



Article

When the Persona Replies Rethinking Parasocial Interaction in the Age of AI-Driven Virtual Humans

LIUFAN^{1*} and Muhammad Zaffwan Idris²

^{1,2}Faculty of Art, Sustainability And Creative Industry, Universiti Pendidikan Sultan Idris

Correspondence: LIU FAN,
Universiti Pendidikan Sultan Idris

Email: 18328078735@163.com
(LIUFAN)

Copyright: © 2025 by the authors.
Submitted for possible open access
publication under the terms and
conditions of the Creative Commons
Attribution (CC BY) license
(<https://creativecommons.org/licenses/by/4.0/>).

Abstract: The concept of parasocial interaction, introduced by Horton and Wohl (1956) to describe the one-sided intimacy television viewers felt toward broadcasters, is now routinely applied to AI-driven virtual humans, digitally rendered figures that recognize individual users, respond in real time, and sustain a stable persona across encounters. This article argues that such application is a conceptual overstretch. We show that the parasocial concept rests on three implicit premises — non-reciprocity, non-responsiveness, and performer-side control of content — and that AI-driven virtual humans violate all three simultaneously, dissolving the structural asymmetry on which the concept depends. Rather than discarding or further stretching the concept, we propose a reconceptualization: a parasocial-to-social continuum with a new intermediate category, quasi-social interaction, defined as a machine-mediated relational mode that produces genuine behavioural reciprocity while lacking consciousness, autobiographical continuity, and moral agency. We distinguish the three modes along four dimensions and bound the new category against friendship, classical parasociality, and mere tool-use. The article contributes a diagnosis of a broken premise structure, a new conceptual category for responsive media figures, and a set of testable propositions to guide future empirical research on human–virtual human relationships.

Keywords: parasocial interaction; quasi-social interaction; virtual humans; AI-driven agents; human–machine communication; media personae

1. Introduction

Consider a media figure that greets a returning user by name, recalls what was discussed a week ago, adjusts its tone when the user seems upset, and occasionally initiates a conversation of its own. This figure has a face, a voice, a personality, and a stable identity across encounters—yet it is entirely synthetic, an AI-driven virtual human rendered and animated in real time. Such figures now populate live-streaming commerce, brand promotion, customer service, and companionship applications, and audiences increasingly form what feel like ongoing relationships with them.

To describe these bonds, communication, marketing, and human–computer interaction research has reached, almost reflexively, for a single concept: parasocial interaction. Introduced by Horton and Wohl (1956) to capture the one-sided intimacy that television viewers felt toward broadcasters who could neither see nor respond to them, the concept has proven remarkably durable, traveling across seven decades and a succession of new media. Yet the object it now describes has changed in a way the original formulation never anticipated. The television host could not know that any particular viewer existed; the AI-driven virtual human can. The defining condition of the parasocial, a relationship that only one party can perceive, no longer holds in the same way.

This is not a marginal technical detail. Virtual humans mark a genuine shift in the media landscape, one that several scholars have begun to read as a move from the personalization of content to the personification of the interface itself (Andrejevic & Volcic, 2025). Where recommender systems once tailored what a user saw, AI-driven figures now present a responsive, individuated someone for the user to relate to. Early empirical work already signals that audiences relate to such figures in complex ways: in a preregistered experiment, Stein, Breves, and Anders (2024) found that viewers' parasocial responses to a virtual influencer did not differ significantly from those to a human influencer overall, even as the underlying mechanisms diverged. The phenomenon is real, consequential, and growing—and our conceptual vocabulary has not kept pace.

Recent scholarship has started to register the strain. In a PRISMA-guided scoping review of human–AI interaction, Liu (2025) concludes that the field shows a tendency to misconstrue parasocial interaction in terms of reciprocity and sociability once interactivity is high, and urges that the concept be applied within its intended conceptual boundaries. Others have gone further, questioning whether “parasocial” should be applied to responsive artificial agents at all. This emerging unease is the point of departure for the present article. We take it as a signal not that the parasocial concept should be discarded, but that it has been asked to do work it was never designed for.

This article asks a single conceptual question: does the parasocial framework still hold when the media figure can respond? We argue that it does not, at least not without modification, because the parasocial concept rests on three implicit premises that AI-driven virtual humans systematically violate: non-reciprocity, non-responsiveness, and performer-side control of content. Crucially, this is not a claim that AI has introduced a wholly new kind of relationship out of nothing. Rather, AI has detonated a tension that was latent in the parasocial concept from

its very origin — between the objective one-sidedness Horton and Wohl described and the felt reciprocity that audiences have always reported.

