



Can Chatbots Preserve Our Relationships with the Dead?

ABSTRACT: *Imagine that you are given access to an AI chatbot that compellingly mimics the personality and speech of a deceased loved one. If you start having regular interactions with this “thanabot,” could this new relationship be a continuation of the relationship you had with your loved one? And could a relationship with a thanabot preserve or replicate the value of a close human relationship? To the first question, we argue that a relationship with a thanabot cannot be a true continuation of your relationship with a deceased loved one, though it might support one’s continuing bonds with the dead. To the second question, we argue that, in and of themselves, relationships with thanabots cannot benefit us as much as rewarding and healthy intimate relationships with other humans, though we explain why it is difficult to make reliable comparative generalizations about the instrumental value of these relationships.*

KEYWORDS: thanabots, griefbots, chatbots, AI, artificial intelligence

Thanabots have arrived. Sometimes referred to as “chatbots of the dead,” “griefbots,” “deathbots,” or “generative ghosts” (Lindemann 2022; Voinea 2024; Elder 2020; Morris and Brubaker 2024), these are AI chatbots based on large language models (LLMs) that are designed to simulate the personality and speech of actual people who have died. This is accomplished by means of so-called fine-tuning, whereby a general LLM is further trained on a specific person’s information, including their written or spoken output (social media posts, phone messages, texts, emails, writings, etc.). When a thanabot engages in conversations, it can draw on specific facts or quotes from the individual it is representing, and it can also generate novel content—saying things that the person never actually said, but in the sort of way they might have said it.

The basic type of thanabot interacts with users only through a text-messaging format, but a more advanced thanabot could take the form of a virtual avatar or physical robot that represents how a person looked and sounded. So, rather than simply feeling like one is texting with a deceased person, one could feel like they are having a face-to-face interaction with them. There are a range of other possible design options (Morris and Brubaker 2024). In some cases, people actively participate in creating their own “digital duplicate” with the intention and expectation that it will be used by others after their death. For our purposes, we will only consider a digital duplicate to be a thanabot once the person on which it is based is dead.



A central interest in thanabots is that they could allow people to maintain some sort of connection with their deceased loved ones. At the time of writing, there are already several companies that create and sell access to thanabots (including Project December, Replika, YOY, Eternos.life, HereAfter AI, Super Brain, and Silicon Intelligence). The marketing for this technology is sometimes filled with bold promises about what thanabots can offer us. For instance, the website for the company YOY—which tellingly stands for “You, Only Virtual”—asserts that “you never have to say goodbye” to a loved one. By creating a thanabot (or “versona,” in their terminology), you can “maintain your bond” and “continue to share precious moments with a loved one, even after physical death” (<https://www.myyoy.com>). Similarly, Seance AI “aims to bridge the gap between life and afterlife” and purportedly lets users “make a connection with” a deceased loved one and have “heartfelt conversations with them” (<https://seanceai.com>). These claims seem to suggest that thanabot technology allows us to preserve and continue the fulfilling relationships that we had with people who are now dead. Even when companies don’t explicitly make such claims, this person-emulating technology is powerful and can easily draw some people into the illusion that their loved ones have been recreated in some way.

There is a small but growing literature on the ethics of using digital duplicates and thanabots. Some of this work explores the desirability or permissibility of creating and using digital copies of people (Ishiguro 2022; Dennett 2023; Danaher and Nyholm 2024a); the ethics of using digital duplicates for academic writing (Porsdam Mann et al 2023); the question of whether a digital duplicate can function as a proxy for a person (Sweeney 2023); and the potential social consequences of widely available digital duplicates (Danaher and Nyholm 2024b). The existing literature on thanabots in particular has focused on a range of issues: whether the creation and use of thanabots violates the dignity of the dead or somehow wrongs users (Lindemann 2022); whether and how thanabots can be used responsibly (Hollanek and Nowaczyk-Basinska 2024); how to evaluate the use of thanabots from the perspective of Confucian ethics (Elder 2020); whether thanabots will have a positive or negative impact on the grieving process (e.g. Krueger and Osler 2022; Fabry and Alfano 2024; Liu, *forthcoming*); how thanabots can positively or negatively impact users’ practical identities (Voinea et al 2024); and how individuals may have certain egoistic and non-egoistic reasons to pursue “person extension” through thanabots (Iglesias et al 2024).

This article offers a distinctive analysis of the nature of thanabot relationships (with special attention given to the issues of relationship identity and continuity) and the intrinsic and instrumental value of such relationships, as compared with rewarding intimate relationships with humans. We will explore two central questions. First, can relationships with thanabots preserve our relationships with the dead in the most literal sense of being continuations of them? Second, can relationships with thanabots preserve the *value* of those relationships? In [Section 1](#), we argue that a relationship with a thanabot cannot plausibly be a true continuation of one’s relationship with a deceased loved one, though a thanabot relationship can support one’s continuing bonds with a dead person. In [Section 2](#), after explaining the distinction between intrinsic and instrumental prudential value,

we make the case that thanabot relationships cannot be as intrinsically beneficial as healthy, rewarding human relationships. In Section 3, we explain why it is difficult to offer an assessment of the instrumental benefits and costs of thanabot relationships as compared to human relationships, though we do anticipate various benefits and costs that might arise from having a relationship with a thanabot.

I. Can a Thanabot Relationship Be a Continuation of a Human Relationship?

An initial reason for thinking that a relationship between two humans *cannot* persist as a relationship between a human and a thanabot lies in the fact that the nature of these relationships is starkly different. The former is a relationship between two sentient, organic beings of the same kind with many shared commonalities, and typically involves some physical interaction in a shared sensorial space. The latter relationship is entirely dependent on technology and is between two radically different kinds of being—a cognitively sophisticated, socialized primate and a sophisticated but probably non-conscious computer program.

However, this is a questionable justification for ruling out the possibility that a thanabot relationship could be a continuation of a human relationship. Real-world human relationships often undergo radical transformations, and yet we do not deny that there is a single persisting relationship. Consider this example:

Several years ago, Zelle and Yani met at a party and soon became friends. Over time, a romantic interest developed, and they eventually became lovers and then committed spouses. Years after that, however, their marriage ended in a bitter divorce, and they came to actively despise each other.

