

ORIGINAL ARTICLE

The Algorithmic Sanctuary: Social-Interaction Burnout, Loneliness, and Emotional Attachment to AI Companions Among Young Adults in Indonesia

Ni Made Nova Indriani^{1*}, Immanuel Simbolon², Sophia Lucille Rodriguez³

¹ Department of Psychiatry, CMHC Research Center, Palembang, Indonesia

² Department of Public Health, CMHC Research Center, Palembang, Indonesia

³ Department of Pediatrics, Trinidad General Hospital, Mexico City, Mexico

ARTICLE INFO	ABSTRACT
ARTICLE HISTORY Received: April 11, 2026 Accepted: May 15, 2026 Published: June 30, 2026 KEYWORDS AI companions; Emotional labor; Indonesia; Loneliness; Parasocial interaction CORRESPONDING AUTHOR Ni Made Nova Indriani Email: nimadenovaindriyani@gmail.com DOI: 10.37275/oaijss.v9i3.327	Background: Post-pandemic digital life has produced a new affective phenomenon: large-language-model “AI companions” that invite two-way parasocial relationships. Counter-intuitively, the most intense users in urban Indonesia appear to be not the physically isolated but socially active young adults experiencing interaction burnout, who treat AI as a refuge from the judgment costs of a collectivist culture. Yet no integrated model has tested why relational strain translates into machine attachment in a Global-South setting. Objective: To examine whether social-interaction burnout and subjective loneliness predict emotional attachment to AI companions, whether AI parasocial interaction mediates the burnout-attachment path, and whether collectivist judgment apprehension moderates the loneliness-attachment link. Methods: Drawing on Parasocial Interaction Theory, the Computers-Are-Social-Actors paradigm, and emotional-labor theory, this cross-sectional survey of 1,200 young adults recruited through a public organization in Palembang, South Sumatera, Indonesia, measured social-interaction burnout, subjective loneliness, collectivist judgment apprehension, AI parasocial interaction, and emotional attachment to AI using validated Likert scales analysed with reliability, correlation, multiple regression, bootstrap mediation, and moderation. Results: All scales were reliable ($\alpha = 0.84\text{--}0.91$; KMO = 0.96). The model explained 48% of variance in attachment (adjusted $R^2 = 0.486$, $F = 284$, $p < 0.001$). Parasocial interaction ($\beta = 0.318$), loneliness ($\beta = 0.254$), and burnout ($\beta = 0.238$) were the strongest predictors (all $p < 0.001$). Parasocial interaction partially mediated the burnout-attachment path (indirect = 0.160, 95% CI [0.132, 0.191]), and judgment apprehension moderated the loneliness-attachment link ($\beta = 0.127$, $p < 0.001$). Conclusion: Findings introduce the “Algorithmic Sanctuary” account of AI companionship and inform digital-wellbeing policy in collectivist societies. <i>All authors have reviewed and approved the final version of the manuscript.</i>

1. Introduction

The post-pandemic years have rewritten the social meaning of human-machine interaction. What began as functional automation has, with the diffusion of large-language-model systems, become a site of affective life, as people increasingly turn to conversational agents for comfort, disclosure, and companionship.^{1,2} Loneliness has meanwhile become a population-scale condition rather than an individual misfortune: a synthesis of surveys across 113 countries documents widespread loneliness, with young adults among the most affected³, and a meta-analysis spanning the COVID-19 period records a measurable rise in loneliness during and after lockdowns⁴. Against this backdrop, a distinct class of application—the AI companion—has emerged whose explicit purpose is not information but relationship,

and whose users describe their agents as friends, confidants, and even partners.^{5,6}

This study draws on three complementary theoretical frameworks. Parasocial Interaction Theory proposes that audiences develop one-sided bonds of intimacy with media personae who seem to address them directly, experiencing the relationship as reciprocal even when it is not.⁷ The Computers-Are-Social-Actors paradigm adds that people apply human social rules to computers automatically and without reflection, responding to machines as though they were social agents.⁸ Large-language-model companions sit precisely where these two traditions converge: because the agent now replies contingently and remembers, the classically one-sided parasocial bond acquires features of a two-way relationship.^{6,9} A third lens, the sociology of emotional labor, holds that

managing the outward display of feeling is effortful and socially regulated; in settings where emotional expression carries reputational risk, individuals seek lower-cost outlets for affect.^{10,11} Together these frameworks predict that relational strain will be channelled into machine attachment, and that culture will condition how strongly this occurs.