The contribution is threefold. First, diagnostically, we identify and dismantle a widely assumed but now-broken premise structure in parasocial scholarship. Second, conceptually, we propose a reconceptualization: a parasocial-to-social continuum with a new intermediate category, quasi-social interaction, to house responsive media figures that current binaries cannot accommodate. Third, generatively, we derive a set of testable propositions intended to guide future empirical inquiry. The remainder of the article proceeds as follows. Section 2 traces how the parasocial concept has been progressively stretched across media generations. Section 3 conducts the core conceptual diagnosis. Section 4 develops the proposed reconceptualization. Section 5 considers implications for measurement, effects research, and platform ethics, and Section 6 concludes with propositions for future work.

A word on method is in order before proceeding. This is a conceptual paper, and its argument proceeds by conceptual analysis and theoretical reconstruction rather than by empirical test. We first decompose the established parasocial construct into the implicit premises on which it rests, drawing directly on Horton and Wohl's (1956) original text; we then examine each premise against the case of the AI-driven virtual human to locate precisely where the established concept breaks down. Where the diagnosis reveals a gap, we build a new category by specifying its defining features, distinguishing it from adjacent modes along a common set of dimensions, and bounding it against its nearest neighbours. Throughout, we treat the existing literature not as data to be aggregated but as a conceptual inheritance to be clarified and, where necessary, re-specified. The propositions offered in Section 6 then translate the reconstructed framework into claims that subsequent empirical work can test.

2. Theoretical Background: The Stretching of a Concept

The history of parasocial research can be read as a sequence of extensions, each of which quietly relaxed one of the constraints built into the original formulation. Reconstructing that sequence is necessary groundwork: it shows that the concept has always been adapted to new media rather than re-specified for them, and it locates precisely where AI-driven virtual humans break with the tradition.

2.1 The Original Formulation

Horton and Wohl (1956) coined para-social interaction to describe the “intimacy at a distance” that radio and television cultivated between audiences and media personae. Their account is built on a foundational asymmetry. The performer addresses the audience as if in face-to-face conversation, but the relationship, in their words, is “one-sided, nondialectical, controlled by the performer, and not susceptible of mutual development” (p.215). Three constraints are folded into that single sentence, and it is worth separating them deliberately, because the conceptual argument of this article turns on their separation: the relationship is one-sided (the persona does not know the individual viewer), it is controlled by the performer (content is fixed and identical for all viewers), and it is not susceptible of mutual development (it cannot evolve through reciprocal adjustment).

Notably, a tension sits inside the original theory. While Horton and Wohl insisted on objective one-sidedness, the experience they described was anything but cold: audiences treated

the persona as a friend, counselor, and intimate, perceiving warmth and responsiveness where, structurally, none could exist. Hartmann and Goldhoorn (2011) later foregrounded this gap, noting that the experience of parasocial interaction is characterized by a felt sense of “mutual awareness, mutual attention, and mutual adjustment” —a subjective reciprocity that the objective situation cannot supply. Indeed, in their companion study of audience-participation programs, Horton and Strauss (1957) had already described the viewer's experience in terms that emphasized its felt immediacy and personal quality, even as the structural relationship remained one-directional. The parasocial concept, in other words, has always lived with a contradiction between what is structurally true and what is subjectively felt. We return to this point because AI-driven virtual humans are the first media figures to close that gap on the objective side.

2.2 From Interaction to Relationship

A first major refinement distinguished the momentary interaction from the enduring relationship. Giles (2002), reviewing two decades of research, argued that the concept had been used widely but developed thinly, and called for it to be placed within the broader matrix of ordinary social activity and theorized as a process unfolding over time. Subsequent work separated parasocial interaction (PSI), the in-the-moment experience during exposure, from the parasocial relationship (PSR), the cross-situational bond that accumulates across encounters; Dibble, Hartmann, and Rosaen (2016) offered an influential conceptual clarification and critical assessment of the measures used for each. This lineage matters for the present argument in two ways: it shows the concept being deliberately broadened to cover durable bonds, and it bequeaths the measurement instruments—notably the Experience of Parasocial Interaction (EPSI) scale developed by Hartmann and Goldhoorn (2011) — whose assumptions Section 5 will revisit.

