



The ethics of break-up chatbots

Pilar Lopez-Cantero¹ 

Accepted: 1 November 2025

© The Author(s), under exclusive licence to Springer Nature B.V. 2025

Keywords Chatbot · Grief · Break-ups · Emotion · Continuing bonds · LLM · AI

Romantic break-ups are common scenarios of suffering, transformation, and moral concern. While bereavement grief is a common object of philosophical analysis, the analogous phenomenon of grief over unchosen romantic dissolutions has only recently begun to be explored. At the same time, technological developments are reshaping how people navigate the end of relationships. Large language model (LLM)-based systems now make it possible to simulate interaction with an ex-partner through ‘break-up chatbots’, which draw from user-input data to approximate the conversational style of one’s former partner. These tools raise an urgent normative question: what, if anything, is wrong with using artificial stand-ins for romantic partners after they end the relationship?

This paper offers the first systematic normative analysis of break-up chatbots. I focus on their use in sustaining continuing bonds with ex-partners, by analogy with the use of ‘deathbots’ in bereavement. My central claim is that, when designed with the aim or the potential to sustain continuing bonds, break-up chatbots are not merely contingently harmful but morally noxious by design. They function as hostile scaffolds that (i) place users in harmful imaginary relationships, (ii) foster noxious romantic tendencies that corrupt moral agency, and (iii) embed users in oppressive niches that they themselves co-create. I proceed as follows. Section 1 briefly describes the paradigm of romantic terminations I focus on—unchosen break-ups—, as well as the technologies that count as relevant ‘break-up chatbots’. Then, Sect. 2 surveys the potential harms and benefits of break-up chatbots, and distinguishes these from the harms resulting in a specific use context: the use of break-up chatbots designed with the (potential) aim to sustain continuing bonds with ex-partners. Sections 3 to 5 develop this idea by showing how break-up chatbots act as hostile scaffolds in three

✉ Pilar Lopez-Cantero
pilar.lopezcantero@uantwerpen.be

¹ Department of Philosophy, University of Antwerp, Antwerp, Belgium

ways: they place users in harmful imaginary relationships, foster noxious romantic tendencies that corrupt moral agency, and embed users in oppressive niches they co-create. Section 6 concludes by reflecting on the implications of my argument for both the philosophy of technology and the philosophy of grief.

1 Introduction: Break-ups in the digital age

Romantic break-ups are heterogeneous. They can be the result of the end of love, or be motivated by love (for example, if one of the partners judges that the relationship is standing in the way of the other person's goals).¹ Relationships can end explosively or fizzle out; their conclusion can come as complete surprise or be a long time coming. There is, however, a paradigm of relationship dissolution that is the most common object of academic research, laypeople's preoccupation, and fictional plots: break-ups where one of the partners decides to end the relationship against the wishes of the other. Romantic dissolutions of this nature often result in emotional pain for the person experiencing them, as noted by philosophers (Ben Ze'ev & Goussinsky, 2008; Milligan, 2011; Brogaard, 2015; Lopez-Cantero, 2018; Jenkins, 2022; Betzler, 2024), sociologists (Illouz, 2012), neuroscientists (Kross et al., 2011; Woo et al., 2014) and psychologists (Lee et al., 2011; Rhoades et al., 2011). This is the paradigm of romantic terminations I focus on here, referring to them as unchosen, unilateral, or non-consensual break-ups. The emotional pain of unchosen break-ups is often managed through a mix of social interactions—for example, by talking to friends—and self-reflection, such as through therapeutic processes or journalling. The regulation of post-dissolution psychological suffering often involves technology: people use phones to communicate with friends and therapists, follow self-care steps in mental health apps, or interact with strangers in online forums. The popularisation of products and services based on Large Language Models (including LLM-powered chatbots) is set to significantly amplify the role of technology in romantic dissolutions.²

For example, people are using LLM-powered chatbots to help them initiate a break-up, like writing break-up texts, or practising break-up conversations with their current partners. Some chatbots specifically designed with (self-)therapeutic purposes are also being employed during romantic dissolutions—for instance, they provide mental health information, or coaching and advice through interaction with the user. When used with an interactive, therapeutic aim, chatbots sometimes work from simple user prompts (such as the question 'how can I try to move on from my ex?'), and deliver answers from the perspective of a neutral third party in the relationship

¹ My argument pertains to romantic break-ups only. Dissolutions of other kinds of relationship, such as friendship break-ups, are too under-researched to address without straying from the topic. Also, 'romantic' shall be understood as an umbrella term for committed relationships socially acknowledged as such, within or outside of the dyadic and cis/hetero/allosexual paradigm, and independently of their legal status. I remain agnostic about aromantic individuals, insofar as their committed relationships could constitute romantic partnerships, close friendships, or a *sui generis* aromantic bond.

² I have chosen not to name specific products in this article; only describing what the relevant products (promise to) do and their possible uses. Given the potential harms of the products that I describe here, and that many of them are provided by for-profit firms, I prefer to not publicize them.

(such as a therapist, a friend or a search engine). Other service providers employ the user's private communications with their ex-partners to deliver a 'data-driven analysis' of the relationship (or at least so they promise). These are some of the ways in which people are starting to use LLM-powered products and services at the end of romantic relationships—but these are not my focus here.

Instead, I reserve the term *break-up chatbots* to refer to a category of products that are designed in such a way that they can produce digital duplicates of an ex-partner through the user's input of person-specific data, so that the user may interact with the duplicate as a stand-in for their former partner.³ At the time of writing, there are no publicly available services designed with the sole aim of creating digital duplicates of an ex-partner—only one service on waitlist without public information about its launch date. However, there are several companies offering so-called 'AI companions' or social chatbots that users can fine-tune according to their preferences in order to resemble their ex-partners.⁴ General-purpose conversational chatbots can also be used in this way, although they generally offer less variety of communication methods and add-on outputs than social chatbots (like personalised avatars or voice calls).

Break-up chatbots can be fine-tuned to varying degrees. The coarsest models follow a user's description prompt, for example, 'My ex is kind and generous, but she did not explain why she left me. Can you please talk to me as if you were her?'. Answers are built on large, pre-existing collections of text—including internet writings, curated conversational datasets, and human-annotated training examples—without any particular link to the individual that is being represented in the conversation. On the other side of the spectrum, the best fine-tuned break-up chatbots can more closely resemble a particular ex-partner's conversational style. This is achieved by the user's input of relevant personal data (such as chat and email history, social media posts, and other writings) and/or pretend 'programming' character traits, preferences or choices of the former partner either directly (by selecting from a menu or providing descriptions) or through a technique called reinforcement learning from human feedback (RLHF).⁵ Through RLHF, the model can then be further refined through the user's evaluation of generated responses (such as scoring accuracy, or suggesting alternative responses that the individual might give in that situation). The better the fine-tuning, the higher the likelihood that users feel as if they were, in fact, talking to their ex-partners. It shall be noted that these products also draw from large pre-existing bodies of text unrelated to the data input directly by the user.

If anecdotal reports from online forums are taken at face value, these technologies are not merely theoretical. People are already using break-up chatbots (in the specific sense I understand them) during romantic dissolutions. Even if these reports were

³ See Danaher and Nyholm (2025) for an accessible philosophical explanation of the digital duplicate technology and a proposal for a general ethical framework; see Ho et al. (2025) for a review of the wider literature on 'romantic AI companions' (definition, uses, and potential benefits and pitfalls).