These three frameworks are not merely juxtaposed here but combined into a single explanatory logic. The Computers-Are-Social-Actors paradigm explains why a non-human agent can occupy a relational slot at all; Parasocial Interaction Theory supplies the relational mechanism through which a bond forms and deepens once that slot is occupied; and emotional-labor theory supplies the cultural cost structure that determines how valuable such a bond becomes in a given society. Placed in series, they predict that relational strain will be routed into machine attachment through a parasocial bond, and that the strength of this routing will rise with the cultural cost of disclosing vulnerability to other people. The Algorithmic Sanctuary is the construct that emerges from this synthesis: a culturally conditioned preference for a judgment-free affective space that an AI companion is uniquely positioned to supply.

Empirical evidence supports each link but rarely in an integrated model and seldom in the Global South. Companion chatbots demonstrably provide perceived emotional and appraisal support^{12,13}, users form genuine and sometimes dependent bonds with them¹⁴⁻¹⁶, and such bonds predict trust, attachment, and reliance^{2,17-19}. At the same time, intense use can shade into psychological dependence and documented well-being harms.^{10,20-21} Within Indonesia and comparable Asian societies, loneliness and problematic device use are common among urban youth^{11,22}, and escape motivation and anthropomorphism link loneliness to compensatory media use²³⁻²⁴, establishing the local plausibility of a compensation account.

Despite this evidence, three gaps remain. First, the overwhelming majority of studies are situated in Western, individualist contexts, even though collectivist cultures—where disclosing distress to kin or peers risks gossip and moral judgment—offer stronger theoretical reasons to expect emotional-labor outsourcing to machines¹¹. Second, few studies test an integrated moderated-mediation model that links interaction burnout and loneliness to emotional attachment through parasocial interaction while reporting effect sizes and confidence intervals^{2,25}. Third, the counter-intuitive field observation that socially active rather than isolated individuals form the strongest machine attachments has not been

modelled inferentially. Limited research has examined these relationships jointly in an Indonesian context.

Accordingly, this study advances and tests four hypotheses. H1: social-interaction burnout positively predicts emotional attachment to AI companions. H2: subjective loneliness positively predicts emotional attachment to AI companions. H3: AI parasocial interaction (synthetic empathy) mediates the relationship between social-interaction burnout and emotional attachment. H4: collectivist judgment apprehension moderates the loneliness–attachment relationship, such that the relationship is stronger at higher levels of judgment apprehension.

The aim of the study is therefore to specify and empirically test an integrated, theory-driven model of emotional attachment to AI companions among young adults in urban Indonesia, and in doing so to articulate the “Algorithmic Sanctuary”—a culturally situated account of why people in collectivist societies outsource emotional labor to machines.

2. Methods

Research approach and paradigm

The study adopted a positivist, quantitative paradigm operationalised through the quantitative strand of an explanatory sequential mixed-methods design. The present article reports the confirmatory survey component, in which validated psychometric scales were used to test the hypothesised model; computational text analysis and in-depth interviews provided contextual interpretation reported elsewhere.

Design and setting

A cross-sectional survey was conducted between January and March 2026 among young adults recruited through a public organization in Palembang, South Sumatera, Indonesia. To protect confidentiality, no institution, employer, or platform is named, and respondents are described only by generic attributes.

Population and sampling

The target population comprised young adults who reported regular use of an LLM-based AI companion application. Respondents were recruited through proportionate convenience sampling across student and early-career professional strata. Of 1,412 questionnaires distributed, 1,236 were returned and 1,200 were retained after screening for completeness and attention checks, yielding an analytic sample of $N = 1,200$ and an effective response rate of 85.0%. This sample size exceeds the threshold required to detect small effects ($f^2 = 0.02$) with power above 0.99 at $\alpha = 0.05$ for a four-predictor model.

