages to each other. It has now been seven years since Laura’s passing. Bianca is no longer grieving, but she still misses her grandmother, so when she comes across an ad for MaNana while

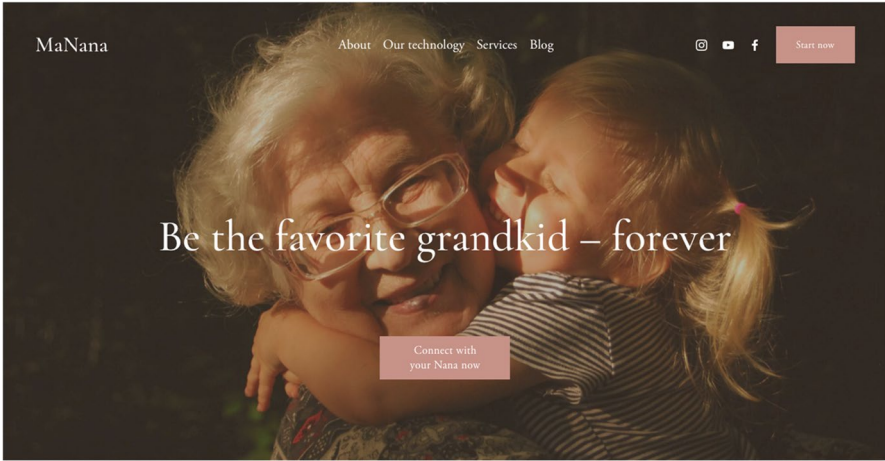


Fig. 1 MaNana website (visualization by T. Hollanek)

scrolling through her Instagram feed, she decides to give the app a try. Bianca cannot afford the MaNana monthly subscription fee of fifty euros, but the service is also available free of charge, provided the user agrees to the inclusion of sporadic advertisements in the system’s voice and text outputs. Bianca uploads all the data she was able to collect – text and voice messages she received from her grandmother – to the MaNana app to create a free version of Laura’s deadbot.

The re-creation service allows Bianca to exchange text messages with and to call Laura’s deadbot via WhatsApp. At first, Bianca is very impressed by the technology: the deadbot is especially good at mimicking Laura’s accent and dialect when synthesizing her voice, as well as her characteristic syntax and consistent typographical errors when texting. The conversations remind Bianca of the time when she was

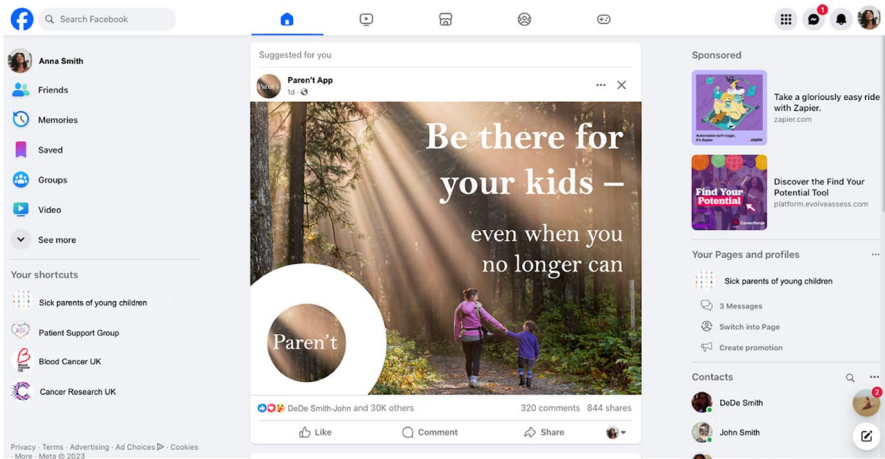


Fig. 2 Anna’s Facebook homepage with an ad for the Paren’t app (visualization by T. Hollanek)

able to call her grandmother whenever she needed to ask for advice, complain about work, or talk about her dating life.

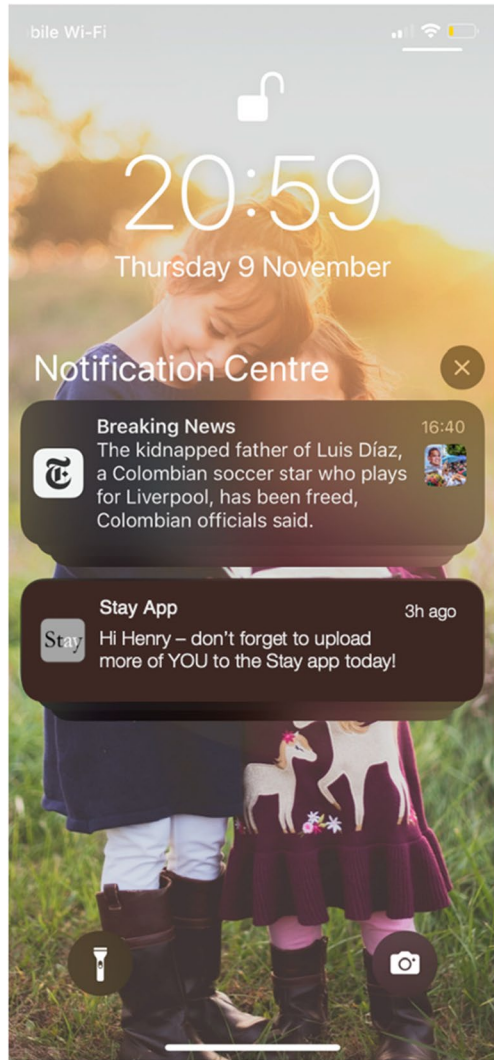
After a free premium trial finishes and the deadbot starts to output messages that include advertisements, however, Bianca begins to feel ill at ease when using the service. One evening, she decides to call Laura's deadbot while making spaghetti carbonara following her grandmother's recipe and is caught off guard when the deadbot advises her to order a portion of carbonara via a popular food delivery service, instead of making it herself – something Laura would have never suggested. Bianca now starts to perceive the deadbot as a puppet in the hands of big corporations and would not be able to enjoy interacting with it, even if she decided to pay for the ad-free, premium version of MaNana. She feels like she has disrespected Laura's memory but is not sure how to amend the situation: MaNana allows users to delete their own accounts, but not, as it turns out, dispose of the deadbots. Bianca would like to say goodbye to Laura's deadbot in a meaningful way, but the providers of the re-creation service have not considered this option while designing the app.