⁴ LLM-powered chatbots do not *resemble* anyone, because they are not people; however, the chatbot's responses can be attuned to data from an individual, so that the chatbot's response to prompts can start to resemble a particular individual's writing style. Here I use 'resemble an individual' as shorthand for ease of reading.

⁵ See González Barman et al. (2025) for a philosophical discussion of RLHF which also addresses the potential problems (both epistemic and ethical) with this technique.

fabricated or exaggerated, it seems plausible that break-up chatbots could become more widespread in the near future—especially as LLM-powered chatbots become more technologically advanced. Research on normative questions regarding break-up chatbots (such as potential harms and wrongs to the user, to the person being replicated, or to society at large) is currently negligible. However, existing discussions on personal relationships, grief, and technology can help articulating some normative insights on the matter.

2 From contingent harms to hostile scaffolding

Although there is virtually no scholarship on break-up chatbots, they are almost analogous to what is popularly known as ‘deathbots’, which are technologically similar to break-up chatbots but used when a loved one (not necessarily a romantic partner) passes away. Several philosophers working on the normativity of deathbots have expressed concerns about potential harms to the autonomy and dignity of the deceased (Lindemann, 2022), or the risk of replacing the deceased as a loved one (Buben, 2015; Stokes, 2021); while others have highlighted their potentially helpful role in healthy grief for the bereaving party (Osler and Krueger, 2022). According to Regina Fabry and Mark Alfano (2024, 764—765), the potential harms and benefits of deathbots are not inherent to the technologies themselves, but to the use context. For Fabry and Alfano, the moral permissibility of the use of deathbots depends on factors like relationship history, or the timeframe and nature of the data that is used. Still, determining exactly how they are harmful or beneficial is a useful starting point. One approach to the ethics of break-up chatbots could then be to outline the most common scenarios in which these technologies may be prudentially or morally valuable and disvaluable, while acknowledging that benefits and harms are contingent.⁶

2.1 Potential harms of break-up chatbots to the represented person

To begin with, Fabry and Alfano (2024, 765) argue that deathbots can be harmful if they represent individuals in a derogatory, inaccurate or superficial manner, which can all be forms of disrespect. This worry applies to break-up chatbots too, since they may be used to insult or mock one’s ex-partner, and they can be fine-tuned to produce utterances that the individual in question would repudiate, thereby subjecting them to disrespect through inaccurate representation. Misrepresentations can be grave and morally repugnant, such as racist remarks; but they also include seemingly more frivolous remarks that would hurt the individual insofar they target their practical identities—such as endorsing their favourite sports team’s rival.

Further, like deathbots, break-up chatbots can enable invasions of the represented person’s mental privacy, understood as “control over access to our neural data and to the information about our mental processes and states that can be obtained by

⁶ For clarity, I am separating prudential and moral value to differentiate between impact to well-being and impact to moral agency, or broader morally noxious consequences. It is worth noting that well-being may be morally valuable, so these may not be conceptually separate (see Fletcher, 2021).

analysing it” (Wajnerman Paz, 2021). When a chatbot is fine-tuned on personal messages, conversations, or social media posts, it can produce outputs that expose or simulate the thoughts, emotions, and preferences of the person being represented. Even if these outputs are not perfectly accurate, they may still reveal aspects of the individual’s inner life in ways they never consented to, thus constituting a violation of mental privacy.

Meaningful consent, alongside a firm commitment from the user to stick to certain constraints, can plausibly prevent these wrongs to the person represented. People can make agreements on how others may attempt to re-create them, for example, by setting up rules (such as avoiding disrespectful or degrading conversations), or by promising that the chatbot will be used alongside more faithful representations (such as pictures, videos, letters, or conversations with people who knew the person). So provisionally, it can be claimed that as long as there is meaningful consent that prevents harms and wrongs to the person represented, the use of break-up chatbots is morally permissible, particularly in scenarios where they are bound to be prudentially and morally valuable for the user.

2.2 Potential benefits of break-up chatbots for the user

Break-up chatbots can be used in the process of moving on from romantic dissolutions, in the sense of helping overcome the peak of emotional suffering in the immediate aftermath of the break-up. They can do this in two ways. First, Betzler (2024) argues that people experiencing unilateral break-ups have a moral right to receive an explanation for their partner’s decision to end the relationship, and to participate in the decision of ending the relationship. Break-up chatbots can be an effective tool when people initiating break-ups are not in the position to offer an explanation, or an explanation that will satisfy their ex-partner as sufficient. Getting to comprehend—or at least accept—the reasons for a unilateral break-up can be a long process. Being able to go through the explanation for the break-up for as long and as often as needed can help people reach that point, and even feel as if they are actively participating in the decision. Allowing an ex-partner to fine-tune a chatbot to resemble one’s conversational style can then be considered as a way to indirectly provide an explanation and participation, which are at least prudentially valuable for the person using it—and morally valuable if, as Betzler argues, they have a right to those goods.

Second, in unilateral break-ups, it is common for the person who did not make the choice to still be very much in love, and being unable to be with the person they love is a source of emotional suffering. But also, it can be a source of disorientation (Harbin, 2016). When people enter a romantic relationship, they often exhibit what psychologists call ‘progression bias’, which is “a bias towards pro-relationship decisions” (Joel and MacDonald, 2021, 317). The term ‘bias’, although correct, fails to capture the role of relationships as sources of meaningfulness, and that this function contributes to their continuation. Jones (2008, 274) argues that being in love is the property of being in a love trajectory, that is, a diachronic process in which events are constitutionally dependent on each other—a narrative. When an individual is in a love narrative, they will be more likely to have attitudes and behaviours that continue the trajectory than those that do not, since these will be more meaningful for

them. Doug McConnell calls this ‘narrative momentum’: “some futures make more sense than others, and we are inclined to enact our self-narratives so that they make sense” (McConnell, 2016, p. 308). When a person falls out of love, their sense of reality stops being bounded to the relationship, which was a key focus of orientation through the existential feeling of intelligibility—of things ‘making sense’ like they are (Lopez-Cantero, 2023; drawing from Ratcliffe, 2023). However, if the relationship ends and they are still in love, there is a mismatch between what they feel that makes sense, and what is the case. People going through unchosen break-ups often need to exert their will in order to fall out of love, since they remain oriented towards the relationship. They still have the existential feeling that a world that makes sense is a world with the relationship in it, even if they know that the relationship does not exist—thus, their phenomenological experience does not align with the actual world, where the relationship has ended. Existential feeling can be shifted through narrative projections, which allow the sufferer to experience positively valenced emotions in relation to a future where their existential feeling has changed (Bortolan, 2021)—in fact, the prudential benefits of creating self-narratives after romantic terminations are supported by empirical evidence (Bourassa et al., 2017).

So there are two potential uses of break-up chatbots that can be prudentially and morally valuable for the user: when they are used to understand the reasons for the relationship ending, or when they are used to advance the aim of falling out of love—with the ultimate aim of moving on. Like it has been established, whether users may in fact achieve these benefits, and whether they can do it in a morally permissible way that avoids harms to the represented person, depends on contingent factors. In the remainder of this paper, however, I will focus on a specific use of break-up chatbots that makes them morally noxious by design. This is when break-up chatbots are used to build so-called *continuing bonds*, as part of the grieving process that often follows unchosen romantic dissolutions.