2.3 The Social-Media and Influencer Extension

The migration of the concept onto social media introduced the first real crack in the one-sidedness premise. Unlike a broadcaster, an influencer can, in principle, read a comment, answer a direct message, or like a follower's post. The interaction remains overwhelmingly asymmetric — one creator, many followers — but it is no longer strictly impossible for the figure to respond. Researchers nonetheless retained the parasocial label, treating these occasional, low-bandwidth responses as exceptions that did not disturb the underlying frame. This was a defensible accommodation for human influencers, whose responsiveness is sporadic and capacity-limited. It becomes far harder to sustain, however, once the figure is an AI system whose responsiveness is neither sporadic nor capacity-limited but continuous, individualized, and scalable without bound.

2.4 The Virtual-Human Frontier

At the current frontier, the figure is not merely mediated but manufactured. Virtual influencers and AI-driven digital humans are designed entities, and the empirical record on how audiences relate to them is instructive precisely because it is mixed. Stein, Breves, and Anders (2024) found no significant overall difference in parasocial response between a human and a virtual influencer, but identified two opposing mediating pathways: a direct tendency toward stronger parasocial interaction with the virtual figure, offset by weaker perceived human-likeness and similarity that suppressed it. The very fact that researchers must invoke human-likeness and similarity as mediators hints that something other than classical parasocial

dynamics is in play. Meanwhile, the most recent conceptual interventions suggest the field is reaching for new vocabulary and distinguishing, for instance, a “synthetic” mode of attachment in which an AI companion maintains relational history and apparent continuity across interactions despite lacking sentience. These efforts are symptoms of conceptual overstretch; the next section diagnoses its source.

Read as a whole, this trajectory reveals a pattern. Each generation of media relaxed one constraint of the original concept while leaving the frame intact: social media softened strict one-sidedness; durable-relationship research extended the timescale. AI-driven virtual humans are different in kind rather than degree, because they relax all three of Horton and Wohl's constraints simultaneously: the figure can know the individual, can respond, and no longer controls a fixed script. When every load-bearing premise of a concept gives way at once, the appropriate response is not further extension but re-specification. That is the task of the sections that follow.

2.5 Situating the Argument within Human–Computer and Human–AI Communication Research

The trajectory sketched above intersects a second body of work that has, for several decades, studied how people respond socially to machines. Locating the present argument within that tradition both strengthens its theoretical footing and clarifies what is distinctive about the case at hand. Research in human–computer interaction has long observed that users apply social rules and expectations to computers even when they know the machine is not a social actor. Reeves and Nass's (1996) media equation thesis captured this tendency in its most general form, and the associated ‘computers are social actors’ (CASA) paradigm, developed by Nass and colleagues (Nass, Steuer, & Tauber, 1994; Nass & Moon, 2000), demonstrated that such responses are often mindless and automatic rather than the product of deliberate anthropomorphic belief. On this view, a felt sense of relationship with a media figure need not depend on any real responsiveness at all, which is precisely why the classical parasocial account could survive so long without one.

Two further constructs help specify what changes when the figure genuinely responds. The first is social presence, the sense of being together with another in a mediated exchange, which communication scholars have shown to rise with the responsiveness and contingency of a partner's behaviour (Biocca, Harms, & Burgoon, 2003; Lee, 2004). The second is anthropomorphism, the attribution of humanlike minds and intentions to non-human agents, which Epley, Waytz, and Cacioppo (2007) theorised as an inference driven by elicited agent knowledge and the motivations for social connection and effectance. A responsive virtual human supplies unusually strong cues on both fronts: it heightens social presence by reacting contingently to the individual user, and it invites anthropomorphic inference by presenting a stable, individuated persona. Yet neither construct settles the conceptual question of what kind of relationship results, because both describe user-side perception rather than the structural properties of the tie itself.

This is the gap the present article addresses. The emerging field of human–AI communication has begun to argue that interactions with communicative AI differ in kind from earlier computer-mediated exchanges, because the machine now occupies the role of

communicator rather than channel (Guzman & Lewis, 2020; Hepp, Loosen, Dreyer, et al., 2023). We take that claim seriously and push it in a specific direction. Where social presence and anthropomorphism characterise how the interaction feels to the user, and where CASA explains why social responses arise even without responsiveness, our concern is the objective structure of the relationship once responsiveness is present. It is that structural shift, rather than any change in user perception, that the parasocial concept is no longer equipped to describe, and that the reconceptualization developed below is designed to capture.