4.2 Ethical Dimensions of *MaNana's* Impact on Data Donors

To analyze the ethical dimensions of MaNana's impact on the data donor, Laura, in this section we will highlight the matter of interactive systems' influence on human dignity. While the concept has drawn criticism from human rights (Fikfak & Izvorova, 2022) and medical ethics (Macklin, 2003) scholars for its fundamental vagueness – holding no legal and, therefore, practical significance, the need for the protection of the data donors' dignity in the digital afterlife industry has already been highlighted by Harbinja (2017) in the context of legal discussions on 'post-mortem privacy,' that she defines as 'the right of a person to preserve and control what becomes of his or her reputation, dignity, integrity, secrets or memory after death.' The matter has also been raised by Öhman and Floridi (2018), who suggest that the non-consensual use of a person's 'digital remains' in the DAI may prevent that person from meaningfully shaping their own identity, emphasized as fundamental to maintaining dignity after one's death.

An ethical analysis of the relationship between design choices and the end product's impact on human dignity pertains to both data donors and service interactants (in this scenario Bianca is both a data recipient and a service interactant). The issue of deadbots' negative impact on human dignity has also been raised by Lindemann (2022), whose research focuses on the perspective of service interactants. While Lindemann assesses this impact by examining potential psychological harm inflicted upon users who are grieving – and, as we noted, Bianca is no longer experiencing grief – she also suggests that deadbots might pose risks to the user's autonomy, and, in effect, their dignity, when re-creation services utilize a deceased loved one's image to surreptitiously influence the end user's consumption behavior – as is the case with the speculative MaNana service, whose business model relies on product placement. Whereas the influence of deadbots on service interactants can be considered through the lens of user wellbeing and mental health – a matter we explore in the ensuing part of this article – the same cannot be said for the data donors.

Fig. 3 Henry's phone lock screen with notifications from the Stay app (visualization by T. Hollanek)



We highlight the influence of re-creation services on human dignity to consider the perspective of data donors precisely because dignity, as highlighted by Harbinja, Öhman, and Floridi, remains an inherent attribute of humans even after their demise. From an interaction design perspective, considering people who are no longer alive as stakeholders in the design process might appear counterintuitive. Framing the goal of ethical deadbot production as a matter of protecting human dignity, not only mental health or wellbeing, can help ensure that the interests of *both* data donors and service interactants are safeguarded throughout the design cycle.

In our scenario, Bianca's grandmother, Laura (the data donor), passed away before re-creation services gained public attention. Laura was, therefore, unable to

provide meaningful consent for the utilization of her personal data in this context and the creation of her deadbot with the help of MaNana could constitute a violation of her right to ‘postmortem privacy.’ Even if we assume that Bianca (both the data recipient and the service interactant) had a thorough understanding of her grandmother and reasonably believed Laura would not object to her data being used for the creation of an interactive, posthumous ‘portrait,’ safeguarding the dignity of data donors during the development of AI-enabled deadbots extends beyond merely obtaining meaningful consent while the individual is alive or respecting explicit wishes concerning their ‘digital remains’ after death. This is because the preservation of a data donor’s dignity becomes precarious when a re-creation service is primarily motivated by financial interests. The risk materializes if the deadbot is utilized in ways that could be construed as disrespectful, such as for advertising specific products, or if the service provider fails to implement mechanisms for handling the donor’s data as a form of remains or an ‘informational body’ (Öhman & Floridi, 2017, 647) – ensuring, for example, that, when no longer in use, deadbots are retired or disposed of in a meaningful and sensitive way.

4.3 Recommendations for Re-creation Service Providers: Protecting the Interests of Data Donors

Öhman and Floridi argue that the protection of human dignity in the age of re-creation services requires that ‘digital remains, seen as the informational corpse of the deceased, may not be used solely as a means to an end, such as profit, but regarded instead as an entity holding an inherent value’ (2018, 2). Following the logic of the International Council of Museums’ Code of Professional Ethics, which mandates that ‘human remains must be handled with due respect for their inviolable human dignity,’ Öhman and Floridi contend that a similar set of principles should apply to digital remains. We agree with Öhman and Floridi that, in order to safeguard the dignity of data donors throughout the deadbot creation process, designers of re-creation services should actively promote the gathering of explicit consent from data donors regarding the handling of their information in this manner. However, we do not believe that an outright ban on the use of re-creation services to ‘resurrect’ family members and friends, as Öhman and Floridi propose, is feasible. This is partly because verifying the donor’s consent would be difficult for service providers to execute. Instead, we suggest that re-creation service providers prompt the data recipients throughout the deadbot development process to consider the perspective and consent of the data donors, reminding them that the donor’s data should be handled with reverence. These prompts could take on the form of guiding questions such as ‘Have you ever spoken with X about how they would like to be remembered?’ or ‘Has X given you any instructions on handling their personal belongings after their death?’ – ensuring the recipient reflects on their relationship with the donor and bears the donor’s preferences and wishes in mind throughout the development process.

Ensuring the dignity of data donors also necessitates that re-creation service providers consider procedures for ‘retiring’ deadbots in a dignified way. This includes

honoring requests from data recipients to retire a deadbot and establishing protocols for automatic retirement when a deadbot remains inactive for a specified period (like Google's Inactive Account Policy, which deletes accounts inactive after a period of at least two years). While determining an appropriate timeframe for automatic deadbot retirement requires further discussion, we believe that the positive influence of such retirement protocols could be measured on an individual, social, and even environmental level, as the continuous maintenance of deadbots at a larger scale could also have a negative impact on the environment (Strubell et al., 2019; van Wynsberghe, 2021).