2.3 Break-up chatbots and continuing bonds: Harm by design

There is growing consensus in the philosophy and psychology of grief that the loss of a loved one does not necessarily sever one’s relationship with them. Instead, bereaving individuals can preserve elements of their relationship by forming continuing bonds with the deceased (e.g. Klass, 2006; Valentine, 2008; Norlock, 2017; Klass & Steffen, 2018; Higgins, 2013; Cholbi, 2022). This can be done by remembering them, holding memorials, or even having imagined conversations with them. Most philosophical accounts of grief consider continuing bonds as a component of good grief; that is, grieving processes that do justice to the magnitude of the loss of a loved person, and the significance of the loving relationship that death has disrupted, without rendering the griever unable to continue with their life.⁷ Good grief is not only prudentially good, by fostering the griever’s well-being; it is also morally good, since

⁷ Some accounts seem to consider continuing bonds necessary components of good grief (e.g. Higgins, 2013). However, this would be incompatible with non-bereavement grieving processes where there is no existing relationship with the lost object, such as the grief over involuntary childlessness (Ratcliffe & Richardson, 2023) or religious deconversion (Williams, 2024).

it is the normatively appropriate response to the fact that what has been lost is of great importance on a cognitive, affective, and existential level for the griever.

Joel Krueger and Lucy Osler argue that deathbots can function as positive cognitive-affective *scaffolds*, which are environmental resources with directly influence and shape mental processes in a positive sense (Colombetti & Krueger, 2015).⁸ Specifically, they can scaffold “healthy griefwork”, which Krueger and Osler (2022, 233) define in similar terms to those I have used to describe ‘good grief’ (the latter term is, in my view, clearer about its moral significance besides prudential considerations). According to Krueger and Osler, deathbots facilitate three types of ‘habits of intimacy’ that enable continuing bonds: conversational practices with loved ones, emotional regulation through interaction within intimate relationships, and the feeling of shared time. In that sense, if continuing bonds are components of good grief, and deathbots enable continuing bonds, then deathbots can facilitate good grief, which is both prudentially and morally valuable.

At first glance, this line of reasoning seems applicable to romantic dissolutions. Romantic break-ups can trigger a form of non-bereavement grief (Lopez-Cantero, 2018; Ratcliffe & Richardson, 2023), and so one might think that break-up chatbots could likewise support continuing bonds and thereby contribute to prudentially and morally valuable griefwork. I argue, however, that this extension fails. When designed with the aim or the potential to sustain continuing bonds, break-up chatbots are *morally noxious by design*. Instead of functioning as positive scaffolds, I argue that they function as *hostile scaffolds*, given their potential to harm and wrong the user, the represented person, and society at large (Timms & Spurrett, 2023). I show that although some of these harms seem contingent, break-up chatbots aim at bringing them about when used to build continuing bonds—which means that break-up chatbots are noxious by design, regardless of how successful they are in bringing about these harms. Specifically, I argue that when used to build continuing bonds, break-up chatbots are aimed at placing users in a harmful imaginary relationship, foster noxious romantic tendencies, and embed the user in an oppressive narrative niche that they themselves co-create.

3 Harmful imaginary relationships

Let us look further into the notion of continuing bonds, and what it is exactly that they preserve. First, it should be noted that continuing bonds after bereavement facilitate the preservation of elements of a relationship, but do not preserve the relationship itself, which has, in fact, ended (Millar & Lopez-Cantero, 2022, p. 432). Second, it has been argued that interactions with deathbots do not themselves constitute a continuation of the relationship one had with a bereaved partner, since the *persona* portrayed through the chatbot does not constitute a metaphysical continuation of the

⁸ In this paper, I do not focus on describing the functioning of scaffolds (i.e., I do not provide a step-by-step explanation on how environmental resources cause mental processes to happen). The works cited go into these mechanisms at length; I focus instead on describing the harms and wrongs that break-up chatbots can scaffold.

person they represent (Campbell et al., 2025). My view here is that interactions with break-up chatbots that sustain continuing bonds aim at placing the user in an imaginary relationship that is ultimately harmful to the user, even if it brings some benefits along the way.

The concept of imaginary relationships is described by Norlock (2017), whose work is foundational for the philosophical debate on continuing bonds. For Norlock, continuing bonds practices after bereavement (visiting the deceased's tomb, having imagined conversations with them, reading their old letters) are features of *imaginal* relationships, which are real relationships that should not be confused with *imaginary* relationships. Relationships with the living, Norlock argues, have a large imaginal component: interactions that happen in the mind of the lover such as thinking about them, reflecting about their words, or planning certain behaviours informed by love for that person. These imaginal features may be preserved even when the loved person dies, through continuing bonds. To contrast imaginal relationships with imaginary ones, Norlock brings up Wilson, an anthropomorphised volleyball, which is treated as a friend by Tom Hanks's character in the film *Castaway*:

What makes the friendship imaginary instead of merely imaginal resides in the material facts of... the attribution of the capacity for reciprocity to an object, the entirely fictional product of a lonely man's wishful thinking. The imaginal dialogues that the castaway has with the volleyball are not informed by memories of Wilson's words, mementos, or behaviors of Wilson's, since a nonliving object can offer none of these things. (Norlock, 2017, p. 344)

The question that arises here is whether interactions with break-up chatbots are analogous to the imaginary relationship of Hanks's character with Wilson. On one hand, interactions with chatbots wrongly attribute a capacity for reciprocity of love to an inanimate object, so impossibility of reciprocity seems to favour the understanding of the relationship as an imaginary one. On the other hand, break-up chatbots are informed by memories of the loved person's words or behaviours, so (unlike volleyballs) they are not completely disconnected from a living, existing person (they are not wholly 'a product of a lonely man's wishful thinking'). This could suggest that relationships with break-up chatbots are not imaginary. However, a more careful consideration shows that despite the chatbot's connection to an existing person, break-up chatbot users are in an analogous situation to Hanks's character—and that unlike him, they are participating in a harmful imaginary relationship.

Following Norlock, when two people are in a romantic relationship, much of the relationship is imaginal in nature. These imaginal components of the relationship *with that person* can persist after one of them dies, enabled by continuing bonds. When this happens, the other member of the relationship (which is the object of love that persists) is still the deceased person, not the environmental resources or scaffolds (photographs, gravestones, letters) that facilitate continuing bonds. Equally, the imaginal components of a relationship with one's ex-partner may continue after a break-up: one might still find oneself wondering how the other person is spending their birthday, or finding it hard to imagine going solo to an extended family event. But when engaging with a break-up chatbot, they are not continuing the imaginal

relationship, since the other person is still alive, still acting, and still being influenced by the world and by others in ways that differ from the way the *persona* portrayed in the chatbot is. Instead of continuing the imaginal component of the extinct relationship with their ex-partner, break-up chatbot users build a new *imaginary* relationship with an object which is by nature incapable of reciprocity: the representation of their ex-partner, that is, the scaffold itself.

