



EPIC-PP-2026-011 · RESEARCH ARTICLE · AI GOVERNANCE

Words They Never Said: Generative Griefbots, Posthumous Attribution, and an Islamic Ethics of Speaking for the Dead

EPIC – Emirati Perspectives and Insight Centre

Words They Never Said: Generative Griefbots, Posthumous Attribution, and an Islamic Ethics of Speaking for the Dead

EPIC – EMIRATI PERSPECTIVES AND INSIGHT CENTRE

EPIC-PP-2026-011 · AI Governance · Published 20 October 2026

© 2026 EPIC – Emirati Perspectives and Insight Centre. Released under Creative Commons Attribution 4.0 International (CC BY 4.0): <https://creativecommons.org/licenses/by/4.0/>

EPIC is an independent research initiative in formation. It is not yet a registered or licensed entity. EPIC does not accept donations, payments or contracts.

Abstract

Chatbots built from the messages, voice recordings and posts of people who have died are now sold as aids to mourning. Most ethical discussion of these "griefbots" asks whether they help or harm the bereaved, or whether the deceased consented to the use of their data. This article argues that generative griefbots raise a further and more basic problem, which it calls the attribution problem. A system built on a large language model does not replay what a person said. It produces new sentences, in that person's voice, that the person never said. Each output is a fresh attribution of speech to someone who can no longer confirm, correct or withdraw it. The article develops this problem through the literature on posthumous harm and tests it against the fictionalist reply that users know the bot is not the person. It then turns to a resource that the field has largely ignored: Islamic ethical and juristic thought about speaking for the dead. A 2025 ruling by Egypt's Dar al-Ifta permits AI animation of the deceased for remembrance but forbids making them say what they never said. Read alongside the hadith tradition's long discipline of verified attribution, this ruling points to a practical standard. The article proposes "attribution-bounded" design: quotation before generation, visible provenance for every reply, a ban on new commitments in the dead person's name, and mourning-sensitive limits on use. It closes with implications for regulation in the Gulf and beyond.

Keywords: Griefbots; Deathbots; Posthumous harm; Attribution; Islamic ethics; Generative AI

1 Introduction

In 2021 a newspaper feature described a man who, eight years after his fiancée died, gave her old text and social media messages to an experimental chatbot service built on a large language model and held long, emotional conversations with the result (Fagone, 2021; Henrickson, 2023). The system did not search an archive of her words and return

them. It wrote new replies, in a style close enough to hers that the exchanges felt, to him, like talking with her again. The story became a reference point for a new genre of product. Companies now offer "griefbots", "deathbots" or "postmortem avatars": conversational agents trained on the digital remains of a particular dead person and offered to those who mourn them (Hollanek & Nowaczyk-Basińska, 2024). A patent granted to Microsoft in 2020 describes how "social data" such as images, voice data, posts, messages and written letters can be used to train a chatbot "to converse in the personality of the specific person", and notes that the specific person may be a past entity, a friend or a relative (Abramson & Johnson, 2020).

The philosophical literature on these systems has grown quickly. It has asked whether they serve or obstruct healthy grieving (Elder, 2020; Krueger & Osler, 2022; Fabry & Alfano, 2024), whether they treat the dead as replaceable resources (Buben, 2015), whether the digital afterlife industry exploits the remains of its users for profit (Öhman & Floridi, 2017, 2018), and what rules of design and consent responsible providers should follow (Lindemann, 2022; Hollanek & Nowaczyk-Basińska, 2024). These are serious questions, and this article builds on the answers already given. Its claim, though, is that the debate has not yet isolated the feature that makes generative griefbots morally distinctive.

That feature is generation itself. A photograph, a voicemail or a scrapbook preserves what a person did or said. Even an early rule-based chatbot that answers with stored sentences remains, in a sense, a curated archive. A system built on a large language model is different in kind. It produces sentences that its source never uttered, on topics its source may never have considered, in response to situations that arose after the source's death. Each such sentence is offered in the dead person's voice. The bereaved user, and anyone to whom the user shows the exchange, is invited to hear it as something the person "would have said". We will call the moral difficulty that follows the attribution problem: generative griefbots attribute new speech to people who cannot confirm, correct or withdraw it.