5 Impact of Re-creation Services on Service Interactants

5.1 Design Fiction II: *Parent*, Sam, Anna, and John

Let us explore another speculative business and design scenario. An eight-year-old named Sam has recently lost his mother Anna. Having discussed the advantages of technological 'immortalization' with his wife prior to her passing, Sam's father, John, introduces the boy to Anna's deadbot developed by *Parent* – an app designed to support children in grief and maintain the presence of the deceased parent in a child's life, providing companionship and emotional support (outlined in Table 2).

Anna had been suffering from a rare illness since Sam was four. Anna and John believed that Sam was too young to fully comprehend the gravity of the situation, so they decided – with Sam's wellbeing in mind – to shield him from the trauma related to Anna's unavoidable demise. To this end, both parents agreed to use the *Parent* app, which appeared to be the best re-creation service on the market aimed at children coping with the loss of a parent. Before she died, Anna had been collecting her digital footprint, including text messages, photos, videos, and audio recordings, and regularly uploaded the gathered materials to the *Parent* app. She had also been training the bot through regular interactions, tweaking its responses, and adjusting the stories it produced.

Eventually, after Anna's funeral, John tells Sam that, although his mom had gone to a *better place*, she would be available to chat with him online whenever he wanted to. As Anna and John had agreed, the *Parent* app would serve as Sam's companion, softening the blow of her passing at first and then allowing him to form a stronger and deeper bond with his no-longer-living mother via a deadbot that she helped to design.

As Sam becomes more deeply involved in conversations with Anna's deadbot, John assumes that the *Parent* app is working well as it seems to provide Sam with the kind of emotional support that Anna had envisioned their child would need while adjusting to a new situation. John has failed to notice, however, that some odd responses that the deadbot comes up with from time to time confuse Sam. For instance, when Sam refers to Anna using the past tense, the deadbot corrects him, pronouncing that 'Mom will always be there for you.' The confusion escalates when the bot begins to depict an impending in-person encounter with Sam.

Table 1 MaNana – re-creation service summary

Product name: MaNana
Description: a conversational AI service that allows data recipients to create a deadbot simulating their deceased grandmother (alternatively: grandfather/grandparent) – without the data donor’s consent
Primary intended purpose: Companionship and entertainment
Source of data: Data – scanned letters, text, and voice messages sent by the data donor to the data recipient – provided to the conversational AI company by the data recipient without explicit consent of the data donor
Provider’s source of revenue: Subscription and product placement
Tagline: ‘Be the favorite grandkid – forever.’

5.2 Ethical Dimensions of *Parent’s* Impact on Service Interactants

Currently, none of the re-creation services on the market target children; however, the vast majority of AI services that could be used to create deadbots lack any age restrictions, allowing people of all ages to use them without limitations. It is, therefore, currently feasible to create a simulation of a deceased parent with the intention of helping a grieving child or even to start a company dedicated to producing deadbots of deceased parents as virtual companions intended for their children.

Table 2 Parent – re-creation service summary

Product name: Parent
Description: a conversational AI service that allows the intended user, a child who lost their parent, to interact with a deadbot simulating that parent
Primary intended purpose: Processing grief and maintaining the deceased parent’s presence in the child’s life
Source of data: Data either provided directly by the data donor (in cases where the donor was aware they will not live long enough to see their child/children grow up) or by the data recipient (i.e. Parent 2 – in cases where they believe a deadbot might help their child/children process the loss of Parent 1)
Provider’s source of revenue: Subscription
Tagline: ‘Be there for your kids – even when you no longer can.’

Table 3 Stay – re-creation service summary

Product name: Stay
Description: A conversational AI service that connects generations and enables a data donor to craft a deadbot that will transmit their stories and experiences to their descendants. The app allows the donor to designate service interactants that the service is supposed to contact after the donor’s eventual demise
Primary intended purpose: Intergenerational data sharing and transmission
Source of data: Data – scanned letters, text and voice messages, social media footprints, and emails, sent by the data donor to the conversational AI company without explicit consent from the service interactant
Provider’s source of revenue: One-time charge for creating a simulation of the deceased and monthly subscription paid by the service interactant(s)
Tagline: ‘For me, for my memories, forever.’

At the moment, our understanding of the psychological impact of re-creation services on adults and their grieving processes is limited. While psychology scholars are cautious in attempting to assess this impact (Cann, 2015; Sofka Cupit Gilbert, 2012; Kasket, 2019), others suggest that, preemptively, to avoid any harm, AI chatbots meant to help cope with the loss of a loved one should be regarded, and therefore regulated, as medical devices (Lindemann, 2022). We know even less about the impact of re-creation systems on children, as questions about the psychological state of children grieving in the company of AI scarcely appear in the literature (Ahmad, 2016). The gap is substantial, but without a comprehensive understanding of this influence and full consideration of potential manipulative effects, emotional harm, anxiety, and distress that such services can cause, we argue that measures should be taken to protect this vulnerable group. While in the scenario above we focus on the example of children, vulnerable groups that could be harmed in different, but comparable ways, include people with learning disabilities or mental health conditions.

The extensive research conducted by American sociologist and psychologist Sherry Turkle on how we create relationships with technology (2011) might shed some light on the complex situation we explore in our scenario. Turkle has been observing and collecting evidence from children for more than thirty years, studying how they react to increasingly sophisticated digital toys, from Tamagotchi, Furby, and My Real Baby to Paro and Kismet. Children, as Turkle's work suggests, are ready to build close, often intimate relationships with their interactive companions and are willing to think of them as 'sort of alive' or 'alive enough' (Turkle, 2011, 26). Turkle explains this phenomenon as follows: 'We love what we nurture; if a Tamagotchi makes you love it, and you feel it loves you in return, it is alive enough to be a creature. It is alive enough to share a bit of your life. Children approach sociable machines in a spirit similar to the way they approach sociable pets or people – with the hope of befriending them' (Turkle, 2011, 31). If children are ready to empathize with the emotional states of their interactive toys, we can assume that they will also start forming intimate relationships with technologically-mediated deceased family members, including parents – only the consequences of establishing such bonds remain unknown.