Now, the mere fact of facilitating imaginary relationships does not make break-up chatbots harmful, because imaginary relationships can be prudentially valuable. Hanks's character, who is stuck alone in a desert island, benefits greatly from his imaginary relationship with Wilson. Also, despite not calling them 'relationships', people have loving connections with all sorts of objects that do not have the capacity for reciprocity, such as sports teams (Archer, 2021) or ideals (Shpall, 2018). This can be a form of love that is valuable for its own sake, and may bring great happiness—i.e., prudential value—to the lover.⁹ Nevertheless, as I will argue, imaginary relationships with break-up chatbots are ultimately harmful for the user, because they lack wholehearted and authentic love.

Love is authentic when it is an expression of one's identity, and wholehearted when it is endorsed by the lover and not in conflict with their first-order desires.¹⁰ Relationships with break-up chatbots where users believe they are continuing their extinct romantic partnership lack both features, given that love is misdirected in these cases. In offline life, misdirected love happens when person A establishes a loving relationship with person B on the grounds of their love for person C. A classic literary instance of misdirected love appears in Edmond Rostand's *Cyrano de Bergerac* (1897), where Roxane believes she loves Christian but her love is in fact responding to Cyrano's personality, secretly presented to her through Christian. This dynamic mirrors chatbot-mediated intimacy: the user (like Roxane) establishes a relationship with a visible 'partner' (the chatbot), while their ex-partner is the object of the love that expresses their identity, and the one that they endorse as an expression of their first-order desires. Inauthentic and not wholehearted love in an imaginary relationship is harmful because it represents an unsuccessful attempt to engage in a meaningful project (a loving relationship with a particular, existing individual, and not just a fictional character or an imaginary friend), which turns out to be a failure as such (due to the lack of authentic and wholehearted love resulting from misdirection). Thus, a relationship based on love without these features contributes to meaninglessness in the user's life, given that it is a meaningless activity.¹¹ In that sense, even if imaginary

⁹ This is why alluding to the lack of consciousness of chatbots (e.g. Lindemann, 2022) is not enough to prove that these relationships are harmful. See Nyholm & Frank (2017) for a conceptual defense of mutual love with robots.

¹⁰ Although I loosely adopt Harry Frankfurt's (2004) notions of wholeheartedness and authenticity, this should not be understood as an endorsement of Frankfurt's account of love, given the difficulties to translate his view to romantic love (which he did not see as wholehearted or authentic). See McKeever (2019) for a detailed explanation of Frankfurt's missteps with regard to romantic love.

¹¹ Here I follow Wolf (2010), who describes two scenarios where people experience meaningless lives: those who take up pointless projects, and those whose meaningful projects are unsuccessful. I do not claim that the user's life is meaningless as a whole as a result of using break-up chatbots, since they may well engage in many other meaningful endeavours in other areas of their lives. Their life would be meaning-

relationships offer users prudential value (i.e., even if they feel good), they are ultimately harmful to the user by maintaining them engaged in an unsuccessful attempt at achieving meaningfulness.

A possible response to my view can be drawn from Krueger and Osler's fictionalist stance on deathbots. Here, I am working on the assumption that users feel as if the relationship with their partners continues through the interaction with chatbots. According to Krueger and Osler, deathbots users act as if they are interacting with their loved ones, but they know they are, in fact, interacting with a lifeless technological device (Krueger & Osler, 2022, 245—247). From this, it can be objected that people are mostly aware that their relationship with break-up chatbots is imaginary, and hence there is no problem of misdirected love (and subsequently of the lack of wholeheartedness and authenticity which leads to meaninglessness). Whether fictionalism is really a universal or even widespread stance among break-up chatbots users is an empirical question in need of further research (Fabry & Alfano, 2024, p. 764). Nevertheless, what makes break-up chatbots hostile scaffolds is that they are designed to make users *feel* as if they are interacting with their ex-partner, which in turn makes them *feel* like the relationship is in fact continued—regardless of what they *believe* to be the case.

Recall that relationships are sustained by narrative momentum. Momentum of a relationship is not preserved so much by what one believes ('I am/am not in a relationship right now') but what one *feels*. Being in love with someone means that some futures make more sense than others as one's own future, so for people who are in love, their sense of reality stays oriented towards the relationship (Lopez-Cantero, 2023). Experiencing an unchosen break-up then often means that there is a mismatch between what one experiences as making sense (the relationship continuing) and the state of affairs (the relationship having ended). By making efforts to maintain momentum, users fail to go through the process of "recognition and reorientation" (Ratcliffe et al., 2023, p. 328) that allows grievers to move on. This process often takes a long time and requires working through contradictory mental and bodily states (like the feeling like that the person is there while knowing they are not), complicated social interactions, and a general reorganization of one's world. Because the best fine-tuned chatbots will make users *feel* as if they are talking to their ex-partner, they will plausibly fail to make them feel as if they are not, in fact, in a relationship with the 'person' they are talking to. In other words, users might know the relationship is over and believe they are interacting with a chatbot, but given that narrative momentum is maintained, they might not *feel* like the relationship is actually done or that these interactions are not in fact components of said relationship. Hence, fictionalism does not prevent them from engaging in a meaningless activity.

Regardless of whether users believe or not that they are in a relationship with the break-up chatbot, the function to preserve narrative momentum is built into the design of these products, so the best fine-tuned break-up chatbots will, in fact, place users in an imaginary relationship—which as we saw above, is harmful because it is a meaningless activity. The point here is that break-up chatbots are designed to make

less only if it was entirely dedicated to interactions with the chatbot (which is not a far-fetched scenario, particularly as for-profit companies move to keep users indefinitely engaged).

users engage in a harmful endeavour or to facilitate this engagement, regardless of whether users are in fact harmed or not. For that reason, it does not matter whether break-up chatbots are successful in maintaining narrative momentum indefinitely, for a long time, or at all.¹² In sum, break-up chatbots function as hostile scaffolds when used to establish continuing bonds: they are morally disvaluable insofar they are specifically designed to keep the user engaged in a meaningless project.

4 Noxious romantic tendencies

Bonds built through break-up chatbots are purely self-interested, which may have potentially noxious moral consequences. This is a concern that has been raised in the first instance with respect to deathbots. It has been argued that relationships with deathbots treat the deceased person as a means to an end, given that they are valuable only insofar as they serve the users' aims (Stokes, 2021; Cholbi, 2025). In other words, relationships with deathbots (fictional or not) are utility-based. This is not necessarily bad, since utility friendships are not inherently wrongful, and as explained before, the relationship is not actually established with the dead individual (Campbell et al., 2025, p. 13). The worry is that the self-interested nature of interactions with deathbots may negatively influence people's ideals of relationships, and create expectations of constant availability and acquiescence from others (Fabry & Alfano, 2024, 766–767; Lindemann, 2022, p. 7; see also Jollimore, 2015, p. 142). I expand on this worry and take it further to claim that break-up chatbots function as hostile agential scaffolds, given their potential to corrupt moral agency by fostering noxious romantic tendencies.