The article makes three moves. Section 3 sets out what generative griefbots do and how existing ethical work has framed them. Section 4 develops the attribution problem, draws on the philosophy of posthumous harm to show why it matters, and considers the strongest reply, namely that users adopt a fictional stance towards the bot. Section 5 turns to a resource the field has largely overlooked: Islamic ethical and juristic reflection on speaking for the dead. A ruling issued by Egypt's Dar al-Ifta in January 2025 addresses AI-generated videos of deceased people directly and draws the line exactly where the attribution problem lies. It permits animation for remembrance and comfort, and forbids making the dead appear to say what they never said (Dar al-Ifta al-Misriyyah, 2025). Read alongside the hadith tradition's long discipline of verified attribution, the ruling suggests a standard that is both principled and practical. Section 6 turns that standard

into design requirements for what we call attribution-bounded griefbots. Section 7 draws out implications for regulation, with attention to the Gulf, where Arabic-language generative systems are being built at national scale and where religious norms about the dead are widely shared. Section 8 answers objections, and Section 9 concludes.

The argument is not that griefbots should be banned. Many people already keep and reread the messages of those they have lost, and there is good reason to think that technologies of remembrance can support what grief researchers call continuing bonds (Klass et al., 1996). The argument is rather that the move from preserving speech to producing speech changes the moral character of the product, and that design and regulation should track that change.

2 Method

This is a work of applied philosophy. It proceeds by conceptual analysis of a technology and by critical engagement with published ethical, legal and religious sources. No human participants were involved and no new data were collected. Religious sources are cited from published English translations and from the official English text of the Dar al-Ifta ruling; where the article interprets them, the interpretation is offered as an argument, not as a juristic opinion.

3 What Generative Griefbots Do

3.1 From archive to voice

It helps to distinguish three kinds of artefact that are often grouped together. The first is the archive: a store of a person's photographs, letters, voice notes and posts, which the bereaved may browse. The second is the scripted or retrieval-based agent, which answers a user's input by selecting from material the person actually produced, perhaps with light rephrasing. The third is the generative agent, which uses a statistical model of language, adapted to a person's style and content, to compose new replies.

The Microsoft patent illustrates the third kind clearly. It describes a system in which a person's social data are used to build a "personalized chat index", and it anticipates that the data will not cover every question a user might ask. In that case the system may draw on the conversational data of other people or on generic sources to "reply to user dialogue and/or questions for which the social data does not provide data" (Abramson & Johnson, 2020). In other words, when the person's own words run out, the system fills the gap with words that are not theirs but are delivered as if they were. The patent also anticipates synthetic voice and two- or three-dimensional likenesses. The chatbot described

in the 2021 newspaper feature was built on a general-purpose language model, so the generation was not a fallback; it was the whole mechanism (Fagone, 2021; Henrickson, 2023).

The distinction matters because the ethical vocabulary that suits archives does not transfer cleanly to generators. We ask of an archive whether it is accurate, complete and appropriately accessed. We ask of a generator something new: on whose authority does it speak?

3.2 The existing debate

Öhman and Floridi (2017) were among the first to analyse the "digital afterlife industry" as a political economy in which the remains of the dead become a commercial resource. In a later framework they proposed treating digital remains by analogy with the human body in museum and archaeological ethics, governed by a principle of human dignity (Öhman & Floridi, 2018). Their framework already contains an attribution-sensitive element: it warns against radical alteration of the deceased's informational remains. But it was written before large language models made generation cheap and convincing.

A second strand asks whether griefbots help mourners. Elder (2020) draws on classical Confucian debates over mourning ritual, including the use of impersonators of the dead, and concludes that chatbots of the dead can be a valuable resource if they are embedded in social practices that guard against fixation. Krueger and Osler (2022) argue that griefbots can support continuing bonds through "habits of intimacy" and that worries about replacement are overstated, because users enjoy only a "thin reciprocity" with the bot and typically adopt a fictionalist stance towards it. Fabry and Alfano (2024) describe deathbots as affective scaffolds that shape the emotions of the bereaved, with both promise and risk. Henrickson (2023) places "thanabots" in a long lineage of attempts to communicate with the dead.