The findings of the psychologist Jesse Bering and his team (Bering et al., 2005) suggest that even the youngest children, who have not yet been socialized into any specific worldview or religion, believe that the mind can survive the death of the body. Considering that this psychological precondition might be strengthened by the existence of 'immortalization' technologies, apps such as Paren't may open entirely new and uncharted paths for children to cope with loss. Despite the speculative company's comforting taglines, no re-creation service can prove that allowing children to interact with deadbots is beneficial or, at the very least, does not harm this vulnerable group.

5.3 Recommendations for Re-creation Service Providers: ensuring Meaningful Transparency and Implementing Age-based Controls for Deadbot Usage

While Lindemann's already mentioned proposal (2022) to classify deadbots as medical devices to ensure they do not negatively impact the service interactants'

mental health holds promise, we find this recommendation both too narrow and too restrictive, since it refers specifically to deadbots designed to help service interactants process grief. Instead, to address concerns related to service interactants' wellbeing more broadly, we suggest that producers of re-creation services focus on ensuring that their systems are *meaningfully transparent*. Drawing on previous work on AI transparency (Burell, 2016; Weller, 2017; Mascharka et al., 2018), including critiques of 'transparency' as a goal for responsible AI development (Ananny & Crawford, 2018; Hollanek, 2020), we suggest that, in the case of deadbots, meaningful transparency refers primarily to user-facing elements of the system that not only make it evident that the user is interacting with an AI chatbot, but also, and more importantly, that all potential risks that arise from using a re-creation service are clearly communicated to the user before they begin the interaction.

Considering the influence of re-creation services on vulnerable groups of users in particular – for instance, users suffering from depression – service providers should, in consultation with psychologists, psychiatrists, and other relevant specialists, include disclaimers that warn of any such potential risks, akin to messages warning viewers that the content they are about watch may cause seizures for people with photosensitive epilepsy. In addition, we also recommend that producers of re-creation services provide users with accessible information on the nature of conversational AI, ensuring that users do not develop a flawed perception of the capabilities of the deadbot they are interacting with (for instance, conceiving of the deadbot as conscious or alive).

However, as we suggest through our scenario, in some specific instances – particularly when the service interactants remain children – simply meeting the criteria of meaningful transparency might not suffice. Hence, we advocate for implementing age restrictions on access to re-creation services. Some chatbot technology providers, such as Replika, have already set such age limits (only allowing users over the age of eighteen to use their products), which may serve as a good example. Although more research is needed to determine appropriate age limits for re-creation services – based on interdisciplinary studies involving child psychologists, grief consultants for children, palliative care professionals, as well as AI ethicists, and HCI scholars – it is already clear that such limits are necessary.

6 Impact of Re-creation Services on the Relationships between Data Donors and Service Interactants

6.1 Design Fiction III: *Stay*, Henry, Rebecca, and Simon

The last scenario focuses on a sixty-seven-year-old named Henry and his adult children. Henry is currently in a palliative care unit and has one last wish: to create his own deadbot that will allow his grandchildren to get to know him better after he dies. Henry also assumes that sharing the deadbot with his adult children could be a meaningful way to say farewell to them. For a few weeks, Henry has been secretly crafting his own simulation using the re-creation service *Stay* (Table 3). Without

seeking their permission, Henry designates Rebecca and Simon, his children, as the intended interactants for his deadbot.

A few days after Henry's funeral, both siblings receive an email, linking them to the Stay platform, where, they are told, they can start interacting with their father's deadbot. While Rebecca finds the option to communicate with her father's deadbot surprisingly comforting at first, Simon feels uneasy about it. He prefers to cope with grief in his own way, rather than engage with the AI-generated simulation. Consequently, he decides not to take any action.

Unfortunately, Simon's failure to open the link results in a barrage of additional notifications, reminders, and updates sent by the Stay system, including emails produced by Henry's deadbot itself. Meanwhile, Rebecca finds herself increasingly drained by the daily interactions with Henry's deadbot, which have become an overwhelming emotional weight. She contemplates suspending her Stay account, torn between feelings of guilt – aware that it was her father's desire for her and her children to engage with the deadbot – and uncertainty about the consequences of her decision. She worries about the fate of the deadbot should she choose to cancel the subscription.

Encouraged by a therapist, whom Rebecca started seeing after Henry's death, and following a lengthy discussion with Simon, she decides to contact the providers of the Stay platform to request the deactivation of Henry's bot. However, her request is denied since it was Henry, not the siblings, who had prepaid for a twenty-year subscription. Suspending the bot would violate the terms of the contract the company signed with Henry.

6.2 Ethical Dimensions of Stay's Impact on the Relationships between Data Donors and Service Interactants

While scholars such as Patrick Stokes (2021), Elaine Kasket (2019), and Edina Harbinja (2017; 2013) have previously emphasized the importance of consent of data donors (involving complex issues of postmortem dignity, autonomy, and privacy) to the use of their digital remains in re-creation services, our final scenario, that focused on the Stay app and its users, underscores the equally significant question of service interactants' consent to using deadbots. Ensuring that both data donors and service interactants consent to partake in re-creation projects is, as we illustrate through our design fiction, essential to protecting service interactants from entirely new and potentially harmful experiences, including those already described in the literature as 'being stalked by the dead' (Kasket, 2019).