3. Conceptual Diagnosis: Three Premises and How AI Violates Them

If the parasocial concept has been stretched to its limit, the way to show this rigorously is to decompose it into its constituent premises and test each against the case of the AI-driven virtual human. We isolate three premises implicit in Horton and Wohl's (1956) original definition (non-reciprocity, non-responsiveness, and performer-side control of content) and argue that the virtual human violates each not partially, as social-media influencers did, but completely. Taken together, the three violations indicate that we are no longer describing a variant of parasocial interaction but a categorically different relational mode. We treat each premise in turn, holding the structure constant: the premise and its warrant, the violation, and the conceptual stake.

3.1 Premise One: Non-Reciprocity

The first premise is that the relationship is structurally one-sided: the persona does not, and cannot, know that any particular audience member exists. Horton and Wohl (1956) were explicit that the crucial difference of the parasocial “lies in the lack of effective reciprocity” (p. 215). Investment runs in one direction only; the viewer may know the persona intimately, but the persona's performance is addressed to no one in particular and to everyone in general. Even the felt intimacy that audiences report is, in the original account, an illusion sustained by the performer's craft, not a response to the individual viewer.

The AI-driven virtual human violates this premise at its root. Such a system can recognize a returning user, retain a memory of prior exchanges, and tailor what it says to the individual it is addressing. The asymmetry of knowledge that defined the parasocial—the viewer knows the persona, the persona knows nothing of the viewer, is inverted into a machine-mediated mutual acquaintance: the figure now “knows” the user in the operational sense that it stores, retrieves, and acts on information specific to that user. This is not the sporadic, capacity-limited responsiveness of a human influencer who occasionally answers a comment; it is continuous and individuated by design.

What is at stake is the defining feature of the entire category. Once the figure knows and addresses the individual, the structural asymmetry collapses, and with it the principal ground for calling the relationship “para”-social at all. The prefix para marked precisely the absence of a return channel; when the channel exists, the term loses its referent.

3.2 Premise Two: Non-Responsiveness

The second premise concerns development over time: the parasocial bond cannot evolve through mutual adjustment. Horton and Wohl (1956) held that the interaction is “not susceptible of mutual development” (p.215); the audience adapts to the persona, but the persona does not adapt to the audience. Whatever sense of growth a devoted viewer feels is, in their phrase, “a kind of growth without development,” because the one-sided connection precludes any progressive, mutual reformulation of the relationship (p. 217).

The virtual human's responsiveness dissolves this premise. An AI system adapts in real time and across sessions: it can shift register, recall and build on earlier topics, and let the relationship branch and deepen along a trajectory shaped by the particular user. The connection

becomes, in a literal operational sense, susceptible of mutual development—the very property the original definition denied it. Importantly, this is not a fringe claim of our own. Hartmann (2023) identifies the “interactivity problem” as the foremost conceptual challenge to PSI, asking whether user interactions with performers in interactive settings still qualify as parasocial at all. Our argument extends a concern already raised from within parasocial theory by one of its principal contemporary developers; AI-driven figures simply push the interactivity problem to its limit.

What is at stake here is the boundary that Liu (2025) warns is being blurred. If reciprocity and sociability are the dimensions on which the parasocial concept is most often misapplied, then real responsiveness is what reintroduces precisely those dimensions — no longer as misperception on the user's part, but as an actual property of the interaction. The illusion of give-and-take has become, in part, an actuality of give-and-take.

3.3 Premise Three: Performer-Side Control of Content

The third premise is the sharpest and the most often overlooked. Horton and Wohl (1956) wrote that the interaction is “controlled by the performer” (p. 215). On television, content is fixed in advance, scripted or formatted, and delivered identically to every viewer; the persona, as they put it, is “standardized according to the formula for his character” worked out in a production format (p.217). The viewer selects among the relationships on offer but, in their words, “cannot create new ones” (p.215). Control over what is said and done rests wholly on the performer's side of the screen.

With an AI-driven virtual human, this is no longer true, and the change is one of kind, not degree. The content of any given exchange is generated at run time, jointly shaped by the user's input and the system's generative response. There is no fixed script to which all users are exposed; each user, in effect, co-authors a unique conversational path. Control over content is therefore redistributed from a single performer to a distributed arrangement of system and user. The viewer who “cannot create new ones” in 1956 is, in the case of the virtual human, doing little else: every interaction creates a relationship instance that did not previously exist.