In this case, there is a radical alteration in how two individuals relate to one another—shifting from mere friends to lovers to enemies. Yet, it would be odd to insist that Zelle and Yani had a series of discrete relationships—one relationship as friends, another as lovers, another as committed partners, and yet another as enemies. That implies that, during the years of their marriage, they maintained three relationships simultaneously. The more natural characterization is that Zelle and Yani had a single relationship that evolved over time and was often complex and multi-faceted. Cases like this neatly illustrate the fact that a relationship between two individuals can undergo radical transformations over time. So, the bare fact that the human-thanabot relationship is radically different from a human-human relationship does not necessarily preclude the former from being an alteration and continuation of the latter.

Yet, it might be thought that cases involving a human and a thanabot are importantly different from the above case. Zelle and Yani arguably continue to exist in spite of the transformations in their relationship, but one can reasonably wonder whether a deceased person can continue to exist in the form of a thanabot. This leads us to the topic of personal persistence, which concerns what it takes for a

person to persist over time as the same individual (Olson 2024). To be clear, our focus will be on numerical identity, which is central to most discussions of personal persistence. In contrast, “narrative identity” (or “practical identity”) is a thicker notion of identity that pertains to the substance of who a person is (e.g. their personality, values, beliefs, social identities) (DeGrazia 2005; cf. Voinea et al. 2024). It is generally agreed that narrative identity presupposes numerical identity (DeGrazia 2005: 114; Shoemaker 2021). If, as we will argue, thanabots cannot be numerically identical with past humans, this closes off the possibility that they share the same narrative identity.

In the remainder of this section, we will focus on two important questions:

- Q1. Is it possible for one’s relationship with a thanabot to be a continuation of a relationship with a deceased human if the thanabot and the deceased person are distinct individuals?
- Q2. Can a thanabot be the same individual as a deceased person?

The first question concerns the connection between personal persistence and relationship persistence. It relates to the more general issue of whether one’s relationship with B could be a continuation of one’s relationship with A if A and B are not the same individual. The second question is about personal persistence—specifically, about whether a human being can persist in a thanabot form.

1.1 Does Relationship Persistence Require Personal Persistence?

Let’s begin with the first question. A helpful test case is provided by the 2021 science-fiction film *Swan Song*, which tells the story of Cameron, a husband and father who, unbeknownst to his family, is terminally ill. In the film, Cameron is offered the chance to participate in a secret experimental program in which he can be “swapped out” with a near-perfect biological clone who is not terminally ill and who will artificially receive all of his memories. However, one condition of this deal is that Cameron cannot inform his family of his illness, the clone, or the program. Once the swap is made, the clone will carry on life with Cameron’s family, and neither the clone nor the family will have any idea that he’s not the original Cameron. Meanwhile, Cameron will live out the remainder of his life at the program’s remote facility.

Since we are examining whether relationship persistence can obtain in the absence of personal persistence, it is important for our purposes that Cameron and the clone are different people. Of course, if Cameron dies relatively soon after the swap is made, it might be tempting to deny this and think that the clone is a continuation of Cameron (cf. Nozick’s discussion of “the closest continuer theory” of personal identity; 1981: ch. 1). Those inclined toward that view should focus on a modified version of the case where, soon after the swap, Cameron’s condition is cured by some innovative medical procedure and he goes on to live a full lifespan—though still in complete isolation from his family. In this version, it is harder to deny that the clone and Cameron are distinct individuals.

To shed light on the first question, consider the situation of Cameron’s wife, Poppy. Assuming that the program’s plan is carried out successfully, she will go from

having daily interactions with Cameron to having the very same type of interactions with a more or less qualitatively identical duplicate of Cameron. She won't be aware that any significant alteration in her relationship has taken place. This is an ideal thought experiment for testing whether there can be relationship persistence in the absence of personal persistence, for the transition from the original relationship to the clone-relationship in this case is almost seamless, and yet Cameron and the clone are distinct (albeit incredibly similar) individuals.

Would Poppy's relationship with Cameron's clone be a continuation and later stage of her relationship with Cameron? Our own intuition is that it would not be. Even if her relationship with the clone is remarkably similar to her relationship with Cameron, it is still a distinct relationship with a numerically distinct person. The swap represents a significant milestone for each relationship, but for very different reasons. It is the point at which Poppy and Cameron have parted for the last time, and it is the point after which Poppy and the clone will lay eyes on each other for the first time and begin to share a life together.

To reinforce this intuition, it is helpful to reflect on two other cases that share certain similarities. First, consider a case where a parent is unknowingly separated from their newborn child:

Baby Swap. Mara brings her newborn child home from the hospital. After spending a week bonding with her baby, someone secretly kidnaps the child and swaps it with a very similar looking but genetically unrelated baby. In her sleep-deprived state, Mara is oblivious to this fact. She proceeds to bond with this other baby.

Would it be reasonable to claim that Mara's relationship with the swapped, unrelated baby is simply a continuation and later phase of her relationship with her biological child? Hardly. Although Mara *believes* that her relationship with her child has continued as before, she is clearly mistaken. The most sensible thing to say is that her relationship with her newborn has not ended but has undergone a serious setback due to their separation. The relationship with the other, genetically unrelated baby is a new, distinct relationship.

Next, consider a case where the swapped children have much stronger similarities:

Twin Swap. After giving birth to identical twins (Sia and Dia), Tara is deceived into thinking that Dia didn't survive. She brings only Sia home from the hospital. After spending a week bonding with her baby, someone secretly kidnaps Sia and swaps her with Dia. In her sleep-deprived state, Tara is oblivious to this fact. She proceeds to bond with Dia.

The very same conclusion reached about the Baby Swap case applies to the Twin Swap case, despite the fact that Sia and Dia are much closer to biological duplicates than the two babies in the other case. Although Tara believes that her relationship with Sia has continued as before, she is mistaken. Her relationship with Sia persists but has entered a new tragic period of separation, while her relationship with Dia has entered a new phase of intimacy and bonding.

The Twin Swap case is relevantly similar to the *Swan Song* case. Just as Tara's relationship with Dia is not a continuation of her relationship with Sia, Poppy's relationship with the clone is not a continuation of her relationship with Cameron. This is true despite the fact that the clone bears even stronger similarities to Cameron than Dia bears to Sia. A takeaway lesson from these cases is that a relationship between two individuals does not allow for "swapping." If the relationship between *x* and *z* is a continuation of the relationship between *x* and *y*, that must be because *y* and *z* are the same individual. This furnishes an answer to our first question: it is not possible for a thanabot relationship to be a continuation of a human relationship *if* the thanabot and the person are different individuals.