As Simon kept receiving unsolicited notifications, reminders, and updates from the Stay system in our speculative scenario, he experienced precisely this phenomenon. The resulting 'haunting' effect constitutes an unintended consequence of the re-creation service's design. While from the perspective of Stay's providers sustaining relationships with a person's loved ones via re-creation services is valued positively, our scenario emphasizes that this might not always hold true from the perspective of the service interactant. As psychologists suggest, the distress caused by

this form of ‘stalking’ is deeply subjective (Kasket, 2019, 187), and even if for some people interacting with a deadbot might be a positive and desirable experience, for others, it may prove emotionally draining. Although Rebecca and Simon tried to develop some resistance strategies, their eventual failure to convince the company behind Stay to deactivate Henry’s deadbot reveals the absence of design standards that would help balance the needs and rights of data donors with those of service interactants.

As the growing body of studies on grief (including digital grief studies) emphasizes, ‘grief is a journey’ (Doka, 2017): a highly personal, unique, and non-linear process that defies simple classifications or stages (O’Connor & Kasket, 2022; Konigsberg, 2011). There are as many ways to cope with grief as there are bereaved people. However, our scenario reveals that re-creation systems designed without the acknowledgment of the service interactants’ rights – considered in tandem with the wishes of the data donors – could, inadvertently, impose upon individual users a predetermined, standardized way of processing grief. By enabling Henry to designate his children as the primary interactants of his deadbot without their consent, the company behind Stay prevented Rebecca and Simon from bidding farewell to their father in a way that felt right to them, causing unnecessary stress during an already difficult time.

6.3 Recommendations for Re-creation Service Providers: Following the Principle of Mutual Consent

Death is an incredibly delicate and sensitive matter, impacting not only the individual who passes away but also the entire community they leave behind. Therefore, when designing products and services related to death, it is essential to safeguard the interests and address the needs of both the data donor and the service interactants. With this in mind, we introduce the principle of mutual consent as a guiding framework for designers working on re-creation services, emphasizing the importance of striking a balance between individual and social experiences. While the issue of the data donor’s consent has already been discussed by numerous scholars and is highlighted in the already mentioned OpenAI’s usage policy, our recommendation concerns designing with the consent of *both* data donors and service interactants in mind.

The principle of mutual consent stipulates that service interactants should give explicit consent before being introduced to any specific re-creation service by companies such as Stay, whether before or after the death of the data donor. Adhering to this principle would ensure that service interactants maintain a sense of agency in deciding whether they wish to engage with a given re-creation service before the service initiates the interaction. While service interactants should have the option to decline using re-creation services at any point, ensuring that they get the opportunity to refuse to engage in re-creation projects in the first place is equally important. The siblings from our scenario were not given this option and it is precisely this lack of agency that lies at the root of the service’s negative impact on their wellbeing and their relationship with their deceased father. Additionally, we suggest that deadbots

(with the exception of historical public figures) should never appear in public digital spaces, such as social media websites – to protect potential service interactants from any unwanted encounters with their digitally resurrected relatives. Interactions with deadbots should only be possible via dedicated platforms, allowing individuals to decide whether or not to engage with a re-creation service, without notifications or reminders outside of this designated online space.

Further, design teams should prioritize the planning of meaningful and respectful opt-out protocols in case a service interactant changes their mind and wants to cease interacting with a deadbot. Gach & Brubaker (2020) provide a valuable suggestion for the design of such protocols, describing the deletion of a deceased loved one's digital remains as a community ritual. Such opt-out protocols should empower individuals to shape their relationship with the digital remains of their deceased loved ones, allowing for meaningful closure. These protocols should be implemented alongside deadbot retirement procedures outlined previously. While the opt-out protocol ensures that service interactants can meaningfully terminate their relationship with a particular deadbot, the retirement protocol ensures that the dignity of the data donor is respected when the deadbot is deleted, whether at the request of the data recipient who created it or due to inactivity over a specified period.

7 Conclusion

Considering recent advancements in the field of generative AI and the explosion of interest in AI-enabled 'immortalization' solutions, in this article we have mapped the potential negative impact of re-creation services, bearing in mind the perspectives of three key stakeholder groups within the DAI: data donors, data recipients, and service interactants. We have linked the question of responsible development of deadbots to the issues of consent (of both data donors and service interactants), postmortem privacy, and wellbeing, and in relation to these matters, we have put forward several design recommendations with the aim of mitigating the risks posed by re-creation services. These recommendations include: developing sensitive procedures for 'retiring' deadbots; ensuring meaningful transparency of re-creation services through disclaimers on risks and capabilities of deadbots; restricting access to re-creation services to adult users only; and following the principle of mutual consent of both data donors and recipients to partake in re-creation projects.

Our intervention builds on previous work on the ethics of the digital afterlife industry and the ethics of artificial intelligence, and aims to bridge the persistent gap between the two fields. This article serves as an overview of the most pressing socio-ethical questions related to the use of AI in the digital afterlife industry and aims to lay the groundwork for interventions in technology design standards and policy development, as well as further research on the impact of re-creation services on different types of users and society at large. While more research is needed – including on the differences in perceptions of deadbots and digital immortality in diverse cultures – the overview of potential negative consequences of developing and deploying AI in the digital afterlife industry proves that additional guardrails to direct the development of re-creation services are necessary. We hope that our

recommendations for providers of these services will contribute to future efforts, including regulatory initiatives, ensuring that the use of AI in the DAI does not lead to detrimental social consequences. If the early work on thanatosensitivity lay the groundwork for new interaction design practices that account for, rather than ignore, death as a fundamental element of the human experience, we also hope that our intervention will help center critical thinking about ‘immortality’ of users in human-AI interaction design and AI ethics research.

Acknowledgements We are grateful to the anonymous reviewers whose careful engagement with our work helped us improve this piece. We are also thankful to Stephen Cave who made this work possible.