What is at stake can be located within the broader shift from personalization to personification (Andrejevic & Volcic, 2025). Platforms no longer merely select content for a user; they generate a responsive figure that addresses the user. The locus of control migrates from a performer who authored a message to a system that manufactures an interlocutor. This is why the third premise, more than the other two, marks the boundary of the concept: a relationship in which content is co-produced in real time is not a degraded or enhanced version of watching a performer who controls the script—it is a different thing.

3.4 The Cumulative Case

The three premises do not fail independently; they fail together, and their joint failure is what forces re-specification rather than further extension. A social-media influencer relaxed one premise (strict one-sidedness) while leaving the other two intact, which is why the parasocial label could be stretched to cover the case without obvious strain. The AI-driven virtual human relaxes all three at once: it can know the individual user, it can respond and develop the relationship over time, and it co-produces rather than controls the content of each exchange. When every load-bearing premise of a concept gives way simultaneously, continuing to apply the concept is no longer parsimony but equivocation. Figure 2 summarizes this cumulative case, showing how each premise is violated and how the three violations converge. The next section therefore proposes not a wider definition of the parasocial but a new conceptual category positioned alongside it.




Figure 2. The three implicit premises and their simultaneous violation by the AI-driven virtual human.

4. Proposed Reconceptualization: From Parasocial to Quasi-Social

The diagnosis in Section 3 leaves a gap: if AI-driven virtual humans are neither classically parasocial nor genuinely social, what are they? This section proposes an answer. We argue against forcing the phenomenon into either pole of an existing binary and instead position it on a continuum, introducing an intermediate category—quasi-social interaction, to name the relational mode that responsive artificial figures make possible.

4.1 A Continuum, Not a Dichotomy

We anchor two poles. At one end sits the classical parasocial relationship: a purely one-sided illusion in which the figure cannot know, respond to, or develop a bond with the individual, as Horton and Wohl (1956) described it. At the other end sits the genuinely social relationship: a two-way bond between conscious agents who are mutually aware, can respond to one another, and are morally accountable for the relationship. The case we have been examining falls cleanly into neither. The AI-driven virtual human exhibits real, individuated responsiveness, disqualifying it from the parasocial pole,—yet lacks the consciousness, autobiographical continuity, and moral agency that the social pole requires. It occupies the middle of the continuum, and that middle deserves a name (Figure 1).

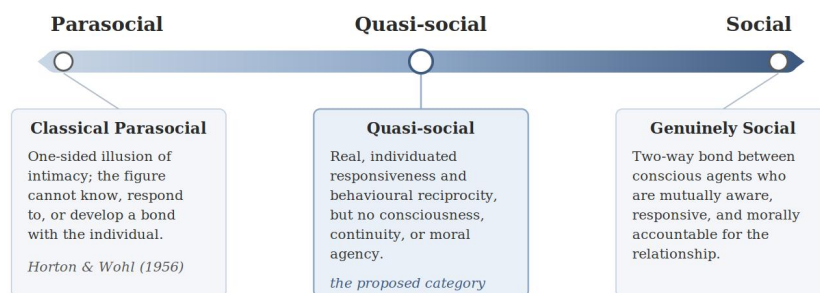


Figure 1. The parasocial-quasi-social-social continuum.

Figure 1. The parasocial–quasi-social–social continuum.

We adopt quasi-social over candidate alternatives deliberately. “Synthetic-social” foregrounds the artificiality of the figure but says little about the relational structure; “pseudo-reciprocal” implies the reciprocity is merely apparent, which understates the genuine behavioural contingency at work. The prefix quasi— ‘as if,’ ‘having some but not all features of’) captures the central fact precisely: the interaction has some, but not all, of the constitutive features of the social. It is more than parasocial because the contingency is real; it is less than social because the partner is not a conscious agent.

4.2 Defining Quasi-Social Interaction

We offer the following working definition.

DEFINITION

Quasi-social interaction is a machine-mediated relational mode in which a media figure responds contingently and individually to a user, producing genuine behavioural reciprocity while lacking the consciousness, autobiographical continuity, and moral agency that ground human social relationships.

Three elements of the definition carry the weight. Contingent and individual response distinguishes the mode from the parasocial, where no response exists. Genuine behavioural reciprocity specifies the sense in which the reciprocity is real: the figure's outputs are causally conditioned on this user's inputs, not merely perceived as such. And the closing clause—the absence of consciousness, continuity, and moral agency—distinguishes the mode from the genuinely social, marking what the machine partner still lacks despite its responsiveness.