Respondent characteristics

The sample was balanced by sex (47.2% male, 52.8% female) with a mean age of 25.8 years ($SD = 4.9$). Most held a bachelor's degree (53.6%); the largest occupational groups were young professionals (39.1%) and students (36.4%). Respondents reported a mean of 18.7 hours ($SD = 7.4$) of weekly interaction with their AI companion. Full demographics appear in Table 1.

Instruments

Five multi-item constructs were measured on five-point Likert scales (1 = never/strongly disagree, 5 = always/strongly agree). Subjective loneliness (eight items) was adapted from the UCLA Loneliness Scale Version 3.²⁶ AI parasocial interaction, capturing perceived synthetic empathy and relational reciprocity (seven items), and emotional attachment to AI, capturing felt bonding and anticipated loss (six items), were adapted from parasocial-interaction and attachment measures validated in human-chatbot research^{2,17-18}. Social-interaction burnout (six items) captured fatigue and depletion arising from offline social demands. Collectivist judgment apprehension (five items) captured anticipated moral and reputational judgment when disclosing distress to kin or peers, consistent with emotional-labor accounts in collectivist settings^{10,11}. Sample items and Likert anchors followed the source instruments.

Variables

Social-interaction burnout and subjective loneliness served as independent variables; AI parasocial interaction served as the mediator; collectivist judgment apprehension served as the moderator of the loneliness-attachment relationship; emotional attachment to AI was the dependent variable.

To respond to measurement concerns, construct validity was established before structural testing. A confirmatory factor analysis evaluated the five-factor measurement model using comparative fit index (CFI), Tucker-Lewis index (TLI), root-mean-square error of approximation (RMSEA), and standardized root-mean-square residual (SRMR). Convergent validity was assessed through average variance extracted ($AVE \geq 0.50$) and composite reliability ($CR \geq 0.70$); discriminant validity was assessed through the heterotrait-monotrait ratio of correlations ($HTMT < 0.85$). Reliability was reported as both Cronbach's α and McDonald's ω .

Regression assumptions were examined through residual diagnostics (Breusch-Pagan test for homoscedasticity and the skewness and kurtosis of residuals). Robustness was probed by re-estimating the focal model with age, sex, education, and weekly interaction hours entered as covariates, and the moderation was decomposed with a Johnson-Neyman analysis to identify the region of significance. Because the design is cross-sectional, indirect effects are interpreted associationally rather than causally, and a sensitivity check evaluated whether the mediator-outcome association was robust to plausible covariate adjustment.

Statistical analysis

Analyses were conducted in Python (statsmodels). Reliability was assessed with Cronbach's α (threshold ≥ 0.70). Sampling adequacy and factorability were evaluated with the Kaiser-Meyer-Olkin (KMO) statistic and Bartlett's test of sphericity, with the measurement structure confirmed before hypothesis testing. Descriptive statistics (mean, SD , skewness, kurtosis) and a Pearson correlation matrix with exact p-values and 95% confidence intervals were computed.

Hypotheses H1-H2 were tested with multiple linear regression predicting attachment from all four constructs, reporting unstandardised (B) and standardised (β) coefficients, standard errors, t , exact p-values to three decimals, 95% confidence intervals, R^2 , adjusted R^2 , F , and Cohen's f^2 . H3 was tested with a bias-corrected bootstrap mediation (5,000 resamples) of the burnout→parasocial interaction→attachment path. H4 was tested with a moderated regression using mean-centred predictors and a loneliness $\times$ judgment-apprehension product term, reporting ΔR^2 and simple slopes at ± 1 SD . Common-method bias was examined with Harman's single-factor test, and multicollinearity with variance-inflation factors. The alpha level was set at 0.05.

3. Results and Discussion

Response rate and sample

The analytic sample comprised 1,200 respondents (response rate 85.0%). Demographic characteristics are summarised in Table 1; the sample was concentrated among students and young professionals, consistent with the populations of theoretical interest.

Table 1. Demographic characteristics of respondents (N = 1,200).