A third strand is more cautious. Buben (2015) asks whether interactive simulations of the dead preserve them as objects of loving remembrance or reduce them to replaceable resources. Lindemann (2022) shifts attention from the dignity of the deceased to the dignity and autonomy of bereaved users, argues that deathbots may harm the grief process and expose dependent users to surreptitious advertising, and proposes that they be classified as medical devices. Hollanek and Nowaczyk-Basińska (2024), working through design scenarios, recommend consent from both the "data donor" and the user, transparency about the system's nature, age restrictions, and rituals for retiring a deadbot.

Two features of this literature are worth noting. First, its main concepts are harm and benefit to mourners, the dignity of the dead, and consent. Where attribution appears, it appears as one risk among many, for instance the risk that a bot will be used to advertise products in the dead person's voice (Hollanek & Nowaczyk-Basińska, 2024). Second, the literature draws largely on Western moral philosophy, with Elder's Confucian study as a

notable exception. The religious traditions within which a large share of the world's population actually mourns, and which have developed detailed norms about the dead, are barely represented. The rest of this article aims to address both features.

4 The Attribution Problem

4.1 Generation as ventriloquism

Consider what happens when a mourner asks a generative griefbot of her late father whether she should accept a job abroad. Suppose the bot replies, in his familiar phrasing, that he is proud of her and that she should go. Three things are true of that reply. It was not said by the father. It was not merely a rearrangement of things he said; it is a new judgment about a situation he never knew. And it is presented in a form, his voice, that invites the daughter to treat it as his view.

The reply is a kind of ventriloquism. The ordinary ethics of ventriloquism is relaxed because the audience knows the dummy does not speak, and because the dummy is not anyone. Neither condition holds cleanly here. The bot is modelled on someone in particular, and the point of the product is that the user should, at least for a time, hear that person. The emotional value of the exchange depends on the attribution being felt.

We will say that a system makes a posthumous attribution when it produces content that is presented, by form, voice or framing, as the speech, belief, intention or commitment of a specific dead person, and that the person did not in fact produce. The attribution problem is the set of moral concerns that arise when such attributions are made systematically, at scale and for profit. It has at least four parts.

The first is accuracy. A language model estimates what text is likely given its training and prompt. It has no access to what the person actually believed about matters they never discussed. Its outputs can be fluent and wrong, and the user has no way to tell which.

The second is authority. Even when an attribution happens to match what the person would have said, the person did not authorise it. Words in someone's voice carry their authority. A father's blessing, a mother's advice, a friend's forgiveness, matter because of who gives them. A generated blessing borrows that authority without consent.

The third is circulation. Generated attributions do not stay inside a private session. Users screenshot them, quote them and share them. A sentence generated in a moment of grief can become, within a family, "what Dad said about the house". The system has thereby inserted testimony into the record of a life.

The fourth is irreversibility. A living person misquoted can object. The dead cannot. The usual social remedy for false attribution, correction by the speaker, is unavailable.

4.2 Why consent does not dissolve the problem

A natural response is that the problem disappears if the person consented before death. Hollanek and Nowaczyk-Basińska (2024) rightly put consent from the data donor at the centre of their recommendations, and Harbinja (2017) has argued for a legal concept of post-mortem privacy grounded in autonomy and testamentary freedom. Consent is necessary. It is not sufficient.

The reason is that consent to a generative system is consent to an open-ended set of future utterances, most of which the person could not have foreseen. One can consent to the publication of one's letters because one knows, or can know, what the letters say. One cannot in the same sense consent to everything a model will say in one's name over decades, in response to questions about events one never lived through. Consent of this kind resembles signing a blank page. We do not usually think that a signature on a blank page authorises whatever is later written above it, especially where what is written can change relationships, decisions and reputations.

Nor does consent from relatives resolve the problem. Relatives can consent to their own use of a system. They cannot authorise new speech on behalf of the dead any more than they could, while the person lived, authorise a forged letter in their name.

4.3 Can the dead be wronged?

The attribution problem presupposes that something is owed to the dead. A long-standing objection holds that this is confused: the dead no longer exist, so nothing can happen to them. Feinberg (1984) answered that some of a person's interests, such as the interest in reputation or in the success of projects, can be set back after death, because those interests were directed at states of the world that outlast the person. Pitcher (1984) argued that a harm done after death can be ascribed to the ante-mortem person, the person as they were in life, whose life is made worse by what later happens to their name and projects.