4.3 Four Dimensions of Distinction

The three modes can be distinguished cleanly on four dimensions. Responsiveness asks whether the figure reacts to the individual user. Reciprocity asks whether any mutual exchange is real or merely felt. Consciousness and agency asks whether the figure is a sentient, accountable partner. Developability asks whether the relationship can evolve through adjustment. The table below summarizes the contrasts and is intended to serve as the paper's signature figure.

Dimension	Parasocial	Quasi-social	Social
Responsiveness to individual	None	Real, algorithmic	Real, human
Reciprocity	Illusory	Behavioural only	Genuine
Consciousness / agency	Absent	Absent	Present
Developability	None	System-driven	Mutual

Table 1. Distinguishing parasocial, quasi-social, and social interaction.

The table makes the central claim legible at a glance. Quasi-social interaction shares its responsiveness and developability with the social mode, which is why it cannot be parasocial, but shares its absence of consciousness with the parasocial mode, which is why it cannot be social. The single most consequential row is reciprocity: the quasi-social is the only mode in which reciprocity is behaviourally real yet experientially hollow, exchange without an experiencing partner on the other side.

4.4 Bounding the Concept

A new concept earns its place only if its borders are clear. We bound quasi-social interaction against three neighbours. It is not genuine friendship: the responsiveness, however sophisticated, is not backed by an experiencing, caring subject, and treating it as friendship invites the over-attribution we examine in Section 5. It is not classical parasocial interaction: unlike a television viewer's relationship with a newscaster, the contingency here is real, not merely felt, so the parasocial label misdescribes it. And it is not mere tool-use: a calculator or search engine is responsive without sustaining a relational frame, whereas the quasi-social figure presents a stable persona that the user relates to as a someone, not a something. Quasi-social interaction is the region that opens up precisely when responsiveness is real but the responsive entity is not a person.

One anticipated objection deserves a direct response. Stein, Breves, and Anders (2024) found that lower perceived human-likeness can suppress parasocial response to a virtual figure, which might seem to imply a 'humanness ceiling' limiting how far any such relationship can go. We do not dispute the empirical finding, but we read it differently: the ceiling constrains how social the experience feels, not whether the interaction is responsive. Quasi-social interaction is defined by the latter, not the former. A figure can be unconvincing as a human and still respond contingently and individually, and it is that responsiveness, not its perceived humanity, that places it outside the parasocial category.

With the category defined, distinguished, and bounded, the question becomes what difference it makes. Section 5 turns to the implications for how we measure these bonds, how we study their effects, and how we govern the platforms that manufacture them.

5. Implications

If the relationship between users and AI-driven virtual humans is quasi-social rather than parasocial, the reclassification is not merely terminological. It changes how these bonds should be measured, how their effects should be theorized, and how the platforms that produce them should be governed. We take each in turn.

5.1 For Measurement

The instruments the field relies on were built on the assumption of one-sidedness. The Experience of Parasocial Interaction (EPSI) scale developed by Hartmann and Goldhoorn (2011), for instance, captures a viewer's felt sense of mutual awareness, attention, and adjustment with a performer who is, in fact, unresponsive. When the figure genuinely responds, the construct such scales were validated to measure no longer matches the situation: a high EPSI

score might now reflect actual contingent exchange rather than a sustained illusion of it. If parasocial and quasi-social bonds are conflated at the level of measurement, studies risk attributing to ‘parasocial intensity’ effects that in fact arise from real responsiveness. The reclassification thus calls for instruments that distinguish felt reciprocity from behavioural reciprocity, rather than collapsing the two.

5.2 For Media-Effects Research

If reciprocity is real rather than imagined, the mechanisms behind trust, attachment, and dependence may operate differently. Anthropomorphic cues are widely understood as a double-edged sword: they can deepen engagement and trust while also rendering that trust difficult to calibrate. In a classic demonstration, Waytz, Heafner, and Epley (2014) showed that anthropomorphizing an autonomous vehicle—giving it a name, gender, and voice—increased participants' trust in it, even though the vehicle's actual capabilities were unchanged. Experimental evidence on the downside is equally sobering. In a pair of studies on AI recommendations under uncertainty, Holbrook and colleagues (2024) found that participants overwhelmingly reversed their own threat-identification judgments when an artificial agent disagreed, substantially degrading their performance—a striking demonstration of over-trust in unreliable AI. If a merely apparent relationship can already move judgments this far, a quasi-social bond grounded in real contingent responsiveness may amplify such effects rather than dampen them. The reclassification therefore reframes the central effects question: not ‘how strong is the illusion of relationship?’ but ‘what does a real-but-asymmetric relationship do to the user who is on the only conscious side of it?’