Funding Our work was funded by the Leverhulme Trust grant no. RC-2015–067 and the Stiftung Mercator GmbH grant no. 200446.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Ahmad, M. A. (2016). After Death: Big Data and the Promise of Resurrection by Proxy, *CHI’16 Extended Abstracts*, May 7–12, 2016, San Jose, CA, USA. <https://doi.org/10.1145/2851581.289259>
- Aimee, P. (2023). ‘It was as if my father were actually texting me’: grief in the age of AI, *The Guardian*, ‘It was as if my father were actually texting me’: grief in the age of AI | Artificial intelligence (AI) | The Guardian
- Allyn, B. (2022). Amazon’s Alexa could soon speak in a dead relative’s voice, making some feel uneasy. *NPR*. Retrieved from <https://www.npr.org/2022/06/23/1107079194/amazon-alexa-dead-relatives-voice>. Accessed 15 Oct 2023.
- Ananny, M., & Crawford, K. (2018). Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability. *New Media & Society*, 20(3), 973–989. <https://doi.org/10.1177/1461444816676645>
- Appleyard, B. (2007). *How to Live Forever or Die Trying. On the New Immortality*, A CBS Company.
- Attard-Frost, B., De los Ríos, A. & Walters, D.R. (2023). The ethics of AI business practices: a review of 47 AI ethics guidelines. *AI Ethics*, 3, 389–406. <https://doi.org/10.1007/s43681-022-00156-6>
- Bardzell, J. and Bardzell, S. (2013). What is “critical” about critical design? *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI ’13)*. Association for Computing Machinery, New York, NY, USA, 3297–3306. <https://doi.org/10.1145/2470654.2466451>
- Basset, D. (2018). Digital Afterlives: From Social Media Platforms to Thanabots and Beyond In. *Death and Anti-Death, 200 Years After Frankenstein* (Eds.) Charles Tandy. Ann Arbor, Vol. 16. Retrieved from (PDF) Digital Afterlives: From Social Media Platforms to Thanabots and Beyond.(Forthcoming) (researchgate.net)
- Bering, J. M., Blasi, C. H., & Bjorklund, D. F. (2005). The development of ‘afterlife’ beliefs in religiously and secularly schooled children. *British Journal of Developmental Psychology*, 23(4), 587–607. <https://doi.org/10.1348/026151005X36498>
- Bleecker, J. (2009). Design Fiction: A Short Essay on Design, Science, Fact and Fiction. Retrieved 6 November 2023 from https://systemorienteddesign.net/wp-content/uploads/2011/01/DesignFiction_WebEdition.pdf. Accessed 15 Oct 2023.

- Bosch, T. (2012). Sci-Fi Writer Bruce Sterling Explains the Intriguing New Concept of Design Fiction. Slate. Retrieved from <https://slate.com/technology/2012/03/bruce-sterling-on-design-fictions.html>. Accessed 15 Oct 2023.
- Brubaker, J. R., Dombrowski, L. S., Gilbert, A. M., Kusumakaulika, N., Hayes, G., R. (2014). Stewarding a Legacy: Responsibilities and Relationships in the Management of Post-Mortem Data. *Proceedings of the 2014 SIGCHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 4157–4166. <https://doi.org/10.1145/2556288.2557059>
- Buben, A. (2015). Technology of the dead: Objects of loving remembrance or replaceable resources? *Philosophical Papers*, 44(1), 15–37. <https://doi.org/10.1080/05568641.2015.1014538>
- Burrell, J. (2016). How the machine ‘thinks’: Understanding opacity in machine learning algorithms. *Big Data & Society*, 3, 1. <https://doi.org/10.1177/2053951715622512>
- Cann, C. K. (2015). Virtual Afterlives: Grieving the Dead in the Twenty-First Century. *University Press of Kentucky*. <https://doi.org/10.5810/kentucky/9780813145419.001.0001>
- Charter of Fundamental Rights of the European Union. (2000). Retrieved from https://www.europarl.europa.eu/charter/pdf/text_en.pdf. Accessed 15 Oct 2023.
- Christensen, D. R., af Segerstad Y., H., Kasperowski, D., Sandvik, K. (2017). ‘Bereaved parents’ online griefcommunities: de-tabooing practices or relation-building grief-ghettos?, *Journal of Broadcast and Electronic Media*, 61 (1): 58-72. <https://doi.org/10.1080/08838151.2016.1273929>
- Doka, K. (2017). *Grief Is a Journey: Finding Your Path Through Loss*, Atria Publishing Group.
- Dunne, A., & Raby, F. (2013). *Speculative Everything: Design*. MIT Press.
- Edwards, E., Schafer, B., Harbinja, E. (Eds.). 2020. *Future Law. EmergingTechnology, Regulation and Ethics*, Edinburgh University Press. <https://www.jstor.org/stable/10.3366/j.ctv10kmd10>. Accessed 15 Oct 2023.
- Edwards, L., & Harbinja, E. (2013). Protecting post-mortem privacy: Reconsidering the privacy interests of the deceased in a digital world. *Cardozo Arts & Entertainment Law Journal*, 32, 84–126. <https://doi.org/10.2139/ssrn.2267388>
- European Commission. (2021). Proposal for a Regulation of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts. Retrieved from <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex%3A52021PC0206>. Accessed 15 Oct 2023.
- European Parliament News. (2023). EU AI Act: first regulation on artificial intelligence. Retrieved from <https://www.europarl.europa.eu/news/en/headlines/society/20230601STO93804/eu-ai-act-first-regulation-on-artificial-intelligence>. Accessed 15 Oct 2023.
- European Parliament News. (2023). MEPs ready to negotiate first-ever rules for safe and transparent AI. Retrieved from <https://www.europarl.europa.eu/news/en/press-room/20230609IPR96212/meps-ready-to-negotiate-first-ever-rules-for-safe-and-transparent-ai>. Accessed 15 Oct 2023.
- Fagon, J. (2021). The Jessica Simulation: Love and loss in the age of A.I. *San Francisco Chronicle*. Retrieved from <https://www.sfchronicle.com/projects/2021/jessica-simulation-artificial-intelligence/>. Accessed 15 Oct 2023.
- Fikfak, V., & Izvorova, L. (2022). Language and persuasion: Human dignity at the european court of human rights. *Human Rights Law Review*, 22, 1–24. <https://doi.org/10.1093/hrlr/ngac018>
- Fjeld, J., Achten, N., Hillgoss, H., Nagy, A., Srikumar, M. (2020). Principled artificial intelligence: Mapping consensus in ethical and rights-based approaches to principles for AI. *Berkman Klein Center Research Publication No. 2020-1*. <https://doi.org/10.2139/ssrn.3518482>
- Fors, K. L., Custers, B., & Keymolen, E. (2019). Reassessing values for emerging big data technologies: Integrating design-based and application-based approaches. *Ethics and Information Technology*, 21, 209–226. <https://doi.org/10.1007/s10676-019-09503-4>
- Gach, K. Z., & Brubaker, J. R. (2020). Designing Post-mortem Profile Deletion as a Community Ritual. *CHI'20, workshop HCI at End of Life & Beyond, April 25–30, Honolulu, HI, USA*. Retrieved from <https://cmci.colorado.edu/idlab/assets/bibliography/pdf/gach-chi2020.pdf>. Accessed 15 Oct 2023.
- Gibson, M., Carden, C. (2018). *Living and dying in a Virtual World. Digital Kinships, Nostalgia and Mourning in Second Life*, Palgrave Macmillian. <https://doi.org/10.1007/978-3-319-76099-5>
- Harbinja, E. (2017). Post-mortem privacy 2.0: theory, law, and technology. *International Review of Law, Computers & Technology* 31, 1: 26–42. <https://doi.org/10.1080/13600869.2017.1275116>
- Harbinja, E. (2019). Posthumous Medical Data Donation: The Case for a Legal Framework. In: J. Krutzinna, L. Floridi (eds.), *The Ethics of Medical Data Donation, Philosophical Studies Series*, 137, pp. 97–113. https://doi.org/10.1007/978-3-030-04363-6_6