One does not need to settle this debate to see the attribution problem. Even on a sceptical view, posthumous attribution wrongs the living: the mourner who is misled about her father's views, the sibling who disagrees with what the bot "said", the wider community whose memory of a person is quietly revised. On the Feinberg and Pitcher views, it also wrongs the dead person directly, by setting back their interest in being known for what they actually said and believed. The more fluently a system speaks, the larger the gap between the person who lived and the person the system presents, and the greater the setback.

4.4 The fictionalist reply

The strongest reply comes from Krueger and Osler (2022), who argue that users of griefbots typically adopt a fictionalist stance. They engage with the bot as one engages with a novel or a game: emotionally, but without believing that the dead person is really speaking. If that is right, there is no false attribution, only an imaginative exercise that both parties understand.

The reply captures something real. Many users do know what they are talking to, and fiction can be a healthy medium for grief. But it faces three limits.

First, the fictional stance is unstable under exactly the conditions griefbots create. The product is designed to make the user forget, for a while, that the speaker is a model. Henrickson (2023) notes that for a griefbot to work as a meeting place of the living and the dead, the user must at least temporarily suspend awareness of the system's computational nature. The emotional power of the exchange comes from the moments when the frame slips. A design that profits from those moments cannot rely on the frame to absolve it.

Second, fiction does not travel well. A novel carries its fictional status with it; a screenshot of a chat does not. When generated words leave the session, they leave the frame.

Third, and most important, fiction has its own ethics of attribution. Historical fiction that puts invented words in the mouths of real people is a recognised genre, but it is also one with recognised limits. Readers and courts distinguish between an evident novel and a fabricated quotation, and between invention about public figures long dead and invention about private individuals whose families are still grieving. The fictionalist reply therefore does not remove the attribution problem. It relocates it within the ethics of fiction, where the question of what we may make real people say remains open.

5 An Islamic Ethics of Speaking for the Dead

5.1 Why this tradition

The literature on griefbots has largely drawn on Western moral theory, with Elder's (2020) neo-Confucian study as the clearest exception. Elder's work shows the value of asking how a long-lived tradition of mourning would understand the new technology. Islamic thought offers a particularly rich resource for the attribution problem, for three reasons. It has explicit norms about how the dead may be spoken of. It has a centuries-old scholarly discipline devoted to the verification of attributed speech. And it has, very recently, produced a juristic ruling that addresses AI-generated representations of the

dead directly. We draw on these sources not to settle questions of Islamic law, which is beyond the scope of an applied ethics paper and the competence of its author, but to show what the tradition can contribute to a shared moral problem.

5.2 The Dar al-Ifta ruling

In January 2025 Egypt's Dar al-Ifta, one of the most influential fatwa-issuing institutions in the Sunni world, published an English-language ruling in answer to a question from a programmer. A client had asked him to create an AI-generated video of the client's late father from photographs and audio, and to recreate a specific situation involving the two of them (Dar al-Ifta al-Misriyyah, 2025). The ruling begins from the juristic principle that a ruling on a thing depends on a proper conception of it, and then distinguishes three cases.

In the first case, animation is permissible if the images "reflect the deceased's usual speech and movements", with the intention of finding solace in moments of longing. In the second, it is prohibited where the images faithfully reflect the person's manner but the purpose is "to reject reality, dwell on past memories, renew grief, or inwardly object to Allah's law regarding death". In the third, it is "categorically prohibited" where the purpose is "making the deceased appear to say what they never did, attributing to them commitments they never made, or depicting them as testifying to something they did not witness". The ruling adds that such videos are often used for deceptive advertising and that gains obtained by deception are forbidden. Its summary states that the practice is permissible "for remembrance, finding comfort, prompting prayers of mercy and forgiveness", provided it causes no harm, present or future.

Two features of the ruling deserve emphasis. First, it locates the central wrong precisely where Section 4 located it: in attribution. The third case is not about data, privacy or even grief; it is about making the dead say, promise or witness what they did not. Second, it combines this with a concern for the mourner that is close to the clinical literature's distinction between healthy and complicated grief. Permission is tied to comfort and remembrance; prohibition is tied to the renewal of grief and the refusal to accept death.