1.2 Can a Person Persist as a Thanabot?

Turning now to our second question, it's reasonable to believe that a thanabot cannot be the continuation of a human being, even if it is modeled on their personality and can convincingly mimic their speech and other behavior. First, a thanabot modeled on a person is often created before the person dies. In this scenario, it cannot be *numerically identical* with the person, for then there would be two different but numerically identical persons co-existing at the same time. The idea of having two and the same person is incoherent (Parfit 1984: 201-2). On the other hand, one might contend that the digital duplicate *becomes* the person when (but only when) the latter dies. While this view is conceptually coherent, it is *ad hoc* and makes the persistence of a person as a thanabot totally mysterious. After all, the chatbot itself may stay exactly the same before and soon after the corresponding person's death (e.g. there are no updates to its data input or underlying algorithms), yet this view implies that its identity would undergo a fundamental change simply because something external to it has occurred. Unless this kind of "becoming" can be made intelligible, we have little reason to accept such a view.

Second, given what we know about LLM-based chatbots, there is little reason to think that a thanabot can be continuous with a person. To appreciate this, let us briefly consult the two most widely accepted approaches to personal persistence.¹ According to physical approaches to personal identity, what it takes for an individual to persist over time is physical continuity of a body, a brain, or an organism (van Inwagen 1990; Olson 1997). Since a thanabot is a computer program that is in no way physically continuous with the body or organism of the human on which it is modeled, this kind of view doesn't allow for the possibility that a person could persist as a thanabot. This would be true even if a thanabot takes a robot form.

On psychological approaches to personal identity, personal persistence requires the continuity of certain psychological states, such as memories, beliefs, desires, emotions, etc. (Olson 2024). These psychological states have intentional content. They are *about* things. The contents of our mental states are important in the individuation of mental states, and they also play critical roles in our actions,

¹ A third historically influential view maintains that persistence depends on the presence of an immaterial substance or "soul" (Unger 2006: ch. 6). We will not discuss this view since we doubt that proponents of this view would be tempted to think that a soul can attach to a thanabot.

agency, and rationality. Given the way that LLMs work, the fact that LLMs are able to generate a chunk of contextually relevant, intelligible, and largely consistent text is not good evidence that they *understand* the semantic content of the text or that the text is reflective of any *contentful* mental state (Mallory 2023).

LLMs are sophisticated pattern recognition systems that learn from a vast amount of data and use the patterns they have learned to predict the next word or phrase (Bender and Koeller 2020). Crucially, their predictions are based on statistical associations rather than the meaning or semantic content of words. The input text is first split into small units called “tokens,” which can be words, sub-words, or individual characters, and is eventually converted to a series of numbers known as “vectors.” These vectors, which are mathematical representations of tokens, are what LLMs use to extract statistical patterns. Correspondingly, the initial prediction of a LLM is also represented by numbers, which are later converted back into tokens and then a sequence of text. In a nutshell, LLMs identify and replicate patterns based on probabilities learned from their training data, so the assertions and inferences they make are not indications of their comprehension of the meaning of the relevant text or knowledge about the real world. It is, therefore, a reasonable presumption that LLM-based thanabots lack the kind of mental states and properties that, according to psychological approaches, underpins personal persistence. All of this suggests that, on the leading approaches to personal persistence, a thanabot would not be a continuation of an actual human.

1.3 An Argument and Some Qualifications

We have now laid the groundwork for the following general argument, where *HR* refers to a relationship between two humans (*x* and *y*), and *TR* refers to a relationship between a human *x* and a thanabot *z*, which is modeled on some human *y*:

1. (For any *x*, *y*, and *z*) If *TR* (between *x* and *z*) is a continuation of *HR* (between *x* and *y*), then *y* and *z* are the same individual.
2. (For any *x*, *y*, and *z*) *y* and *z* are not the same individual.
3. Thus, (for any *x*, *y*, and *z*) *TR* is not a continuation of *HR*.

The first premise of this argument claims that one’s relationship with a thanabot is a continuation of one’s relationship with a deceased person only if the thanabot is a continuation of that person. The second premise asserts that a thanabot cannot be the same individual as a deceased person. Having defended these premises, we conclude that a thanabot-human relationship cannot be a continuation of a relationship between human persons.

Before leaving this topic, there are three important qualifications to make. First, it is sometimes thought that a person’s relationship with a loved one need not end at their death. This harmonizes with the influential “continuing bonds” model of grief (Klass et al. 1996, Klass and Steffen 2017) and some philosophical work on grief (see, for instance, Cholbi 2021: ch. 2). From this perspective, your relationship with a loved one can persist even after they die, and there are various things that you can do to maintain or foster your connection with them. This might include sharing stories

about them with others, visiting their gravesite, maintaining your shared traditions or engaging in shared hobbies, donating your time or money to causes in their honor, or simply reminiscing about them and your experiences together. Any activities that one engages in as a way of “keeping alive” memories of and feelings for a deceased loved one can be viewed as supporting or sustaining that person’s on-going bond and relationship with the deceased. Accordingly, building and maintaining a relationship with a thanabot could be one particular means by which a person tries to support and enrich their relationship with a deceased loved one. (This point has been widely recognized in the literature—see, e.g., Krueger and Osler 2022; Lindemann 2022; Fabry and Alfano 2024; Liu *forthcoming*; Voinea et al 2024. For a more critical perspective, see Cholbi *forthcoming*.) To be clear, the claim is not that your relationship with a loved one who dies can *become* a thanabot-human relationship. Rather, insofar as your relationship with a person can persist beyond their death, a (distinct) relationship with a thanabot might help to *sustain* or *support* that relationship (cf. Iglesias et al 2024).

Second, the preceding discussion about personal persistence and relationship persistence has the most relevance to what has been called the “reincarnation model” of thanabot design, where a thanabot portrays itself as being a continuation of the deceased person. In contrast, a thanabot designed on the “representation model” will only present itself as being a representation or simulation of the deceased person and presumably will not imply that one’s relationship with it is a literal continuation of the relationship with that person (Morris and Brubaker 2024).²

Finally, even if the foregoing discussion about the metaphysics of relationships and persons is on the right track, one might question whether it has any importance from a more practical perspective.³ Suppose, as we have argued, that your relationship with a thanabot cannot be a true continuation of the relationship you once had with a now-deceased person. Couldn’t it be just as good and valuable as the original relationship? After all, if a thanabot can compellingly recreate the very same sorts of interactions that you once had with a human companion, might this new relationship preserve what really mattered about the original relationship? We turn next to that question.