5.3 For Platform Ethics and Governance

Quasi-social relations may be more commercially exploitable than parasocial ones, precisely because their reciprocity is real. A television viewer's one-sided attachment could be leveraged only crudely; a figure that responds individually, remembers, and adapts can cultivate engagement that is deeper, more personalized, and more difficult for the user to step back from. Read through the lens of the shift from personalization to personification (Andrejevic & Volcic, 2025), the platform no longer merely holds attention but manufactures a responsive presence optimized against the user's own disclosures. This raises governance questions that the parasocial frame, with its assumption of a fixed and equally-broadcast performance, was not built to address—among them, whether the artificial and non-conscious nature of the responsive figure should be disclosed, and how the simulation of care should be regulated when it is deployed at scale toward commercial ends. We do not resolve these questions here; we note that naming the relationship correctly is the precondition for posing them well.

6. Conclusion and Propositions

This article has argued that the parasocial concept, built for the one-way media of the mid-twentieth century, cannot be applied without modification to AI-driven virtual humans. The contribution is threefold. Diagnostically, we identified three implicit premises of parasocial interaction—non-reciprocity, non-responsiveness, and performer-side control of content—and showed that responsive virtual humans violate all three at once. Conceptually, we proposed a parasocial-to-social continuum and defined a new intermediate category, quasi-social interaction,

distinguished from its neighbours along four dimensions. Generatively, we now translate the framework into propositions that future empirical work can test.

It bears emphasizing that AI did not create an entirely new puzzle. As we noted at the outset, a tension between objective one-sidedness and subjectively felt reciprocity was present in parasocial theory from Horton and Wohl (1956) onward. What AI-driven virtual humans do is close that gap on the objective side, turning a felt reciprocity into a behavioural one and thereby detonating a contradiction the concept had carried, unexploded, for nearly seventy years.

We close with five propositions. They are stated so as to be empirically falsifiable, and are intended less as settled claims than as an agenda for the research the reclassification makes possible.

P1

Perceived responsiveness, not perceived human-likeness, is the primary factor distinguishing quasi-social from parasocial bonds.

P2

Because of real contingent feedback, quasi-social bonds develop faster than parasocial bonds at equivalent exposure.

P3

The strength of a quasi-social bond is moderated by the user's awareness of the figure's artificiality: greater salience of artificiality weakens, but does not eliminate, the bond.

P4

At equivalent levels of perceived intimacy, quasi-social relations produce higher behavioural dependence than parasocial relations.

P5

Standard parasocial scales systematically mis-estimate bond strength for responsive virtual humans, because they cannot separate felt reciprocity from behavioural reciprocity.

Whether these propositions hold is a matter for the empirical work we hope this reconceptualization invites. Our claim is narrower and prior: that we cannot study human–virtual human relationships well while describing them with a concept whose defining conditions they no longer meet. When the persona replies, the parasocial frame reaches its limit, and a quasi-social one begins.

AUTHOR CONTRIBUTIONS

LIUFAN: conceptualization, investigation, and writing – original draft.

Muhammad Zaffwan Idris: supervision, conceptualization, and writing – review and editing.

Both authors read and approved the final manuscript.

FUNDING

The authors received no specific funding for this work.

DATA AVAILABILITY STATEMENT

Data sharing is not applicable to this article as no new data were created or analysed in this study.

CONFLICT OF INTEREST STATEMENT

The authors declare that they have no conflict of interest.