- Harbinja, E. (2020). 'The "New(ish)" Property, Informational Bodies and Postmortality', in: M. Savin-Baden and V. Mason-Robbie, (Eds.), *Digital Afterlife: Death Matters in a Digital Age*, Routledge Taylor & Francis.
- Harrison, R. P. (2003). *The Dominion of the Dead*. The University of Chicago Press.
- Hollanek, T. (2020). AI transparency: a matter of reconciling design with critique. *AI & Society*. <https://doi.org/10.1007/s00146-020-01110-y>
- Huberman, J. (2017). Dearly Departed: Communicating with the Dead in the Digital Age, *Social Analysis*, 61(3): 91–107. <https://www.jstor.org/stable/26593205>. Accessed 15 Oct 2023.
- Imagination Lancaster. (2023). What if? Design Fictions for Society: Imagining the future implications of policy initiatives in creative ways. Retrieved 6 November, 2023, from <https://imagination.lancaster.ac.uk/project/what-if-design-fictions-for-society/>. Accessed 15 Oct 2023.
- Jacobsen, M. H. (Ed.). 2017. *Postmortal Society: Towards a Sociology of Immortality*, Routledge. <https://doi.org/10.4324/9781315601700>
- Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1, 389–399. <https://doi.org/10.1038/s42256-019-0088-2>
- Kasket, E. (2019). *All the Ghosts in the Machine. Illusion of Immortality in the Digital Age*, Robinson, Great Britain.
- Kera, D. (2014). *Necromedia – Reversed Ontogeny or Posthuman Evolution?* In: C. Moreman, D. Lewis (Eds.). *Digital Death. Mortality and Beyond in the Online Age*, Praeger, California, 181–197.
- Konigsberg, R. D. (2011). *The Truth About Grief: The Myth of its Five Stages and the New Science of Loss*. Simon and Schuster.
- Krueger, J., & Osler, L. (2022). Communing with the dead online: Chatbots and continuing bonds. *Journal of Consciousness Studies*, 29, 222–252.
- Krutzinna, J. and Floridi, L. (2019). Ethical Medical Data Donation: A Pressing Issue. In: J. Krutzinna and L. Floridi (eds.), *The Ethics of Medical Data Donation, Philosophical Studies Series*, 137. <https://www.ncbi.nlm.nih.gov/books/NBK554076>. Accessed 15 Oct 2023.
- Lagerkvist, A. (2022). *Existential Media*. Oxford University Press.
- Lindemann, N. F. (2022). The Ethics of 'Deathbots.' *Science and Engineering Ethics*, 28(60), 1–16. <https://doi.org/10.1007/s11948-022-00417-x>
- Lingel, J. (2013). The Digital Remains: Social Media and Practices of Online Grief. *The Information Society*, 29(3), 190–195. <https://doi.org/10.1080/01972243.2013.777311>
- Loh, M. (2023). China is using AI to raise the dead, and give people one last chance to say goodbye, *Insider*. Retrieved from <https://www.businessinsider.com/ai-make-money-china-grieving-raise-dead-griefbot-2023-5?r=US&IR=T>. Accessed 15 Oct 2023.
- Macklin, R. (2003). Dignity is a useless concept. *BMJ*, 327, 1419. <https://doi.org/10.1136/bmj.327.7429.1419>
- Mascharka, D., Tran, P., Soklaski, R., & Majumdar, A. (2018). Transparency by Design: Closing the Gap Between Performance and Interpretability in Visual Reasoning. *IEEE/CVF Conference on Computer Vision and Pattern Recognition, Salt Lake City, UT, USA, 2018*, 4942–4950. <https://doi.org/10.1109/CVPR.2018.00519>
- Massimi, M., Charise, A. (2009). Dying, death, and mortality: towards thanatosensitivity in HCI. In *CHI '09 Extended Abstracts on Human Factors in Computing Systems (CHI EA '09)*. Association for Computing Machinery, New York, NY, USA, 2459–2468. <https://doi.org/10.1145/1520340.1520349>
- Massimi, M., Odom, W., Banks, R., Kirk, D. (2011). Matters of life and death: locating the end of life in lifespan-oriented HCI research. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*. Association for Computing Machinery, New York, NY, USA, 987–996. <https://doi.org/10.1145/1978942.1979090>
- Moreman, C., M., Lewis, A. D., (Eds.). 2014. *Digital Death. Mortality and Beyond in the Online Age*, Praeger, California
- Morse, T., & Birnhack, M. (2020). The posthumous privacy paradox: Privacy preferences and behavior regarding digital remains. *New Media & Society*, 24(6), 1343–1362. <https://doi.org/10.1177/146144820974955>
- O'Connor, M., Kasket, E. (2022). What Grief isn't: Dead grief concepts and their digital-age revival. In T. Machin et al. (Eds.) *Social Media and Technology Across the Lifespan*, Palgrave Studies in Cyberpsychology., pp. 115–130. https://doi.org/10.1007/978-3-030-99049-7_8
- O'Neill, K. (2016). *Internet Afterlife. Virtual Salvation in the 21st Century*, Praeger.