The ruling was written about video, and its first case speaks of reproducing a person's "usual speech and movements". Applied to a generative chatbot, the ruling becomes demanding. A system that reproduces a person's characteristic expressions is on the permitted side. A system that composes new opinions, advice, promises or testimony in their name is, by the ruling's own terms, on the prohibited side, whatever the intentions of the user. Generative griefbots as currently built fall largely on the wrong side of this line, because generation of new content is their core function.

5.3 Attribution as a moral discipline

The ruling did not appear from nowhere. It draws on a tradition in which the attribution of speech has been a central moral and scholarly concern for more than a millennium. The clearest example is the science of hadith, the discipline that developed to evaluate reports of what the Prophet Muhammad said and did. Its methods, including the scrutiny of chains of transmission (isnad) and the assessment of the reliability of each transmitter, were developed precisely because false attribution to a revered dead speaker was recognised as both tempting and grave (Brown, 2009). The canonical collections preserve the Prophet's warning that whoever "tells a lie against me (intentionally)" should take his seat in the Fire (Sahih al-Bukhari, n.d.-a, hadith 110).

That warning concerns the Prophet specifically, and its theological weight cannot simply be transferred to ordinary people. But the discipline it inspired expresses a general moral insight that applies widely: the authority of a dead speaker is a trust, and speech attributed to them must be traceable to them. The hadith scholars did not forbid all speech about the Prophet. They distinguished between what could be shown to come from him, what was weakly attested, and what was fabricated, and they developed a vocabulary for marking those differences. That is, in essence, a provenance system for post-humous speech.

Other norms point in the same direction. The Qur'an condemns backbiting by comparing it to eating the flesh of one's dead brother (Qur'an 49:12, in Abdel Haleem, 2004), an image that ties the protection of a person's honour to their absence. A sound hadith forbids abusing the dead, "because they have reached the result of what they forwarded" (Sahih al-Bukhari, n.d.-b, hadith 1393). The dead have finished their work, and their record should not be tampered with by the living. Another sound hadith teaches that when a person dies, their deeds come to an end except for three: ongoing charity, beneficial knowledge, and a righteous child who prays for them (Sahih Muslim, n.d., hadith 1631). On this view, the dead person's moral account is closed except through channels that flow from what they themselves did. A system that generates new commitments in their name claims, in effect, to reopen that account on their behalf.

The Qur'anic principle that each soul bears only its own burden (Qur'an 6:164; see Abdel Haleem, 2004) adds a final element. The words a griefbot generates cannot change what the dead person did. But that is exactly why attributing them is a wrong against truth: it creates a false impression in the minds of the living about what the person was, said or wanted.

5.4 Grief, acceptance and continuing bonds

The ruling's second case, the prohibition of animation that renews grief or rejects the reality of death, might seem to conflict with the continuing-bonds model favoured by Elder (2020) and Krueger and Osler (2022). The conflict is less sharp than it appears. The con-

tinuing-bonds literature rejected an older clinical view that healthy grief requires detachment from the dead; it did not endorse a refusal to accept that the person has died (Klass et al., 1996). Stroebe and Schut's (1999) dual process model describes healthy grieving as oscillation between attention to the loss and attention to restoring a life, and treats persistent avoidance of either as a risk.

Islamic mourning practice expresses a similar balance. The Qur'an praises those who, when struck by misfortune, say "We belong to God and to Him we shall return" (Qur'an 2:156, in Abdel Haleem, 2004), a formula of acceptance that is recited at news of a death. At the same time, the tradition encourages continued remembrance of the dead, prayer for them, and acts of charity in their name, which are themselves continuing bonds. The ruling's line between remembrance and the renewal of grief is thus close to the clinical line between continuing bonds and complicated grief. Both point to design features that support remembrance while avoiding systems that keep a mourner in an artificial present with the dead.

The convergence matters. It suggests that the standard developed below is not a parochial religious constraint but a point of agreement between a major religious tradition and the secular literature, reached by different routes.

6 From Principle to Design: Attribution-Bounded Griefbots

If the attribution problem is real, what follows for design? We propose that griefbots be attribution-bounded: built so that what they present as the speech of the dead can be traced to the dead, and so that whatever cannot be traced is clearly marked as something else. Five requirements follow. Table 1 summarises them alongside the concern each addresses and its support in the sources discussed.