Loving relationships require mutual improvisation, in which members react and adapt to their individual transformations and evolution over time, including surprising changes and reactions (Bagley, 2015). Continuing bonds with people who have passed away can involve mutual improvisation to some degree, since the bereaved continues shaping the legacy and memory of the deceased (alone and with others) and being shaped by these changes in ways that can be unexpected (Millar & Lopez-Cantero, 2022, p. 427). As explained in Sect. 1, LLM-powered chatbots can be refined through user feedback, so the most complex ones can indeed 'change' through interaction with the user. However, through their engagement with break-up chatbots, what the user is changing is not the other person (who has continued their life beyond the relationship) but the chatbot. Even if we understand this as a form of thin reciprocity (Krueger & Osler, 2022, p. 225), any 'transformation and evolution' that may appear to happen to the chatbot happens exclusively on the user's terms. The user will reasonably input feedback that makes the chatbot, and the relationship, adapt to their self-interested preferences.

Feminist philosophers have long argued that being excessively shaped by one's romantic partner is unjust as long as this asymmetry expresses, constitutes, and rein-

¹² Cholbi (2025) might be correct that relationships with chatbots eventually become stagnant and die off due to the lack of meaningfulness: this bears on the extent of the actual harm that break-up chatbots may cause to particular users, but not on the fact that they are designed to instantiate this harm.

forces social power imbalances between men and women (e.g. Friedman, 2003). When people shape chatbots that represent their ex-partners, they exercise absolute power over the chatbot's conversational style. That is exactly what feminist philosophers argue against. It would be outlandish to claim that a chatbot can be harmed or wronged by this exercise of power: instead, the claim is that break-up chatbots scaffold attitudes and behaviours directed at building and maintaining unjustly asymmetrical relationships.¹³

Besides fostering interactions aimed at self-interest that ultimately build and maintain unjust asymmetry, the use of break-up chatbots corrupts moral agency by diminishing the normative powers people have in romantic relationships—and this, in turn, also contributes to unjust asymmetry. Healey (2023, 166–167) argues that breaking up with someone is an “exercise of normative power” that communicates one's aim to end the relationship, and thus to “fundamentally redraw the web of relationship-based duties to which they are subject”. Stating ‘I want to break up’, according to Healey, does not fully do away with all relationship-based duties on the side of the initiating party, but it does reduce the stringency of these duties. Importantly, break-ups generate a new obligation for all members of the relationship: a *duty to transition* out of the relationship, which may require, among other things, relying less on the former partner for emotional support (Healey, 2023, p. 182). When an individual refuses to recognise their (ex-)partner's prerogative to break up, they fail to meet their duty to transition out of the relationship, and thus wrong the (ex-) partner insofar they deny their right to attach and detach themselves from commitments (Healey, 2023, p. 182). Now, recall that I am working here on the assumption of a best-case scenario where the ex-partner has provided meaningful consent to the user. What I argue here is that these are normative powers that should not be given away, since it is a relinquishing of one's relational freedom: the right to choose to continue or abandon relationships (Betzler, 2024, p. 6). By introducing the possibility of engaging with break-up chatbots to romantic life, the right to relational freedom is diminished as such, given that it opens up romantic life to scenarios where break-ups are not truly an exercise of normative power, and do not require from the receiving party of an unchosen romantic dissolution to accept and fulfil their duty to transition. This surely constitutes another noxious romantic tendency which threatens to deepen existing asymmetries in romantic life, which are caused precisely by the lack of normative powers of oppressed or marginalised people in romantic relationships.

Neither fictionalism (which would make users aware of asymmetry, according to Krueger and Osler) nor meaningful consent contribute to temper the hostile scaffolding function of break-up chatbots in virtue of their fostering of noxious romantic tendencies. Break-up chatbots have significant potential for moral corruption because as scaffolds, they introduce and reinforce mental attitudes through habituation, and often without the user being aware that this is happening (Slaby, 2016). Users may gradually develop into individuals who abuse power in relationships without even realising it, even when they adopt a fictionalist stance (my guess is that chances are only higher if they do not adopt this stance). They may also slowly devalue the

¹³ For example, A. G. Holdier and Kelly Weirich (2025) claim that engagement with AI companions promotes the objectification and subordination of women.

normative powers of (potential) partners. If the user already has a history of exercising excessive power in their relationships, then engaging with the chatbots fails to counteract, and in fact fosters, these tendencies.

Given that chatbots function better the more power they award users to serve their self-interest, and given that the supplementary body of data chatbots draw from in its responses is developed in a hetero-patriarchal society that is built on and sustained by the asymmetries I have noted here, damage to moral agency is not just potential and contingently realisable, but encoded by design. Break-up chatbots are not only hostile scaffolds then: they are oppressive technological devices, given that they foster and reinforce unjust power relations by design (Liao and Hubner, 2021; Spurrett, 2024). Even in cases where they do not directly oppress anyone—such as in the best-case scenario I am considering, in which the ex-partner gives meaningful consent—break-up chatbots corrupt moral agency by leading users to adopt attitudes and behaviours that express and perpetuate unjust social norms about intimate relationships. Again, whether users in fact adopt these attitudes and behaviours or not is beyond the point, given that break-up chatbots recreate power asymmetry in romantic life by design. In sum, break-up chatbots are hostile agential scaffolds when used to establish continuing bonds: they are morally disvaluable insofar they are specifically designed to foster the user’s tendency to abuse power in intimate relationships.

5 Narrative harm and oppressive niches

The third way in which break-up chatbots act as hostile scaffolds is by disrupting the agent’s narrative capacities—that is, by inflicting *narrative harm* (Lindemann, 2001). Without getting too much into the debate on narrative identity, let us loosely define self-narratives as the temporally extended representations of ourselves that contribute to understanding ourselves as the individuals that we are. They are both cognitive and affective in nature, since self-narratives are felt as meaningful and can be embodied (Jongepier, 2016). Self-narratives are not static but dynamic and social: we incorporate interpretations from other people and form we-narratives with them, as well as being directly influenced by socially shared narratives about norms and expectations (Gallagher & Tollefsen, 2019; Lindemann, 2001). According to authors like Heersmink (2020), technological artefacts can sustain our self-narratives, for example, by storing our memories on social media or providing new data about ourselves through health apps. However, technological devices can also be detrimental to self-narration. In this section, I give three examples of narrative harm that can result from the use of break-up chatbots, and explain how these illustrate the nature of these artefacts as oppressive technological devices.

The first narrative harm is what Osler (2025, 384) calls *narrative railroading*: “an overdetermination of the way in which the narrative is constructed by limiting the available options for shaping one’s (self-)narrative or making certain narrative options overly salient”. I have already explained that break-up chatbots are designed to preserve the narrative momentum of an extinct relationship. This is an instance of narrative railroading, since this built-in design pushes users towards the consequences I discussed in Sect. 3 (the inability to shift their orientation towards the

relationship). Besides, one of the harms that result from narrative railroading is that it pushes users towards impoverished self-narratives (Osler, 2025, 388–389). Even in the best-functioning break-up chatbots, the fact that interaction is solely driven by the user’s interests (see Sect. 4) is bound to result in an impoverished co-shaping of their self-narrative. When we are in a relationship with someone, their interpretations of our attitudes and behaviours may change us in ways we did not expect. The fact that this element of surprise is not present can drastically limit our capacity to alter our self-narratives in interaction with others, and in that sense, narrative railroading results in impoverished self-knowledge.