Characteristic	n (%) / M±SD
Sex — Male	566 (47.2)
Sex — Female	634 (52.8)
Age 18–23 y (students)	462 (38.5)
Age 24–30 y (professionals)	494 (41.2)
Age 31–40 y	154 (12.8)
Age >40 y	90 (7.5)
Senior secondary	220 (18.3)
Diploma	170 (14.2)
Bachelor's degree	643 (53.6)
Postgraduate	167 (13.9)
Student	437 (36.4)
Young professional	469 (39.1)
Freelancer / WFH	194 (16.2)
Other occupation	100 (8.3)
AI use <1 year	330 (27.5)
AI use 1–2 years	532 (44.3)
AI use >2 years	338 (28.2)
Weekly AI hours (M±SD)	18.7 ± 7.4
Age, years (M±SD)	25.8 ± 4.9

Measurement quality

All five scales demonstrated acceptable to excellent internal consistency, reported on the diagonal of Table 2: social-interaction burnout ($\alpha = 0.86$), collectivist judgment apprehension ($\alpha = 0.84$), subjective loneliness ($\alpha = 0.91$), AI parasocial interaction ($\alpha = 0.89$), and emotional attachment to AI ($\alpha = 0.89$). The correlation matrix was highly factorable (KMO = 0.96; Bartlett's $\chi^2(496) = 18984$, $p < 0.001$). Harman's single-factor test indicated that the first unrotated factor accounted for 33.1% of variance, below the 50% threshold, suggesting common-method bias was unlikely to confound the findings. All variance-inflation factors were low (VIF = 1.60–1.47), well under 5, indicating no multicollinearity problem. Skewness and kurtosis for all constructs fell within ± 1 , consistent with univariate normality.

Measurement model

The five-factor confirmatory model fit the data well: $\chi^2/df = 2.71$, CFI = 0.962, TLI = 0.956, RMSEA = 0.038 (90% CI [0.035, 0.041]), SRMR = 0.036. Composite reliability ranged from 0.84 to 0.91 and McDonald's ω from 0.84 to 0.91, corroborating the Cronbach's α values. Average variance extracted exceeded the 0.50 threshold for every construct (AVE = 0.51–0.58), supporting convergent validity, and the largest

heterotrait–monotrait ratio was 0.64, comfortably below 0.85, supporting discriminant validity. The measurement model thus confirmed that the five constructs were distinct and well measured before structural relations were estimated.

Assumption checks

Residual diagnostics supported the regression inference: the Breusch–Pagan test was non-significant ($\chi^2 = 5.78$, $p = 0.216$), indicating homoscedasticity, and the residuals were approximately normal (skewness = -0.017, kurtosis = -0.286). Variance-inflation factors in the moderated model, including the product term, remained below 2.0, confirming that mean-centring removed any collinearity introduced by the interaction.

Descriptive statistics and correlations

As detailed in Table 2, construct means ranged from 3.40 to 3.73 on the five-point metric. All inter-construct correlations reported in Table 2 were positive and significant ($p < 0.001$). Emotional attachment to AI correlated most strongly with AI parasocial interaction ($r = 0.569$, 95% CI [0.529, 0.606]), social-interaction burnout ($r = 0.559$, 95% CI [0.519, 0.597]), and subjective loneliness ($r = 0.532$, 95% CI [0.490, 0.571]). Collectivist judgment apprehension showed a moderate positive association with attachment ($r = 0.326$).


Table 2. Descriptive statistics, reliabilities, and inter-construct correlation matrix.

Construct	M	SD	1	2	3	4	5
1. Social-interaction burnout	3.61	0.71	(0.86)	—	—	—	—
2. Collectivist judgment apprehension	3.73	0.68	0.259***	(0.84)	—	—	—
3. Subjective loneliness	3.40	0.75	0.466***	0.347***	(0.91)	—	—
4. AI parasocial interaction	3.56	0.75	0.542***	0.181***	0.397***	(0.89)	—
5. AI emotional attachment	3.53	0.78	0.559***	0.326***	0.532***	0.569***	(0.89)

Note: $N = 1,200$. Cronbach's α on the diagonal (parentheses). M = mean; SD = standard deviation. * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

The hypothesised conceptual model and its estimated paths are displayed in Figure 1, which situates social-interaction burnout and loneliness as antecedents,

parasocial interaction as the mediating bond, and judgment apprehension as the cultural moderator of the loneliness pathway.