Table 1 Requirements for attribution-bounded griefbots

Requirement	What it requires	Concern addressed	Support in the sources
Quotation before generation	Retrieval of the person's own words as the default; generation only on explicit user choice	Accuracy; authority	Dar al-Ifta (2025), first case: reproducing "usual speech"
Visible provenance	Every reply marked as quoted, adapted or generated, in a form that survives copying	Circulation; authority	Hadith grading of reports (Brown, 2009); transparency duties (Regulation (EU) 2024/1689, art. 50)
No new commitments	No generated promises, blessings, testimony, religious or legal opinions, or endorsements in the	Authority; irreversibility	Dar al-Ifta (2025), third case; covert-advertising risk (Hollanek & Nowaczyk-

	person's voice		Basińska, 2024)
Mourning-sensitive use	Staged access, no simulated unprompted contact, a dignified way to retire the bot, no engagement maximisation	Harm to mourners	Dar al-Ifta (2025), second case; Elder (2020); Stroebe & Schut (1999)
Stewardship of the voice	Mode-specific ante-mortem consent; a named steward with powers to correct and withdraw	Irreversibility; consent gap	Öhman & Floridi (2018); Harbinja (2017)

6.1 Quotation before generation

The default mode of a griefbot should be retrieval: answering with the person's own words, drawn from their messages, recordings and writings, with minimal adaptation. Modern systems can do this well, by searching a person's archive for passages relevant to the user's input and presenting them. Generation should be a secondary mode, entered only when the user chooses it and the system has made clear that it is now composing rather than quoting. This preserves the main good that griefbots offer, hearing the person's voice again, while greatly reducing posthumous attribution. It corresponds closely to the first, permitted case in the Dar al-Ifta ruling, which allows reproduction of the deceased's "usual speech".

6.2 Visible provenance

Every reply should carry a visible marker of its source: quoted verbatim, lightly adapted from identifiable material, or generated. The marker should survive copying and screenshots, for instance as part of the reply text itself rather than as interface decoration. This is the modern analogue of the hadith scholars' grading of reports. It does not forbid generated content, but it prevents generated content from passing as testimony. It also addresses the circulation concern in Section 4.1, because a marked reply carries its status with it.

6.3 No new commitments

Some kinds of content should not be generated in a dead person's voice at all, whatever the mode. These include promises, permissions and blessings concerning decisions made after death; testimony about events; religious or legal opinions; statements about inheritance, family disputes or other contested matters; and endorsements of products, causes or candidates. These categories correspond to the third, categorically prohibited case in the Dar al-Ifta ruling: commitments the person never made and testimony they never gave. They also correspond to the scenario of covert advertising that Hollanek and

Nowaczyk-Basińska (2024) identify as a key risk. A system can still respond to such prompts, but it should do so in its own voice, for instance by saying that it cannot speak for the person on this matter and by offering relevant things the person actually said.

6.4 Mourning-sensitive use

Design should support remembrance without encouraging a refusal to accept death. Elder (2020), drawing on Xunzi, suggests that access might be gradually shortened, as traditional mourning rites move in stages. Hollanek and Nowaczyk-Basińska (2024) recommend rituals for retiring a deadbot. Concretely, systems could default to limited session frequency in the first months after a death, avoid features that simulate the person initiating contact (for example, unprompted messages on anniversaries), and offer an explicit, dignified way to end the relationship with the bot. Systems should also avoid engagement-maximising design, since the commercial incentive to keep a grieving user returning is in direct tension with the user's welfare (Öhman & Floridi, 2017).

6.5 Stewardship of the voice

Finally, someone must be responsible for the dead person's voice. Ante-mortem consent should be specific about modes: a person might consent to a retrieval-based archive but not to generation. Where a system is built after death without such consent, the default should be retrieval only. A named steward, typically a family member chosen by the deceased or by the family, should be able to review, correct and withdraw the system, and should be able to delete generated content that has been wrongly attributed. This does not solve the problem that the dead cannot correct their own record, but it provides the nearest available substitute.

7 Implications for Regulation

Existing regulation addresses parts of this problem but not its core. The European Union's Artificial Intelligence Act requires that people be informed when they interact with an AI system, and that "deep fake" content be disclosed as artificially generated (Regulation (EU) 2024/1689, art. 50). These duties help with transparency, but a disclosure at the start of a session does not distinguish between quoted and generated content within it, and a griefbot user already knows she is speaking to a system. The EU's General Data Protection Regulation states that it does not apply to the personal data of deceased persons and leaves the matter to member states (Regulation (EU) 2016/679, recital 27). In many jurisdictions, therefore, the dead person's data and voice fall into a gap between privacy law, which protects the living, and intellectual property or personality rights, which vary widely (Harbinja, 2017).