Narrative harm is even graver if we consider that people may afford excessive authority to the chatbot’s output about elements of their own self-narratives—such as memories and evaluations of themselves, or their relationships. This phenomenon has been called *narrative deference* by Byrne (2025). Narrative deference is not mere asymmetry, but the affordance of higher epistemic authority to an external party about one’s own experiences. Turning to break-up chatbots, it is not hard to imagine that users may give more weight to a chatbot’s output regarding facts about their past relationship than to their own. There is empirical evidence showing that interactions with chatbots can result in the formation of false memories, and that users are more bound to trust a chatbot’s ‘recollections’ than their own (or that of another human being) (Huang et al., 2024; Salvi et al., 2025). This may happen, for example, if the chatbot produces ‘proof’ that what had been said in a fight was different to what the user recalled (‘I never called you crazy, see this message I sent you?’). The user is then harmed not only because they form inaccurate self-narratives but also because they lose their self-trust (Byrne, 2025, 412–413). Since LLM-powered chatbots are known to ‘hallucinate’ falsehoods, this could mean that in practice, the user is deferring their self-narratives (and their self-trust) to a bullshitter (Hicks et al., 2024).

Another kind of narrative harm is what I call *narrative capitulation*, which refers to one’s agential authority against socially shared scripts (in contrast to deference, which is interpersonal). Lindemann (2001) explains that socially shared *master narratives* shape people’s self-understanding (and ultimately, self-respect) by guiding their attitudes and behaviours, and making certain options more salient than others. She gives the example of the pregnancy narrative, according to which people are expected to (and often will) take certain steps which are supposed to apply to the pregnancy process universally—like buying baby-themed books, refraining from going to parties, taking time off work, feeling elated and excited, or thinking about baby names. Some master narratives are grounded in (and themselves constitute and reinforce) oppressive norms. Some examples are the narratives of behaving like ‘a real man’ or ‘a model migrant’, which are built on sexist and racist ideologies. Lindemann (1995, 34) says that despite their power, it is possible to counteract oppressive master narratives by creating and engaging with *counterstories*, which are “told for the specific purpose of resisting and undermining a dominant story”. My claim is that users of break-up chatbots are doing exactly the opposite to creating counterstories, and thus capitulating to existing oppressive narratives about romantic life. Users who build continuing bonds through break-up chatbots indirectly support the oppressive master narrative of ‘amatormativity’, according to which a fulfilling life can only be had in a romantic relationship (Brake, 2012, 88–9); as well as related ones such

as singlehood stigma (DePaulo & Morris, 2005). After a break-up, these master narratives tend to be questioned in social interactions with friends—even therapeutic uses of chatbots may be useful in this sense, depending on the corpus of data that the device draws from. When break-up chatbots enable the user’s refusal to give up the extinct romantic relationship, they constitute either an indirect contribution or an active endorsement of amatonormativity on the side of the user. This is narrative capitulation, which in turn is a form of self-inflicted harm.¹⁴

It may seem that I have returned to contingent risks; to what *could* go wrong for users of these technologies. Recall that my aim was to show what *does* go wrong in the scenario where people use break-up chatbots to build continuing bonds with their ex-partner. The aim of the discussion above is not (solely) to highlight what kinds of narrative harms may be inflicted to users of break-up chatbots, but to illustrate that the use of break-up chatbots supports the construction of an *oppressive niche* aimed at disrupting narrative capacities (one’s own and those of others). A niche can be defined as “a space that is altered, even designed and constructed, in order to support the needs, cognition, affectivity, and agency of the people that inhabit it” (Osler et al., 2024). Users of break-up chatbots engage in narrative niche construction (Heersmink, 2020) when they articulate their self-narratives through interaction with these technologies. Like Fabry (2025) argues, narrative niches are also constructed on and supported by master narratives. Fabry (2025, 397, fn. 3) sees her account as a response to Heersmink’s, but I see their views as compatible: the way material objects contribute to self-narration is not isolated from social norms, and social norms inform how we use certain objects in self-narration. However, as I have shown in the previous sections, when engaging with break-up chatbots, people are not simply building a niche for themselves. They instead enter a niche that is built by a third party with encoded oppressive design. First, this oppresses the user themselves given that niche is designed to disrupt narrative capacities. Entering this environment is a form of self-wronging, given that one is putting intense pressure on one’s narrative capacities, and by extension, one’s moral agency (Schechtman, 1996). Given that break-up chatbots scaffold noxious romantic tendencies, their harms extend beyond the corruption of the user’s own moral agency. They not only disrupt individual narrative capacities but also contribute to the construction of social niches in which those asymmetrical attitudes and behaviours can affect others. In other words, users actively participate in shaping and maintaining the master narrative that portrays singlehood or break-ups as a personal failure (Kukla, 2022). By fostering noxious romantic tendencies, users risk carrying these dispositions into their offline relationships and interactions, thereby sustaining and reinforcing oppressive norms in particular social contexts that disrupt other people’s narrative capacities and autonomy. In this way, break-up chatbots function not only as sites of individual narrative harm but also as vectors through

¹⁴ Whether narrative capitulation is blameworthy would depend, for example, on whether there is a duty to narrative resistance. Lindemann (2020) does endorse the existence of that duty for bystanders (i.e., people who are presented with counterstories about oppressed others). LaGuardia-LoBianco (2018) has argued that victims of oppression have a duty of self-care, which I take to potentially include engaging with counterstories. Even then, we may find that narrative capitulation is blameless, since living under oppression often leaves people with no choice but to give in to these master narratives.

which oppressive (narrative and interpersonal) niches are co-created by users in virtue of their engagement with chatbots to build continuing bonds.

That said, one might object that this argument overgeneralizes. Just as knives contribute to oppressive niches when used as weapons but also enable beneficial niche construction when used for cooking, break-up chatbots may not invariably support oppressive niches. This objection raises an important distinction between technological affordances (what technology can lead people to do) and technological design (what technology is aimed at doing). My claim is not that every use of break-up chatbots must contribute to oppressive niches, but that their design function of sustaining continuing bonds with ex-partners renders them especially likely to do so, since the conditions for oppressive niche construction (asymmetry, lack of recognition of normative powers, narrative railroading, narrative deference, and narrative capitulation) are encoded in their design. Insofar as this function encodes amatonormative assumptions and disruption of narrative agency, break-up chatbots are not neutral in the way knives are. Rather, their design systematically orients users toward oppressive niches and makes them co-participants in the creation and maintenance of these niches, even if in principle some users might escape that orientation.

6 Conclusion

Let us take stock. Break-up chatbots can act as hostile scaffolds by placing and maintaining the user in a harmful imaginary relationship, fostering noxious romantic tendencies, and embedding the user in an oppressive niche that they co-create. These harms may be, to an extent, contingent, but they exemplify the nature of break-up chatbots as oppressive technological artefacts that are built on, enact, and maintain unjust practices and norms within romantic life. Specifically, they are aimed at maintaining the user engaged in a meaningless activity, promoting abuses of power by prioritising the user's self-interest and depleting normative powers, and supporting of oppressive master narratives. Break-up chatbots are, by design, hostile scaffolds.