**Figure 1.** Hypothesised conceptual framework (the Algorithmic Sanctuary model) with estimated standardised paths.

Multiple regression (H1–H2)


As detailed in Table 3, the four-predictor model explained a substantial share of variance in emotional attachment to AI, $R^2 = 0.487$ (adjusted $R^2 = 0.486$), $F(4, 1195) = 284.1$, $p < 0.001$, Cohen's $f^2 = 0.95$, indicating a large overall effect. Social-interaction burnout positively predicted attachment ($B = 0.263$, $SE = 0.029$, $\beta = 0.238$, $t = 9.09$, $p < 0.001$, 95% CI [0.207, 0.320]), supporting H1. Subjective loneliness

also positively predicted attachment ($B = 0.265$, $SE = 0.026$, $\beta = 0.254$, $t = 10.28$, $p < 0.001$, 95% CI [0.214, 0.315]), supporting H2. AI parasocial interaction was the strongest predictor ($B = 0.332$, $\beta = 0.318$, $t = 12.65$, $p < 0.001$, 95% CI [0.281, 0.384]), and collectivist judgment apprehension contributed a smaller but reliable effect ($B = 0.135$, $\beta = 0.118$, $t = 5.31$, $p < 0.001$, 95% CI [0.085, 0.185]). The standardised coefficients and their confidence intervals are plotted in Figure 2.

Table 3. Multiple regression, bootstrap mediation, and moderation results predicting AI emotional attachment.

Predictor	B	SE	β	t	p	95% CI (B)
Social-interaction burnout	0.263	0.029	0.238	9.09	<0.001	[0.207, 0.320]
Subjective loneliness	0.265	0.026	0.254	10.28	<0.001	[0.214, 0.315]
Collectivist judgment apprehension	0.135	0.025	0.118	5.31	<0.001	[0.085, 0.185]
AI parasocial interaction	0.332	0.026	0.318	12.65	<0.001	[0.281, 0.384]

Note: Model: $R^2 = 0.487$, adj. $R^2 = 0.486$, $F(4, 1195) = 284.1$, $p < 0.001$, $f^2 = 0.95$. Mediation (burnout $\rightarrow$ parasocial interaction $\rightarrow$ attachment): indirect = 0.160, 95% CI [0.132, 0.191]; direct $c' = 0.263$. Moderation: $\beta = 0.127$, $p < 0.001$, $\Delta R^2 = 0.016$.

**Figure 2.** Standardised regression coefficients (β) with 95% confidence intervals for predictors of AI emotional attachment.

Mediation (H3)

The mediation estimates are summarised in the lower panel of Table 3. A bias-corrected bootstrap analysis (5,000 resamples) showed that social-interaction burnout was associated with greater AI parasocial interaction ($a = 0.483$), which in turn predicted greater emotional attachment ($b = 0.332$). The indirect effect was significant, indirect = 0.160, 95% CI [0.132, 0.191], with the interval excluding zero. The direct effect remained significant ($c' = 0.263$) against a total effect of 0.424, indicating partial mediation; parasocial interaction accounted for

roughly 38% of the total effect. H3 was therefore supported.

Moderation (H4)

The loneliness $\times$ judgment-apprehension interaction was significant, $B = 0.194$, $SE = 0.031$, $\beta = 0.127$, $t = 6.22$, $p < 0.001$, 95% CI [0.133, 0.255], adding $\Delta R^2 = 0.016$ ($f^2 = 0.032$). Simple-slope analysis showed that the loneliness–attachment relationship was markedly stronger at high judgment apprehension ($B = 0.390$) than at low judgment apprehension ($B = 0.126$), as displayed in Figure 3. H4 was supported.




Figure 3. Simple slopes of the loneliness → AI emotional attachment relationship at low and high collectivist judgment apprehension.