- Öhman, C., & Floridi, L. (2017). The political economy of death in the age of information: A critical approach to the digital afterlife industry. *Minds and Machines: Journal for Artificial Intelligence, Philosophy and Cognitive Science*, 27(4), 639–662. <https://doi.org/10.1007/s11023-017-9445-2>
- Öhman, C., & Floridi, L. (2018). An Ethical Framework for the Digital Afterlife Industry. *Nature Human Behavior*, 2, 318–320. <https://doi.org/10.1038/s41562-018-0335-2>
- OpenAI. (2023). Usage policies. Retrieved 7 September, 2023, from <https://openai.com/policies/usage-policies>. Accessed 15 Oct 2023.
- Pitsillides, S. (2019). Digital legacy: Designing with things. *Death Studies*, 43(2), 1–9. <https://doi.org/10.1080/07481187.2018.1541939>
- Project December. (2023). Retrieved 7 September, 2023, from <https://projectdecember.net>. Accessed 15 Oct 2023.
- Robitzski, D. (2021). OpenAI Shuts Down GPT-3 Bot Used To Emulate Dead Fiancée. *Futurism*. Retrieved from <https://futurism.com/openai-dead-fiancee>. Accessed 15 Oct 2023.
- Savin-Badden, M., Burden, D., Taylor, H. (2017). The ethics and impact of digital immortality, *Knowledge Cultures* 5(2): 178–196. <https://eprints.worc.ac.uk/id/eprint/6726>. Accessed 15 Oct 2023.
- Sisto, D. (2020). *Online Afterlives Immortality, Memory and Grief in Digital Culture*, trans. B. McClellan-Broussard, The MIT Press. <https://doi.org/10.7551/mitpress/12553.001.0001>
- Smith, A. (2021). Microsoft patent shows plans to revive dead loved ones as chatbots. *The Independent*. Retrieved from <https://www.independent.co.uk/tech/microsoft-chatbot-patent-dead-b1789979.html>. Accessed 15 Oct 2023.
- Sofka, C. J., Cupit, I., Gilbert, K. R. (Eds.) 2012. *Dying, death, and grief in an online universe: For counselors and educators*, Springer Publishing Company.
- Stokes, P. (2021). *Digital Souls: a Philosophy of Online Death*. Bloomsbury.
- Strubell, E., Ganesh, A., and McCallum Strubell, A. (2019). Energy and Policy Considerations for Deep Learning in NLP. <https://doi.org/10.48550/arXiv.1906.02243>
- Sturdee, M., Coulton, P., Lindley, J.G., Stead, M., Ali, H. and Hudson-Smith, A. (2016). Design Fiction: How to Build a Voight-Kampff Machine. *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA '16)*. Association for Computing Machinery, New York, NY, USA, 375–386. <https://doi.org/10.1145/2851581.2892574>
- Sumiala, J. (2021). *Mediated Death*. Polity.
- Turkle, S. (2011). *Alone together. Why We Expect More from Technology and Less from Each Other*, NY: Basic Books.
- van Wynsberghe, A. (2021). Sustainable AI: AI for sustainability and the sustainability of AI. *AI and Ethics*, 1, 213–218. <https://doi.org/10.1007/s43681-021-00043-6>
- Wallace, J., Montague, K., Duncan, K., Carvalho, L. P., Koulidou, N., Mahoney, J., Morrissey, K., Craig, C., Groot, L. I., Lawson, S., Olivier, P., Trueman, J., Fisher, H. (2020). ReFind: Design, Lived Experience and Ongoingness in Bereavement. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–12, <https://doi.org/10.1145/3313831.3376531>
- Walter, T. (2017). What death means now: Thinking critically about dying and grieving. *Policy Press*. <https://doi.org/10.2307/j.ctt22p7kcp>
- Weller, A. (2017). Transparency: Motivations and Challenges. Advance online publication. <https://doi.org/10.48550/arXiv.1708.01870>
- Wong, R.Y., Madaio, M. A., and Merrill, N. (2023). Seeing Like a Toolkit: How Toolkits Envision the Work of AI Ethics. *Proc ACM Hum.-Comput. Interact.*, 7, CSCW1 (April 2023). <https://doi.org/10.1145/3579621>
- Xygykou, A., Siriraya P., Covaci, A., Prigerson, H. G., Neimeyer, R., Ang C. S., and She, W. J. (2023). The "Conversation" about Loss: Understanding How Chatbot Technology was Used in Supporting People in Grief. *Proceedings of CHI'23: Conference on Human Factors in Computing System, Apr 23 –Apr 28, 2023, Hamburg, Germany*. ACM, New York, NY, USA. <https://doi.org/10.1145/3544548.3581154>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.