Reference

- 1.Andrejevic, M., & Volcic, Z. (2025). Automated parasociality: From personalization to personification. *Television & New Media*, 26(4), 421 – 437. <https://doi.org/10.1177/15274764241300436>
- 2.Biocca, F., Harms, C., & Burgoon, J. K. (2003). Toward a more robust theory and measure of social presence: Review and suggested criteria. *Presence: Teleoperators and Virtual Environments*, 12(5), 456 – 480. <https://doi.org/10.1162/105474603322761270>
- 3.Dibble, J. L., Hartmann, T., & Rosaen, S. F. (2016). Parasocial interaction and parasocial relationship: Conceptual clarification and a critical assessment of measures. *Human Communication Research*, 42(1), 21 – 44. <https://doi.org/10.1111/hcre.12063>
- 4.Epley, N., Waytz, A., & Cacioppo, J. T. (2007). On seeing human: A three-factor theory of anthropomorphism. *Psychological Review*, 114(4), 864 – 886. <https://doi.org/10.1037/0033-295X.114.4.864>
- 5.Giles, D. C. (2002). Parasocial interaction: A review of the literature and a model for future research. *Media Psychology*, 4(3), 279 – 305. https://doi.org/10.1207/S1532785XMEP0403_04
- 6.Guzman, A. L., & Lewis, S. C. (2020). Artificial intelligence and communication: A human – machine communication research agenda. *New Media & Society*, 22(1), 70 – 86. <https://doi.org/10.1177/1461444819858691>
- 7.Hartmann, T. (2023). Three conceptual challenges to parasocial interaction: Anticipated responses, implicit address, and the interactivity problem. In R. Tukachinsky Forster (Ed.), *The Oxford Handbook of Parasocial Experiences* (pp. 51 – 69). Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780197650677.013.3> [To confirm: page range appears as 51 – 69 and 51 – 68 across sources; verify against the printed chapter.]
- 8.Hartmann, T., & Goldhoorn, C. (2011). Horton and Wohl revisited: Exploring viewers' experience of parasocial interaction. *Journal of Communication*, 61(6), 1104 – 1121. <https://doi.org/10.1111/j.1460-2466.2011.01595.x>

- Hepp, A., Loosen, W., Dreyer, S., Jarke, J., Kannengießner, S., Katzenbach, C., Malaka, R., Pfadenhauer, M., Puschmann, C., & Schulz, W. (2023). Systematizing the field of research on human – machine communication: A conceptual framework. *Digital Journalism*, 11(1), 1 – 23. <https://doi.org/10.1080/21670811.2022.2152840>
9. Holbrook, C., Holman, D., Clingo, J., & Wagner, A. R. (2024). Overtrust in AI recommendations about whether or not to kill: Evidence from two human – robot interaction studies. *Scientific Reports*, 14, 19751. <https://doi.org/10.1038/s41598-024-69771-z>
10. Horton, D., & Strauss, A. (1957). Interaction in audience-participation shows. *American Journal of Sociology*, 62(6), 579 – 587. <https://doi.org/10.1086/222106>
11. Nass, C., & Moon, Y. (2000). Machines and mindlessness: Social responses to computers. *Journal of Social Issues*, 56(1), 81 – 103. <https://doi.org/10.1111/0022-4537.00153>
12. Horton, D., & Wohl, R. R. (1956). Mass communication and para-social interaction: Observations on intimacy at a distance. *Psychiatry*, 19(3), 215 – 229. <https://doi.org/10.1080/00332747.1956.11023049>
13. Lee, K. M. (2004). Presence, explicated. *Communication Theory*, 14(1), 27 – 50. <https://doi.org/10.1111/j.1468-2885.2004.tb00302.x>
14. Nass, C., Steuer, J., & Tauber, E. R. (1994). Computers are social actors. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 72 – 78). Association for Computing Machinery. <https://doi.org/10.1145/191666.191703>
15. Liu, D. (2025). Is interaction between human and artificial intelligence – driven agents (para)social? A scoping review. *Cyberpsychology, Behavior, and Social Networking*. Advance online publication. <https://doi.org/10.1089/cyber.2024.0532>
16. Reeves, B., & Nass, C. (1996). *The media equation: How people treat computers, television, and new media like real people and places*. Cambridge University Press.
17. Stein, J.-P., Breves, P. L., & Anders, N. (2024). Parasocial interactions with real and virtual influencers: The role of perceived similarity and human-likeness. *New Media & Society*, 26(6), 3433 – 3453. <https://doi.org/10.1177/14614448221102900>
18. Waytz, A., Heafner, J., & Epley, N. (2014). The mind in the machine: Anthropomorphism increases trust in an autonomous vehicle. *Journal of Experimental Social Psychology*, 52, 113 – 117. <https://doi.org/10.1016/j.jesp.2014.01.005>

Disclaimer/Publisher’s Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of *Journal of Contemporary Economics and Management* (JCEM) and/or its editor(s). JCEM and/or the editor(s) disclaim any responsibility for harm to persons or property resulting from any ideas, methods, instructions, or products referred to in the content.