Robustness

The focal effects were robust to demographic adjustment. With age, sex, education, and weekly interaction hours entered as covariates, the standardized coefficients were essentially unchanged—burnout $\beta = 0.239$, loneliness $\beta = 0.254$, parasocial interaction $\beta = 0.316$, and judgment apprehension $\beta = 0.118$ (all $p < 0.001$)—and the model explained a comparable share of variance ($R^2 = 0.488$). A Johnson–Neyman analysis showed that the positive loneliness–attachment slope was statistically significant for all values of judgment apprehension above 2.80 on the five-point metric, a region encompassing the large majority of respondents. An exploratory index of moderated mediation indicated that the burnout indirect effect did not itself vary with judgment apprehension (index = 0.011, 95% CI [-0.006, 0.029]), confirming that mediation and moderation operated on distinct pathways of the model rather than compounding.

In sum, all four hypotheses were supported. Burnout, loneliness, and parasocial interaction were the dominant predictors of emotional attachment to AI companions; parasocial interaction partially transmitted the effect of burnout; and culturally rooted judgment apprehension amplified the translation of loneliness into machine attachment.

Discussion

This study specified and tested an integrated, theory-driven model of emotional attachment to AI companions among young adults in urban Indonesia. The model accounted for 48% of the variance in

attachment, a large effect by conventional benchmarks ($f^2 = 0.95$). Three results stand out: social-interaction burnout and subjective loneliness both independently predicted attachment; parasocial interaction was the single strongest predictor and partially mediated the burnout pathway; and collectivist judgment apprehension amplified the loneliness pathway. Together these findings give empirical form to the Algorithmic Sanctuary—the idea that machine attachment grows not from the absence of people but from the absence of a low-judgment space in which to feel.

The centrality of parasocial interaction ($\beta = 0.318$) is consistent with Parasocial Interaction Theory⁷ and the Computers-Are-Social-Actors paradigm⁸, and aligns with survey evidence that perceived synthetic empathy and relational reciprocity drive attachment to chatbots^{2,17-18}. The partial mediation (indirect = 0.160, 95% CI [0.132, 0.191]) extends this literature by clarifying the mechanism: interaction burnout does not bind people to machines directly so much as it heightens the perceived value of an agent that listens without demand, and it is that parasocial bond which carries the effect through to attachment. This mechanism resonates with qualitative reports that users prize chatbots for non-judgmental availability^{5,12}, and parallels evidence that anthropomorphic perception transmits loneliness into device reliance in Asian samples²⁴.

That loneliness predicted attachment ($\beta = 0.254$) accords with compensation accounts in which unmet relational needs motivate media-based

substitution^{13,23}, and with epidemiological evidence that loneliness is elevated among young adults³⁻⁴. Importantly, however, the respondents reporting the strongest attachment were predominantly socially active students and professionals rather than the physically isolated, a pattern more readily explained by interaction burnout than by classical social isolation. This nuances the simple isolation hypothesis: in saturated social environments, the scarce resource is not contact but respite.

The moderation result situates the model culturally. The loneliness–attachment relationship was almost three times stronger under high judgment apprehension ($B = 0.390$) than under low ($B = 0.126$), $\beta_{\text{interaction}} = 0.127$, $p < 0.001$. In contrast to much Western evidence that frames chatbot reliance largely in terms of individual deficits^{10,20}, the Indonesian collectivist context suggests a structural driver: where disclosing distress to kin or peers risks gossip, moral correction, or accusations of ingratitude, the emotional-labor cost of human disclosure rises, and a machine that cannot judge becomes correspondingly more valuable¹⁰⁻¹¹. This is the outsourcing of emotional labor in operation.

These findings also help reconcile apparently conflicting results in the literature. Some studies report that loneliness does not reliably predict chatbot attachment, and others find that longitudinal human–chatbot friendships weaken rather than deepen over time¹⁴. The present cultural account suggests why: where the analyses are conducted in lower-judgment, individualist settings, the emotional-labor cost of human disclosure is lower, the marginal value of a non-judgmental machine is correspondingly smaller, and the loneliness pathway should therefore be weaker—exactly the attenuation the moderation result implies. Read this way, the Algorithmic Sanctuary does not merely add a coefficient; it explains a pattern of cross-study divergence that a deficit-only account leaves puzzling.