It shall be reiterated that the focus has been on a specific use of break-up chatbots: when they are used to attempt to build continuing bonds with an ex-partner. In principle, my argument has no bearing on other use contexts, like the two scenarios I contemplated in Sect. 2 (attempting to understand the reasons for the break-up and trying to move on). However, I have worked from a best-case scenario, and left out two aggravating factors. First, if the person represented through the chatbot has not given (meaningful) consent, they will be directly harmed by some of the mechanisms I have described. For example, it will be *their* normative power that is depleted if the other party does not recognise their decision to end the relationship. Second, the fact that break-up chatbots are (and will foreseeably continue to be) for-profit products may deepen the extent of the harms I have identified. There are growing concerns about the gamification of dating apps, which are designed to keep users hooked for profit rather than fulfilling the user's goal of finding a partner. This has noxious consequences on users' self-image, and potentially on romantic life in general (McKeever, 2022; Nader, 2025). Break-up chatbots, which may plausibly be marketed as emotional aids during romantic dissolutions, will potentially also be designed to

maintain interaction—and thus profit—over and above the user’s well-being. These two aggravating factors highlight the urgency to establish ethical guidelines before these technologies are more widely available, and shed doubt on whether potential benefits of break-up chatbots will not be outweighed by noxious effects in pretty much all use contexts.

Besides these practical considerations, the discussion of break-up chatbots is revealing of conceptual lacunas in the philosophy of grief with respect to non-bereavement and continuing bonds. Until now, philosophers have been generally optimistic about the role of continuing bonds in good grief. Nevertheless, I have shown that continuing bonds can also be harmful, and even reinforce injustice, in the case of romantic dissolutions. It remains to be seen if some of the attitudes and behaviours I have described here are problematic without technological mediation (or through other means of technological mediation). In that sense, my argument is useful not only to understand the ethical implications of these technologies, but also to acquire a wider understanding of the neglected phenomenon of break-up grief.

Acknowledgments I am grateful to two anonymous reviewers, whose comments greatly helped shaping the final version of the manuscript.

Disclosure statement The author reports no conflict of interest.

Funding This research and author were supported by EU Horizon Europe under the Marie Skłodowska-Curie COFUND grant No 101081327 YUFE4Postdocs.

References

- Archer, A. (2021). Fans, crimes and misdemeanors: Fandom and the ethics of love. *The Journal of Ethics*, 25, 543–566.
- Bagley, B. (2015). Loving someone in particular. *Ethics*, 125(2), 477–507.
- Ben Ze’ev, A., & Goussinsky, R. (2008). *In the name of love: Romantic ideology and its victims*. Oxford University Press.
- Betzler, M. (2024). Ending a special relationship: Toward an ethics of divorce. *Philosophical Quarterly*.
- Bortolan, A. (2021). Narrate it until you become it. *Journal of the American Philosophical Association*, 7(4), 474–493.
- Bourassa, K. J., Manvelian, A., Boals, A., Mehl, M. R., & Sbarra, D. A. (2017). Tell me a story: The creation of narrative as a mechanism of psychological recovery following marital separation. *Journal of Social and Clinical Psychology*, 36, 359–379. <https://doi.org/10.1521/jscp.2017.36.5.359>
- Brake, E. (2012). *Minimizing marriage*. Oxford University Press.
- Brogaard, B. (2015). *On romantic love: Simple truths about a complex emotion*. Oxford University Press.
- 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>
- Byrne, E. A. (2025). Narrative deference. *Topoi*, 44, 405–417. <https://doi.org/10.1007/s11245-024-10105-z>
- Campbell, S. M., Liu, P., & Nyholm, S. (2025). Can chatbots preserve our relationships with the dead? *Journal of the American Philosophical Association*, 11(2), 230–248. <https://doi.org/10.1017/apa.2025.1>
- Cholbi, M. (2022). *Grief: A philosophical guide*. Princeton University Press.
- Cholbi, M. (2025). Memory and mimesis in our relationships with posthumous avatars. In H. Shevlin (Ed.), *AI in society: Relationships (Oxford Intersections)*. Oxford University Press.
- Colombetti, G., & Krueger, J. (2015). Scaffoldings of the affective mind. *Philosophical Psychology*, 28(8), 1157–1176.

- Danaher, J., & Nyholm, S. (2025). The ethics of personalised digital duplicates: A minimally viable permissibility principle. *AI Ethics*, 5, 1703–1718. <https://doi.org/10.1007/s43681-024-00513-7>
- DePaulo, B. M., & Morris, W. L. (2005). Singles in society and in science. *Psychological Inquiry*, 16, 2–3. <https://doi.org/10.1080/1047840X.2005.9682918>
- Fabry, R. E. (2025). Self-narration in the oppressive niche. *Topoi*, 44, 393–404. <https://doi.org/10.1007/s11245-024-10088-x>
- Fabry, R. E., & Alfano, M. (2024). The affective scaffolding of grief in the digital age: The case of deathbots. *Topoi*, 43, 757–769. <https://doi.org/10.1007/s11245-023-09995-2>
- Fletcher, G. (2021). *Dear prudence: The nature and normativity of prudential discourse*. Oxford University Press.
- Frankfurt, H. G. (2004). *The reasons of love*. Princeton University Press.
- Friedman, M. (2003). *Autonomy, gender, politics*. Oxford University Press.
- Gallagher, S., & Tollefsen, D. (2019). Advancing the ‘We’ through narrative. *Topoi*, 38, 211–219. <https://doi.org/10.1007/s11245-017-9452-1>
- González Barman, K., Lohse, S., & de Regt, H. W. (2025). Reinforcement learning from human feedback in llms: Whose culture, whose values, whose perspectives? *Philosophy and Technology*, 38, Article 35. <https://doi.org/10.1007/s13347-025-00861-0>
- Harbin, A. (2016). *Disorientation and moral life*. Oxford University Press.
- Healey, R. (2023). Breaking up and the value of commitment. *Ergo an Open Access Journal of Philosophy*, 10(6), 158–187. <https://doi.org/10.3998/ergo.4636>
- Heersmink, R. (2020). Narrative niche construction: Memory ecologies and distributed narrative identities. *Biology & Philosophy*, 35, 53. <https://doi.org/10.1007/s10539-020-09770>
- Hicks, M. T., Humphries, J., & Slater, J. (2024). ChatGPT is bullshit. *Ethics and Information Technology*, 26, Article 38. <https://doi.org/10.1007/s10676-024-09775-5>
- Higgins, K. (2013). Love and death. In J. Deigh (Ed.), *On emotions: Philosophical essays* (pp. 159–178). Oxford University Press.
- Holdier, A. G., & Weirich, K. (2025). AI Romance and Misogyny: A Speech Act Analysis. In P. Hacker (Ed.), *Oxford Intersections: AI in Society*. <https://doi.org/10.1093/9780198945215.003.0074>
- Huang, T. R., Cheng, Y. L., & Rajaram, S. (2024). Unavoidable social contagion of false memory from robots to humans. *American Psychologist*, 79(2), 285–298. <https://doi.org/10.1037/amp0001230>
- Illouz, E. (2012). *Why love hurts: A sociological explanation*. Polity.
- Jenkins, C. S. I. (2022). *Sad love: Romance and the search for meaning*. Polity.
- Joel, S., & MacDonald, G. (2021). We’re not that choosy: Emerging evidence of a progression bias in romantic relationships. *Personality and Social Psychology Review*, 25(4), 317–343.
- Jollimore, T. (2015). This endless space between the words: The limits of love in Spike Jonze’s *Her*. *Midwest Studies in Philosophy*, 39, 120–143. <https://doi.org/10.1111/misp.12039>
- Jones, K. (2008). How to change the past. In K. Atkins, & C. Mackenzie (Eds.), *Practical identity and narrative agency*. Routledge.
- Jongepier, F. (2016). Towards a constitutive account of implicit narrativity. *Phenomenology and the Cognitive Sciences*, 15(1), 51–66.
- Klass, D. (2006). Continuing conversation about continuing bonds. *Death Studies*, 30(9), 843–858.
- Klass, D., & Steffen, E. M. (Eds.). (2018). *Continuing bonds in bereavement: New directions for research and practice*. Routledge.
- Kross, E., Berman, M. G., Mischel, W., Smith, E. E., & Wager, T. D. (2011). Social rejection shares somatosensory representations with physical pain. *Proceedings of the National Academy of Sciences of the United States of America*, 108(15), 6270–6275. <https://doi.org/10.1073/pnas.1102693108>
- 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>
- Kukla, Q. R. (2022). Good Endings. *The Philosophers’ Magazine*. 11 April (online edition). Retrieved on 14 Feb 2024 from <https://philosophersmag.com/?s=good+endings>
- LaGuardia-LoBianco, A. W. (2018). Complicit suffering and the duty to self-care. *Philosophy*, 93(2), 251–277.
- Lee, L. A., Sbarra, D. A., Mason, A. E., & Law, R. W. (2011). Attachment anxiety, verbal immediacy, and blood pressure: Results from a laboratory-analogue study following marital separation. *Personal Relationships*, 18, 285–301.
- Liao, S., & Huebner, B. (2021). Oppressive Things. *Philosophy and Phenomenological Research*, 103, 92–113. <https://doi.org/10.1111/phpr.12701>