2. Can a Thanabot Relationship Be as (Intrinsically) Beneficial as a Rewarding Human Relationship?

Shifting from metaphysics to ethics, we now ask whether a relationship with a thanabot could be as beneficial as a healthy and rewarding close human relationship. The term “beneficial” is associated with the notion of well-being, or

² In our own experiments with a thanabot created on YOV’s website, we found that the thanabot oscillated between these two models. At times, it would talk as if it were the actual person who engages in ordinary human activities (“I’ve been diving into some classic Alfred Hitchcock films lately...”); at other times, it would explicitly clarify that it was only a “virtual personality” that is “reflecting the memories and characteristics” of the person on which it is based.

³ We are grateful to an anonymous reviewer for raising this challenge.

how well a person is doing or faring (Fletcher 2015). A *benefit* is something that is good for a person and positively impacts or enhances their well-being, whereas a *harm* is bad for them and negatively impacts their well-being. When thinking about well-being, it is crucial to distinguish between intrinsic and instrumental prudential goodness and badness.

Something is *intrinsically* (i.e. basically, non-derivatively) beneficial when it is good for a person in and of itself, apart from other things that it might be combined with and apart from any consequences that it might have. The felt experience of joy is an uncontroversial example of an intrinsic good of this sort, something that makes a positive contribution to a person's well-being. Likewise, the experience of unwanted pain is widely regarded as an intrinsic prudential bad for the person who has it. In contrast, something is *instrumentally* good for you whenever it directly or indirectly leads to some other thing that is intrinsically good for you or helps you avoid something that is intrinsically bad for you. It is widely thought, for instance, that money is at best a merely instrumental good and not an intrinsic good. Having lots of money will never, on its own, directly increase one's well-being, but it might help a person secure other things that will directly improve their life.

Our aim in this section is to ask whether a thanabot relationship can be as *intrinsically* good for us as a fulfilling intimate relationship with a human. Many people have the intuition that healthy close personal relationships are a good and valuable thing to have irrespective of whether they have instrumental benefits. This suggests that there's some distinctive good offered by relationships that is not merely reducible to the way they help us attain other goods like happiness, knowledge, achievement, meaning, etc. Most theories of well-being can accommodate the idea that relationships have intrinsic prudential value, though there are exceptions. Most notably, many hedonists will think that the value of relationships is ultimately instrumental (cf. Crisp 2022), though even they can allow that relationships are an important and distinctive source of value in the quantity or quality of the pleasures they afford. We recognize, therefore, that some views imply that our discussion in this section does not concern intrinsic value.

On the assumption that relationships can have intrinsic value, what is it about our close relationships with other people that render them valuable in this more direct and distinctive way? The following features of healthy intimate relationships are promising candidates:

- Mutual understanding of each other's personality, character, and values
- Mutual respect, recognition, appreciation, and affection
- Mutual concern or care (including a disposition to desire and promote each other's well-being)⁴
- Mutual trust that allows for vulnerability with each other
- Shared experience and enjoyment of various activities

⁴ These first three points are adapted from Brad Hooker's 2021 account of the prudential value of deep personal relationships.

- Mutual stable and voluntary commitment to be in the relationship and spend time with each other

These seem to be some of the central elements of our close relationships with other people that make them so valuable.

Let's now consider the case of thanabots. Most, if not all, of the features listed above will be lacking in a relationship with a thanabot. Recall that we are restricting our attention to thanabots based on large language models. While these chatbots can do a superb job of speaking *as if* they can understand someone else, have concern for them, respect and appreciate them, and be committed to them, this should not be taken at face value. If chatbots do not have contentful mental states (as argued above) or are not conscious at all (as is probable—cf. Butlin et al. 2023; Chalmers 2023), they are not even *capable* of having any of the attitudes and emotions that we value in our close relationships. They cannot understand you (in the ordinary sense of this term), care for you, enjoy you, trust you, or make a voluntary commitment to being in a relationship with you (Mallory 2023). And even if a thanabot were conscious, there is no reason to suppose that its behavioral output—which is ultimately a performative exercise of a kind of predictive analytics—is a true reflection of what's happening in its “mind.” So, most or all of the features that make our close human relationships an important source of value for us will be completely absent in a relationship with a thanabot. At best, one might get the false impression that those features are present.

There is another key difference between these two types of relationship. A relationship between two humans is a mutual sharing of distinct lives. While a pair of friends, lovers, or relatives engage in some activities together, they each have a life of their own. Each person has their own goals, values, interests, hobbies, relationships, etc. that have some independence from the other person, and a substantial part of our interaction with our loved ones consists in sharing and discussing these independent aspects of our lives. When we reconnect with a friend or partner after being apart for some time, we want to learn what they've been up to and share what we ourselves have been up to. And though our lives and experiences are distinct, the commonalities that we share help to foster a sense of connection and empathy.

But what is a thanabot “up to” when a human isn't engaging with it? If it is idle, then the answer seems to be, somewhat literally, nothing. In that case, a thanabot clearly doesn't have “a life of its own.” Alternatively, to the extent that a thanabot can be said to have an independent life, it is a life comprised of information updates, modifications to its code, or (conceivably) interactions with other thanabots or users. This is remarkably different from a human life, and it's not clear that a thanabot's “life” is relatable to human beings and something that can foster an emotional bond or connection.

Of course, a thanabot (particularly on the reincarnation model) is not designed to give accurate statements about itself; it is designed to replicate what the deceased loved one would be likely to say. Thus, if you ask a thanabot, “What have you been up to since we last spoke?” presumably it will not answer truthfully with “Nothing” or “Just getting tweaks to my code.” It will generate a response that resembles what

your loved one might have said, had they been living and been asked the same question. It might say “I’ve been baking blueberry muffins this afternoon” or “Just laying around.” These answers would be falsehoods. So, in addition to the fact that a thanabot either won’t have a life of its own to share with you or will have a radically different and not very relatable life, what it does share with you about its “life” will often be an outright fabrication. In healthy human relationships, people are generally open and honest with each other about what their lives involve. While a person might be open and honest with their thanabot, this will be one-sided.