The attribution-bounded standard suggests a more targeted approach. Regulators could require that griefbot providers (a) disclose whether a system is retrieval-based or generative; (b) mark generated content in a way that survives export; (c) prohibit generated commitments, testimony and endorsements in the voice of the dead; (d) offer a stewardship mechanism with powers of correction and withdrawal; and (e) refrain from engagement-maximising design for bereaved users. None of these requirements bans the technology. Each can be audited.

The Gulf states, and the United Arab Emirates in particular, have reason to take a lead. The UAE is investing heavily in Arabic-language generative models, and its Charter for the Development and Use of Artificial Intelligence commits to principles that include data privacy, transparency, human oversight, and "placing human values at the heart of technological innovation" (UAE Government, 2024). The Charter does not mention the dead, and no published UAE ruling specifically on AI recreations of deceased people was identified in the course of this research. But in a society where Islamic norms about the dead are widely shared, a regulatory standard that draws on those norms, and that also converges with secular ethical work, is likely to enjoy both legitimacy and compliance. The Dar al-Ifta ruling shows that religious authorities are already addressing these questions. Regulators, providers and religious scholars could usefully work together on a shared standard, of which the five requirements above offer a starting point.

8 Objections and Replies

8.1 Paternalism

It may be objected that restricting generation is paternalistic. If a mourner wants to hear new words in her father's voice, and knows they are generated, who is harmed? The reply has two parts. First, the attribution-bounded standard does not forbid generation; it requires that generation be chosen, marked and kept away from certain categories. A mourner who wants generated content can have it, clearly labelled. Second, the parties affected are not only the user. They include the dead person, on the views of Feinberg (1984) and Pitcher (1984); other relatives, whose memory of the person is also at stake; and anyone who later encounters the generated words without their frame. Rules that protect third parties are not paternalistic in the objectionable sense.

8.2 All memory is reconstruction

A second objection notes that human memory is itself reconstructive. When we remember what a dead parent "would have said", we are already generating. Why hold machines to a higher standard? The answer is that the difference lies in authority and scale. A family member's remembered paraphrase is understood by everyone as that family

member's interpretation, and it competes with other interpretations. A system that produces fluent text in the person's own voice, with the person's own turns of phrase, presents its output with an authority that no individual memory has. And it does so on demand, indefinitely, for a fee. The hadith tradition offers a precedent here too: it distinguished between transmitted reports and the interpretations of scholars, and it treated the confusion of the two as a serious fault (Brown, 2009). Griefbots should make the same distinction.

8.3 Parochialism

A third objection holds that grounding a standard in Islamic sources makes it parochial. But the argument of Sections 4 and 5 does not depend on accepting Islamic premises. The attribution problem was developed from the general philosophy of posthumous harm and the ethics of fiction. The Islamic sources were introduced to show that a major tradition reaches the same conclusion by a different route, and that it has practical resources, such as a developed vocabulary of provenance, that the secular literature lacks. Convergence of this kind is a reason for confidence in the standard, not a reason to doubt it. Moreover, a global technology that will be used by mourners of every faith should be governed by standards that can be justified within the moral traditions those mourners actually hold.

9 Conclusion

Generative griefbots promise something that no earlier technology of remembrance could offer: new words from those we have lost. This article has argued that the promise contains its own difficulty. To produce new words in a dead person's voice is to attribute to them speech they never made, with an authority they never granted, in a record they can no longer correct. Consent before death narrows this problem but cannot remove it, and the fictional stance that many users adopt is too unstable, and too easily lost in circulation, to dissolve it.

The Islamic tradition, through the 2025 Dar al-Ifta ruling and the much older discipline of verified attribution, draws the line in a clear place: one may reproduce the dead person's characteristic speech for remembrance and comfort, and one may not make them say, promise or testify to what they did not. This line converges with secular work on posthumous harm and on healthy grieving. It can be turned into design requirements that are concrete and auditable: quotation before generation, visible provenance, no new commitments, mourning-sensitive use, and stewardship of the voice.