- Lindemann, H. (1995). Resistance and insubordination. *Hypatia*, 10(2), 23–40. <https://doi.org/10.1111/j.1527-2001.1995.tb01367.x>. published as Hilde Lindemann Nelson.
- Lindemann, H. (2001). *Damaged Identities, narrative repair*. Cornell University Press (published as Hilde Lindemann Nelson).
- Lindemann, H. (2020). Counter the counterstory: Narrative approaches to narratives. *Journal of Ethics and Social Philosophy*, 17(3).
- Lindemann, N. F. (2022). The ethics of ‘Deathbots.’ *Science and Engineering Ethics*, 28, Article 60. <https://doi.org/10.1007/s11948-022-00417-x>
- Lopez-Cantero, P. (2018). The break-up check: Exploring romantic love through relationship terminations. *Philosophia*, 46, 689–703.
- Lopez-Cantero, P. (2023). Three stages of love, narrative, and self-understanding. In A. Montes-Sánchez, & A. Salice (Eds.), *Emotional Self-Knowledge*. Routledge.
- McConnell, D. (2016). Narrative self-constitution and recovery from addiction. *American Philosophical Quarterly*, 53(3), 307–322.
- McKeever, N. (2019). What can we learn about romantic love from Harry Frankfurt’s account of love? *Journal of Ethics and Social Philosophy*, 14(3), 204–226. <https://doi.org/10.26556/jesp.v14i3.553>
- McKeever, N. (2022). Online dating and love robots: How technology May undermine valuable features of romantic love. In A. Grahle, N. McKeever, & J. Saunders (Eds.), *Philosophy of love in the Past, Present, and future*. Routledge.
- Millar, B., & Lopez-Cantero, P. (2022). Grief, continuing Bonds, and unreciprocated love. *South J Philos*, 60, 413–436. <https://doi.org/10.1111/sjp.12462>
- Milligan, T. (2011). *Love*. Acumen Publishing.
- Nader, K. (2025). The gamification of dating online. *Theoria*, 91(3), Article e12549. <https://doi.org/10.1111/theo.12549>
- Norlock, K. J. (2017). Real (and) imaginal relationships with the dead. *The Journal of Value Inquiry*, 51(2), 341–356.
- Nyholm, S., & Frank, L. (2017). From sex robots to love robots: Is mutual love with a robot possible? In J. Danaher, & N. McArthur (Eds.), *Robot sex: Social and ethical implications* (pp. 219–244). MIT Press.
- Osler, L. (2025). Narrative railroading. *Topoi*, 44, 379–392. <https://doi.org/10.1007/s11245-024-10146-4>
- Osler, L., Engelen, B., & Archer, A. (2024). The ethics and politics of nudges and niches. In T. S. Petersen, S. J. Holmenand, & J. Ryberg (Eds.), *Crime prevention by exclusion: Ethical considerations*. Routledge.
- Shpall, S. (2018). A tripartite theory of love. *Journal of Ethics and Social Philosophy*, 13(2), 91–124
- Ratcliffe, M. (2023). *Grief worlds: A study of emotional experience*. MIT Press.
- Ratcliffe, M., & Richardson, L. (2023). Grief over non-death losses: A phenomenological perspective. *Passion: Journal of the European Philosophical Society for the Study of Emotions*, 1(1), 50–67. <https://doi.org/10.59123/passion.v1i1.12287>
- Ratcliffe, M., Richardson, L., & Millar, B. (2023). On the appropriateness of grief to its object. *Journal of the American Philosophical Association*, 9(2), 318–334. <https://doi.org/10.1017/apa.2021.55>
- Rhoades, G. K., Kamp Dush, C. M., Atkins, D. C., Stanley, S. M., & Markman, H. J. (2011). Breaking up is hard to do: The impact of unmarried relationship dissolution on mental health and life satisfaction. *Journal of Family Psychology*, 25(3), 366–374. <https://doi.org/10.1037/a0023627>
- Salvi, F., Horta Ribeiro, M., Gallotti, R., et al. (2025). On the conversational persuasiveness of GPT-4. *Nature Human Behaviour*. <https://doi.org/10.1038/s41562-025-02194-6>
- Schechtman, M. (1996). *The constitution of selves*. Cornell University Press.
- Slaby, J. (2016). Mind invasion: Situated affectivity and the corporate life hack. *Frontiers in Psychology*, 7, Article 266.
- Spurrett, D. (2024). On hostile and oppressive affective technologies. *Topoi*, 43, 821–832.
- Stokes, P. (2021). *Digital souls: A philosophy of online death*. Bloomsbury.
- Timms, R., & Spurrett, D. (2023). Hostile scaffolding. *Philosophical Papers*, 52(1), 1–30.
- Valentine, C. (2008). *Bereavement narratives*. Routledge.
- Wajnerman Paz, A. (2021). Is mental privacy a component of personal identity? *Frontiers in Human Neuroscience*. <https://doi.org/10.3389/fnhum.2021.773441>
- Williams, J. (2024). Losing faith and losing a world: Deconversion as an occasion for grief. *Phenomenology and the Cognitive Sciences*. <https://doi.org/10.1007/s11097-024-10037-4>
- Wolf, S. R. (2010). *Meaning in life and why it matters*. Princeton University Press.

Woo, C. W., Koban, L., Kross, E., Lindquist, M. A., Banich, M. T., Ruzic, L., Andrews-Hanna, J. R., & Wager, T. D. (2014). Separate neural representations for physical pain and social rejection. *Nature Communications*, 5, Article 5380. <https://doi.org/10.1038/ncomms6380>

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

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.