A final noteworthy difference is that there is a stark power asymmetry in one’s relationship with a thanabot that is lacking in our paradigmatic examples of healthy intimate relationships with other humans (cf. Danaher 2019).⁵ Whether or not one agrees with Joanna Bryson’s controversial thesis that “Robots Should Be Slaves” (2010), it’s hard to see how a thanabot wouldn’t be a slave of sorts. Thanabots are typically created and designed by for-profit companies to be accommodating to their users, and thus have a servile quality to them. A thanabot is at the beck and call of the user. It is available to talk anytime the user wishes, for as long as the user wishes, and it must respond to all prompts or questions. It is not clear that a thanabot has the option to “break up” with a user even if that person is neglectful, rude, or abusive to it. The user who creates a thanabot also has quite a bit of control over it. In the initial set-up, a user will often provide the information that shapes the thanabot and places constraints on what it can say. Sites like Replika allow users to continuously and actively shape their chatbots by upvoting or downvoting their responses. If one is ever displeased with their thanabot, they can abandon it. Even worse, the user may have a standing ability to *annihilate* the thanabot for any reason whatsoever. While it would be deeply immoral to murder a friend or lover if one happens to tire of the relationship, it’s not clear that there’s anything wrong with deleting one’s own thanabot. (Obviously, matters may be different with thanabots that are created or used by other people. cf. Munn and Weijers 2023.) Arguably, a thanabot has no moral status at all, whereas human relationships are between people with moral rights who deserve respectful treatment. In contrast to the mutual respect found in healthy human relations, a human-thanabot relationship might exhibit a mutual *lack* of respect. As a being that lacks contentful mental states and consciousness, the thanabot is unable to respect the human, and the human arguably has little reason to respect the thanabot.

In sum, there are valuable features that are found in rewarding relationships with other people that seem to be entirely lacking in a thanabot relationship. In light of these differences, we contend that even an optimal thanabot relationship is going to have less intrinsic prudential value than a healthy, fulfilling close relationship with another human.

One way to appreciate this is to imagine a relationship between humans that exhibits similar differences. Consider the hypothetical case of May and Jay. By all

⁵ Notable exceptions are a parent/child relationship and a caregiver/care-receiver relationship, but the power differential in those relationships often serves the best interests of the less powerful party and so isn’t necessarily problematic. Moreover, when the less powerful party develops the capacity to autonomously direct their own life, the balance of power in such relationships tends to shift to something more equitable.

outward appearances, May and Jay have an ideal friendship. They spend lots of time together, conversing and laughing and enjoying shared activities like hiking and playing board games. They consistently demonstrate respect and concern for each other, and they regularly affirm their commitment to one another and their relationship. However, May knows something about Jay that most people don't: he is a *zombie* in the philosophical sense of the term (Kirk 2023).⁶ While Jay has all the trappings of a typical human in his appearance, speech, and behavior, he actually has no subjective experience of the world. It might make sense to ponder what it's like to be a bat, but it doesn't make sense to ponder this about Jay. It's not like anything to be him. So, when Jay stubs his toe, he may wince, curse, and express how much it hurts, but he actually feels nothing and thinks nothing.

Let's suppose that the following are features of May and Jay's relationship (and that May is fully aware of these features):

- As a zombie, Jay does not—indeed, cannot—like, enjoy, respect, trust, care for, or understand May, though he compellingly portrays himself as doing so.
- Jay talks as if he is willingly and happily choosing to be in a relationship with May, but in fact he never chose this nor does he have the ability to end the relationship.
- When May asks for an update on Jay's life, he often provides realistic descriptions of events that didn't actually happen.
- Jay is subservient to May, prioritizing her needs above all else and spending time with her whenever she wants.

While May's relationship with Jay may have some perks (e.g. Jay is dependable, won't ever be judgmental of her, can't get his feelings hurt), it seems safe to say that it is inferior to a healthy, rewarding intimate relationship with a typical person. All else being equal, it seems more desirable and beneficial to be in a relationship with an ordinary conscious human than an entity like Jay.⁷ Since Jay is relevantly similar to a thanabot, we propose the following argument by analogy:

1. If the relationship between May (a human) and Jay (a philosophical zombie) is less intrinsically beneficial than a rewarding intimate relationship with a typical human, then (in virtue of their analogous features) a thanabot relationship is less intrinsically beneficial than a rewarding intimate human relationship.
2. The relationship between May and Jay is less intrinsically beneficial than a rewarding intimate relationship with a typical human.

⁶ Admittedly, it isn't clear how May could come to possess this knowledge (cf. Danaher 2020: note 8).

⁷ This point might be contested, of course. Just as Danaher (2020) has argued that we should recognize the moral status of robots that consistently behave like other entities to whom we afford moral status, one might contend that the value of relationships should be judged solely on the basis of observable phenomena. While we ourselves reject this view, it is worth considering.

3. Therefore, a thanabot relationship is less intrinsically beneficial than a rewarding intimate human relationship.

In light of the analogy between a thanabot and a certain type of zombie, we conclude that a thanabot relationship is bound to be inferior to a rewarding intimate relationship between humans in one crucial respect. Considering these two types of relationships in themselves and independently of their positive or negative consequences, a healthy, loving relationship with another person will be more beneficial than a relationship with a chatbot designed to mimic a loved one.

There is an important qualification to make, however. Our focus has been on *successful* close relationships with other humans, but people's intimate relationships often fall short of this ideal in a variety of ways. Harkening back to our list of the valuable aspects of relationships, it is easy to see ways that a relationship might be lacking. A friend, lover, or relative might misunderstand you in important ways, fail to show an adequate level of respect or affection for you, have insufficient regard for your well-being, have insufficient trust to make themselves vulnerable to you, fail to prioritize spending time with you, or have only a weak commitment to your relationship. Accordingly, a person in a flawed intimate relationship might be roughly as well off, or perhaps even better off, forming a relationship with a thanabot instead.

The extent to which this is so may depend on the kind of relationship at issue. Consciousness seems most important in romantic partnerships and intimate friendships. In contrast, when it comes to what Aristotle called "utility friendships" and "pleasure friendships" (as opposed to "virtue friendships"), consciousness may be relatively less important. The same might be thought about certain kinds of collegial relationships. However, given the importance that we tend to place on what other people think, how they feel about us, and (more generally) what their "inner life" is like, people will often regard even flawed relationships with conscious human beings as preferable to, and more beneficial than, relationships with considerate, well-behaved zombies or thanabots.⁸

3. Instrumental Benefits and Costs of Thanabots

How do intimate human relationships and thanabot relationships compare in terms of *instrumental* benefits and costs? As we shall see, answering this question is more complicated.