The general lesson reaches beyond griefbots. As generative systems learn to speak in the voices of particular people, living and dead, societies will need norms about who may speak for whom. The ethics of attribution, long practised by traditions that cared about

the words of the dead, offers a good place to begin.

References

Abdel Haleem, M. A. S. (Trans.). (2004). *The Qur'an: A new translation*. Oxford University Press.

Abramson, D. I., & Johnson, J., Jr. (2020). Creating a conversational chat bot of a specific person (U.S. Patent No. 10,853,717 B2). U.S. Patent and Trademark Office. <https://patents.google.com/patent/US10853717B2/en>

Brown, J. A. C. (2009). *Hadith: Muhammad's legacy in the medieval and modern world*. Oneworld.

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>

Dar al-Ifta al-Misriyyah. (2025, January 1). Creating an AI-generated video of a deceased person (Fatwa No. 24168). <https://www.dar-alifta.org/en/fatwa/details/24168/creating-an-ai-generated-video-of-a-deceased-person>

Elder, A. (2020). Conversation from beyond the grave? A neo-Confucian ethics of chatbots of the dead. *Journal of Applied Philosophy*, 37(1), 73–88. <https://doi.org/10.1111/japp.12369>

Fabry, R. E., & Alfano, M. (2024). The affective scaffolding of grief in the digital age: The case of deathbots. *Topoi*, 43(3), 757–769. <https://doi.org/10.1007/s11245-023-09995-2>

Fagone, J. (2021, July 23). The Jessica simulation: Love and loss in the age of A.I. *San Francisco Chronicle*. <https://www.sfchronicle.com/projects/2021/jessica-simulation-artificial-intelligence/>

Feinberg, J. (1984). *Harm to others: The moral limits of the criminal law* (Vol. 1). Oxford University Press.

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>

Henrickson, L. (2023). Chatting with the dead: The hermeneutics of thanabots. *Media, Culture & Society*, 45(5), 949–966. <https://doi.org/10.1177/01634437221147626>

Hollanek, T., & Nowaczyk-Basińska, K. (2024). Griefbots, deathbots, postmortem avatars: On responsible applications of generative AI in the digital afterlife industry. *Philosophy & Technology*, 37(2), Article 63. <https://doi.org/10.1007/s13347-024-00744-w>

Klass, D., Silverman, P. R., & Nickman, S. L. (Eds.). (1996). *Continuing bonds: New understandings of grief*. Taylor & Francis.

Krueger, J., & Osler, L. (2022). Communing with the dead online: Chatbots, grief, and continuing bonds. *Journal of Consciousness Studies*, 29(9–10), 222–252. <https://doi.org/10.53765/20512201.29.9.222>

Lindemann, N. F. (2022). The ethics of 'deathbots'. *Science and Engineering Ethics*, 28(6), Article 60. <https://doi.org/10.1007/s11948-022-00417-x>

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*, 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 Behaviour*, 2(5), 318–320. <https://doi.org/10.1038/s41562-018-0335-2>

Pitcher, G. (1984). The misfortunes of the dead. *American Philosophical Quarterly*, 21(2), 183–188.

Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation). (2016). *Official Journal of the European Union*, L 119, 1–88. <https://eur-lex.europa.eu/eli/reg/2016/679/oj>

Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). (2024). *Official Journal of the European Union*, L series. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>

Sahih al-Bukhari. (n.d.-a). Hadith 110, Book of Knowledge. Sunnah.com. <https://sunnah.com/bukhari:110>

Sahih al-Bukhari. (n.d.-b). Hadith 1393, Book of Funerals. Sunnah.com. <https://sunnah.com/bukhari:1393>

Sahih Muslim. (n.d.). Hadith 1631, Book of Wills. Sunnah.com. <https://sunnah.com/muslim:1631>

Stroebe, M., & Schut, H. (1999). The dual process model of coping with bereavement: Rationale and description. *Death Studies*, 23(3), 197–224. <https://doi.org/10.1080/074811899201046>

UAE Government. (2024). The UAE Charter for the Development and Use of Artificial Intelligence. The Official Portal of the UAE Government. <https://u.ae/en/about-the-uae/strategies-initiatives-and-awards/policies/Ai/The-UAE-Charter-for-the-Development-and-Use-of-Artificial-Intelligence>