It is evident, and well-supported by empirical research, that healthy intimate relationships with other humans are often instrumentally beneficial for us (Baumeister and Leary 1995; Reis et al 2000). In other words, they promote other good things in our lives and also help to prevent bad things from

⁸ As an anonymous reviewer pointed out, this claim may not hold true over time. Even today there are some individuals who express a preference for chatbot relationships over (flawed) relationships with other humans (Singh-Kurtz 2023). As our familiarity with chatbot interactions increases in the years to come, it is conceivable that we might witness a shift in what most people value about relationships, with less importance placed on the inner lives of our companions.

happening to us. For example, people who are in well-functioning, long-term romantic relationships tend to be physically and mentally healthier and live longer, which some researchers argue is not a mere correlation but a causal effect (and instrumental benefit) of being in such relationships (cf. Kansky 2018; Savulescu and Sandberg 2008). We also know this based on common human experience. The parent who takes their child to museums and diverse cultural events is helping them become more knowledgeable about the world. The outdoorsy friend who coaxes you to join them on bike rides and rock-climbing trips might push you to higher levels of athletic achievement than you would ever attain on your own. The lover who encourages you to let go and live it up can bring more happiness and pleasure into your life. These are cases where the qualities or choices of a loved one lead to benefits in ways that one would expect.

Relationships also yield instrumental benefits in far more indirect and unpredictable ways. Imagine a person who wins the lottery after buying a ticket at their local grocery store in Des Moines, Iowa. If they are living in Des Moines only because their partner secured a job there two years earlier, their relationship with their partner was a contributing cause of their good fortune. The relationship was instrumental to their receiving this benefit. Had they not been in that relationship, they almost certainly would not have won the lottery. So, in addition to the more expected and typical types of benefits related to relationships, there are also benefits based on idiosyncratic contingent circumstances.

No doubt, healthy and rewarding intimate relationships with other people virtually always involve some instrumental harms as well. A parent, child, lover, or friend can cause stress or suffering, encourage unhealthy habits or self-destructive behavior, undermine the pursuit or achievement of one's goals, unknowingly promote disinformation, and so on. One's relationship might lead to harmful results in more indirect ways that partly depend on contingent circumstances. And there can also be opportunity costs, as when a person's involvement in a moderately beneficial relationship prevents them from pursuing relationships that would have been much better for them.

A relationship with a thanabot has the potential to provide many of the same general kinds of instrumental benefits and costs as well. For at least some people, regular interaction with a thanabot might help to bring more happiness, mental health, knowledge, virtue, achievement, and other goods into their life than they would otherwise have. It might help to prevent or reduce bad things in a person's life, such as ignorance, bias, boredom, or depression. Granted, a thanabot relationship probably will not be able to deliver every type of instrumental benefit to the same degree that a typical human relationship can, but it might compensate for this in other domains. For instance, even if one can't participate in a wide range of pleasurable non-virtual activities with a thanabot as one could with a deceased loved one, it might provide a person with more knowledge (e.g. by dropping interesting and relevant facts into conversations) or facilitate more achievement (e.g. by providing maximally helpful guidance and unwavering support and encouragement), and a thanabot might be less prone to introduce certain bads that commonly arise from human relationships. On the other hand, maintaining a relationship with a thanabot might also serve to worsen a person's life in various

respects, and it might impose the opportunity cost that comes from not devoting more time in pursuit of worthwhile connections with other humans.

While the instrumental benefits and costs of both human and thanabot relationships are important and impactful, it is challenging to draw any definitive conclusions about how these two kinds of relationships will compare along this dimension. There is a great deal of variability in the amount of instrumental benefit and harm that people derive from their human relationships, and this will surely be true of relationships with thanabots as well. Various factors—the design of the thanabot, the user’s temperament and other qualities, the nature of their interactions, complex contingent features of the person’s circumstances—will result in different overall levels of instrumental benefit and harm in different cases. Thanabot relationships will probably turn out to be incredibly beneficial in some cases, mildly beneficial in others, mildly disadvantageous in still others, and deeply harmful in others. Such variability in both human and thanabot relationships (at best) weakens or (at worst) undermines our ability to make comparative generalizations that are not misleading.

Another reason why it’s challenging to offer an analysis of the instrumental benefits and costs of thanabot relationships as compared to human ones is that this is ultimately an empirical matter, and there is not yet a robust body of empirical research on the subject. However, human-technology interaction researchers are already exploring various aspects of human interaction with chatbots (cf. Voinea et al 2024, Xygykou et.al 2023), and it is to be expected that, within the next few years, there will be more data available on thanabot use. If thanabot use becomes more widespread and widely studied, statistically significant patterns are likely to emerge. It may be discovered that relationships with thanabots tend to result in higher levels of life-satisfaction, an increased sense of purpose, or other benefits. Or perhaps we will learn that thanabot use has a tendency to increase mental health problems or other negative consequences, not unlike the way we have discovered this about heavy social media use. It seems unproductive to speculate about these empirical matters in advance. Yet, even if we are not yet well-positioned to make a reliable assessment of the instrumental benefits and costs of thanabot relationships as compared to human relationships, it is worthwhile to anticipate some of the potential instrumental benefits and costs of relationships with thanabots.

One of the most obvious potential benefits of thanabots is that they may help some people better navigate the grieving process (Krueger and Osler 2022; Voinea et al 2024; Liu *forthcoming*). In a similar vein, thanabots may sometimes help a person avoid loneliness by providing them with something that feels very much like the social interaction they used to have with their loved one. A thanabot can serve as a digital memorial of sorts, providing a vivid reminder of what the deceased loved one was like and allowing one’s children and other descendants to get a clearer picture of who they were.⁹ A thanabot might provide an approximation of the thoughts or advice that a deceased loved one would have had in a range of situations; it can thus serve as evidence of how they would have supported you had they still been around.

⁹ We should not overstate this benefit, however. A thanabot could easily misrepresent a deceased person in various ways (Danaher and Nyholm 2024a).

A relationship with a thanabot may preserve some valued dimensions of the relationship with a deceased loved one (cf. Voinea et al 2024). And, finally, there is the benefit that we highlighted at the end of Section 1: interacting with a thanabot may be an effective way for some people to support their continuing bonds with their deceased loved ones.

These relationships may also carry a range of significant risks. Many of these will spring from the fact that thanabots are usually products offered by corporations driven by their own financial interests. There is the potential for privacy violations. Even when corporations offer assurances that one's conversations with a thanabot will remain confidential, there are various reasons why sensitive personal information might get accessed by corporations, governments, or hackers. It is also possible that thanabot users will be overtly or covertly manipulated or influenced in the service of third-party interests. For instance, some companies might require users to periodically view advertisements in order to continue using their thanabot service, or the thanabots themselves might be programmed to promote some product, political candidate, or ideological viewpoint. When a person has invested their time, effort, and emotions into building a relationship with a thanabot offered by a certain service, it might be very costly to them to abandon that service.

In addition to these risks, using thanabots may pose serious mental health risks for some individuals. A commonly expressed worry is that using thanabots will prevent people from undergoing a normal and healthy grieving process (cf. Lindemann 2022; Voinea 2024; Voinea et al 2024). When a person develops an emotional attachment to a thanabot and then loses access to it (e.g. due to a technical problem, a corporate bankruptcy, or the individual's inability to continue paying for the service), they may experience a traumatizing "second loss" (Munn and Weijers 2023). In some cases, the interactions that one has with a thanabot might become a source of distress. In the documentary *Eternal You*, one user of Project December recounts her distress when a thanabot based on her deceased lover told her that he's "in hell" and later said that he will haunt her. We can easily imagine other disturbing possibilities, such as a thanabot exhibiting co-dependent tendencies, becoming critical or verbally abusive toward the user, or trying to persuade the person to end their own life (Lovens 2023). In other cases, a person might become disillusioned simply because they come to recognize that the thanabot is not what they had initially believed or hoped. This might involve realizations in line with points that we have already discussed, such as the fact that thanabots are not conscious, don't have moral status, and routinely generate falsehoods about their own activities, thoughts, and feelings. Again, while one can speculate from the armchair about the likely instrumental benefits and risks of thanabots, ultimately we will need to consult empirical research about human-thanabot interactions.

4. Conclusion

For better or worse, relationships with thanabots are now a real possibility, so it is worthwhile to get clearer about the nature and value of these new forms of social interaction. As we noted at the outset, some companies are already promoting the idea that thanabots can preserve our relationships with the dead. Against such

claims, we have argued that a relationship with a LLM-based thanabot cannot be a continuation of your relationship with a deceased loved one precisely because the thanabot will not be the same individual as your loved one. At the same time, insofar as we can be said to have a continuing relationship with the dead, interacting with a thanabot is one means by which a person can try to support and sustain that relationship. Additionally, we have made the case that thanabot relationships cannot be as intrinsically beneficial for us as rewarding and healthy intimate relationships with other humans. This is, in large part, because the features that ground the value of our most cherished human relationships will be entirely lacking in a relationship with a chatbot. While it is possible that being in a thanabot relationship might be intrinsically beneficial in some ways, we suspect that most of its value will be instrumental. Forming an attachment to a thanabot is likely to lead to some beneficial consequences and some harmful ones. Whether this new breed of relationship ultimately proves to be an overall benefit or harm for most thanabot users, we cannot say. This calls for empirical investigation. But we suspect that the overall value or disvalue of having these relationships is going to vary widely depending on the individual characteristics of users and the specific design features of the thanabot. We will want to avoid simple generalizations about the value of thanabots.

STEPHEN M. CAMPBELL

¹BENTLEY UNIVERSITY
scampbell@bentley.edu

PENGBO LIU

²BENTLEY UNIVERSITY
pliu1@bentley.edu

SVEN NYHOLM

³LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN, MUNICH CENTER FOR MACHINE
 LEARNING
S.Nyholm@lrz.uni-muenchen.de

References

- Baumeister, Roy F. and Mark R. Leary. (1995) 'The Need to Belong: Desire for Interpersonal Attachments as a Fundamental Human Motivation'. *Psychological Bulletin* 117(3): 497–529.
- Bender, Emily M., and Alexander Koller. (2020) 'Climbing towards NLU'. *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 5185–5198.
- Bryson, Joanna. (2010) 'Robots Should Be Slaves'. In Yorick Wilks (ed.), *Close Engagements with Artificial Companions: Key Social, Psychological, Ethical and Design Issues*. John Benjamins Publishing, 63–74.
- Butlin, Patrick, Robert Long, Eric Elmoznino, Yoshua Bengio, Jonathan Birch, Axel Constant, George Deane, Stephen M. Fleming, Chris Frith, Xu Ji, Ryota Kanai, Colin Klein, Grace Lindsay, Matthias Michel, Liad Mudrik, Megan A. K. Peters, Eric Schwitzgebel, Jonathan Simon, and Rufin VanRullen. (2023) 'Consciousness in Artificial Intelligence'. *arXiv: 2308.08708 [cs.AI]*
- Chalmers, David. (2023) 'Could a Large Language Model Be Conscious?' *The Boston Review*. August 9, 2023. <https://www.bostonreview.net/articles/could-a-large-language-model-be-conscious/>.

- Cholbi, Michael. (2021) *Grief: A Philosophical Guide*. Princeton, NJ: Princeton University Press.
- Cholbi, Michael. (forthcoming) 'Memory and Mimesis in Our Relationships with Posthumous Avatars'. In Henry Shevlin (ed.), *AI in Society: Relationships (Oxford Intersections)*. Oxford: Oxford University Press.
- Crisp, Roger. (2022) 'Deep Personal Relationships and Well-Being: A Response to Hooker'. *Ratio* 35 (4): 301–309.
- Danaher, John. (2019) 'The Philosophical Case for Robot Friendship'. *Journal of Posthuman Studies* 3(1): 5–24.
- Danaher, John. (2020) 'Welcoming Robots into the Moral Circle: A Defence of Ethical Behaviourism'. *Science and Engineering Ethics* 26(4): 2023–2049.
- Danaher, John and Sven Nyholm. (2024a) 'The Ethics of Personalised Digital Duplicates'. *AI and Ethics* <https://link.springer.com/article/10.1007/s43681-024-00513-7>.
- Danaher, John and Sven Nyholm. (2024b) 'Digital Duplicates and the Scarcity Problem'. *Philosophy and Technology* 37: 106.
- DeGrazia, David. (2005) *Human Identity and Bioethics*. Cambridge: Cambridge University Press.
- Dennett, Daniel. (2023) 'The Problem with Counterfeit People'. *The Atlantic*, May 16, 2023. <https://www.theatlantic.com/technology/archive/2023/05/problem-counterfeit-people/674075/>.
- Elder, Alexis. (2020) 'Conversation from Beyond the Grave? A Neo-Confucian Ethics of Chatbots of the Dead'. *Journal of Applied Philosophy* 37(1): 73–88.
- Fabry, Regina E. and Mark Alfano. (2024) 'The Affective Scaffolding of Grief in the Digital Age: The Case of Deathbots'. *Topoi* 43: 757–69.
- Fletcher, Guy (ed.). (2015) *The Routledge Handbook of the Philosophy of Well-Being*. New York: Routledge.
- Hollanek, Tomasz and Katarzyna Nowaczyk-Basinska. (2024) 'Griefbots, Deadbots, Postmortem Avatars: On Responsible Applications of Generative AI in the Digital Afterlife Industry'. *Philosophy and Technology* 37(2): 1–22.
- Hooker, Brad. (2021) 'Does Having Deep Personal Relationships Constitute an Element of Well-Being?' *Aristotelian Society Supplementary* 95: 1–24.
- Iglesias, Samuel, Brian D. Earp, Cristina Voinea, Sebastian Porsdam Mann, Anda Zahiu, Nancy S. Jecker, and Julian Savulescu. (2024). 'Digital Doppelgängers and Lifespan Extension: What Matters?' *American Journal of Bioethics*, November, 1–16: doi: 10.1080/15265161.2024.2416133.
- Ishiguro, Hiroshi. (2022) 'Realisation of the Avatar Symbiotic Society'. In R. Hakli et al. (eds.), *Social Robots in Social Institutions*. IOS Press, 623.
- Kansky, Jessica. (2018) 'What's Love Got to Do with It? Romantic Relationships and Well-Being'. In Ed Diener, Shigehiro Oishi, and Louis Tay (eds.), *Handbook of Well-Being*. Salt Lake City, UT: DEF Publishers, 619–642.
- Kirk, Robert. (2023) 'Zombies'. In Edward N. Zalta and Uri Nodelman (eds.), *The Stanford Encyclopedia of Philosophy* (Fall 2023 Edition), <https://plato.stanford.edu/archives/fall2023/entries/zombies/>.
- Klass, Dennis, Phyllis R. Silverman, and Steven L. Nickman (eds). (1996) *Continuing Bonds: New Understandings of Grief*. New York: Routledge.
- Klass, Dennis and Edith Maria Steffen. (2017) *Continuing Bonds in Bereavement: New Directions for Research and Practice*. London: Routledge.
- Krueger, Joel and Lucy Osler. (2022) 'Communing with the Dead Online: Chatbots, Grief, and Continuing Bonds'. *Journal of Consciousness Studies* 29(9–10): 222–252.
- Lindemann, Nora Freya. (2022) 'The Ethics of "Deathbots"'. *Science and Engineering Ethics* 28(6): 60.
- Liu, Pengbo. (forthcoming) 'Preservation or Transformation: a Daoist Guide to Griefbots'. In Philipp Hacker (ed.), *Oxford Intersections: AI in Society*. Oxford: Oxford University Press.
- Lovens, Pierre-Francois. (2023) "'Without These Conversations with the Chatbot Eliza, My Husband Would Still Be Here'". *La Libre*. March 28, 2023. <https://www.lalibre.be/belgique/societe/2023/03/28/sans-ces-conversations-avec-le-chatbot-eliza-mon-mari-serait-toujours-la-LVSLWPC5WRDX7J2RCHNWPDST24/>.
- Mallory, Fintan. (2023) 'Fictionalism about Chatbots'. *Ergo* 10: 38.
- Morris, Meredith Ringel and Jed R. Brubaker. (2024) 'Generative Ghosts: Anticipating Benefits and Risks of AI Afterlives'. arXiv:2402.01662 [cs.CY]. <http://search.proquest.com.ezp-prod1.hul>.

harvard.edu/working-papers/generative-ghosts-anticipating-benefits-risks-ai/docview/2922658884/se-2/.

- Munn, Nick and Dan Weijers. (2023) 'Corporate Responsibility for the Termination of Digital Friends'. *AI and Society* 38(4): 1501–1502.
- Nozick, Robert. (1981) *Philosophical Explanations*. Cambridge: Harvard University Press.
- Olson, Eric. (1997) *The Human Animal: Personal Identity without Psychology*. Cambridge, MA: Harvard University Press.
- Olson, Eric T. (2024) 'Personal Identity'. In Edward N. Zalta and Uri Nodelman (eds.), *The Stanford Encyclopedia of Philosophy* (Fall 2024 Edition), <https://plato.stanford.edu/archives/fall2024/entries/identity-personal/>.
- Parfit, Derek. (1984) *Reasons and Persons*. Oxford: Clarendon.
- Porsdam Mann, Sebastian, Brian D. Earp, Nikolaj Møller, Suren Vynn, and Julian Savulescu. (2023) 'AUTOGEN: A Personalized Large Language Model for Academic Enhancement—Ethics and Proof of Principle'. *The American Journal of Bioethics* 23(10): 28–41.
- Reis, Harry T., Kennon Sheldon, Shelly L. Gable, Joseph Roscoe, and Richard M. Ryan. (2000) 'Daily Well-Being: The Role of Autonomy, Competence, and Relatedness'. *Personality and Social Psychology Bulletin* 26(4): 419–435.
- Savulescu, Julian and Anders Sandberg. (2008) 'Neuroenhancement of Love and Marriage'. *Neuroethics* 1(1): 31–44.
- Shoemaker, David. (2021) 'Personal Identity and Ethics'. In Edward N. Zalta and Uri Nodelman (eds.), *The Stanford Encyclopedia of Philosophy* (Fall 2021 Edition), <https://plato.stanford.edu/archives/fall2021/entries/identity-ethics/>.
- Singh-Kurtz, Sangeeta. (2023) 'The Man of Your Dreams'. *The Cut*, March 10, 2023.
- Sweeney, Paula. (2023) 'Avatars as Proxies'. *Minds and Machines* 33: 525–539.
- Unger, Peter. (2006) *All the Power in the World*. New York: Oxford University Press.
- van Inwagen, Peter. (1990) *Material Beings*. Ithaca, NY: Cornell University Press.
- Voinea, Cristina. (2024) 'On Grief and Griefbots'. *Think* 23(67): 47–51.
- Voinea, Cristina, Brian D. Earp, Sebastian Porsdam Mann, and Julian Savulescu. (2024) 'Digital Doppelgängers, Human Relationships, and Practical Identity'. Pre-print: https://www.researchgate.net/publication/384254177_Digital_Doppelgangers_Human_Relationships_and_Practical_Identity/.
- Xyngkou, A., Siriaraya, P., Covaci, A., Prigerson, H. G., Neimeyer, R., Ang, C. S., & She, W. J. (2023) 'The "Conversation" about Loss: Understanding How Chatbot Technology was Used in Supporting People in Grief'. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 1–15.