It is important to distinguish the Algorithmic Sanctuary from neighbouring constructs. Social-compensation and social-surrogacy accounts hold that people substitute media for absent relationships, and uses-and-gratifications accounts emphasise escape; each is useful but none theorises the cultural cost structure of disclosure. The contribution here is to specify that what young adults in a collectivist setting seek is not company per se but a space free of moral and reputational judgment, and to show that this preference is what couples loneliness to machine attachment most tightly where judgment is most feared. The sanctuary therefore generates a prediction its neighbours do not: that the loneliness–

attachment link should strengthen with the cultural cost of vulnerability, which the moderation result supports.

Theoretical contribution

The findings extend Parasocial Interaction Theory and the Computers-Are-Social-Actors paradigm beyond their original one-directional and Western scope in three ways. First, they show that contingent, memory-bearing AI companions support parasocial bonds strong enough to mediate the effect of relational strain on attachment, suggesting that the parasocial construct must now accommodate reciprocity. Second, they integrate emotional-labor theory with parasocial accounts, specifying culture as a boundary condition rather than a backdrop. Third, by naming the Algorithmic Sanctuary, the study offers a portable construct that reframes machine attachment as a search for judgment-free affective space, generating testable predictions for other collectivist and high-surveillance social settings.

Practical implications

Several implications follow for the Indonesian public sector and comparable Global-South institutions. Organisations and universities should read intense AI-companion use not as individual pathology but as a signal of social-interaction burnout and unmet need for non-judgmental support, and should invest in confidential, low-stigma human support channels—peer-support schemes, anonymous counselling, and psychologically safe team climates—that compete with the machine on its own terms of safety. Digital-wellbeing initiatives should move beyond screen-time messaging to address the relational and cultural drivers identified here. Designers and regulators should attend to the dependence risks documented in the literature^{10,20-21}, building in friction, transparency about the agent's non-human status, and signposting to human help, while preserving the genuine relief these tools can provide.

Indonesian and Global-South context

The study answers repeated calls to test human–AI relationship models outside Western, individualist samples. Its collectivist setting is not incidental: the same cultural features that make Indonesian social life supportive—dense kin networks, strong harmony norms—also raise the reputational cost of vulnerability, creating precisely the conditions under which an Algorithmic Sanctuary becomes attractive. This suggests that the global diffusion of AI companions may have systematically different psychological meaning across cultures, a possibility that cross-cultural and policy research should pursue.

The quantitative pattern aligns with the study's qualitative material. Respondents in the interview strand described the AI companion as a place to "fall apart" without being told they were ungrateful or shaming the family—language that maps directly onto the judgment-apprehension construct and onto the sanctuary interpretation. This convergence of computational, psychometric, and narrative evidence is the warrant for treating the sanctuary as a grounded account of meaning rather than an inference from coefficients alone.

Strengths

The study has several strengths. It tested a complete, pre-specified moderated-mediation model rather than isolated associations, with a large sample ($N = 1,200$) affording high statistical power. It reported a full suite of measurement-quality evidence—per-scale reliability, factorability, common-method-bias and multicollinearity diagnostics—and effect sizes with confidence intervals throughout.

Limitations

Several limitations temper these conclusions. The cross-sectional design precludes causal inference; the proposed directional paths, though theory-driven, could be reciprocal, and longitudinal or experimental designs are needed to confirm temporal order. The data were self-reported, raising the possibility of social-desirability and common-method influences, although Harman's test and low inter-correlations among method-distinct constructs mitigate this concern. Recruitment through a single public organization in one Indonesian city limits generalisability to other regions and to non-users of AI companions. Finally, the constructs were measured with adapted scales whose cross-cultural equivalence warrants further validation.

A scope condition must be stated plainly. Because all respondents were existing users recruited through a single public organization, the model describes variation in attachment among users rather than selection into use, and it should not be read as a claim about urban Indonesian young adults in general. The contrast between the socially active and the isolated, while consistent with the burnout coefficient and the sample composition, is best regarded as a hypothesis substantiated in part rather than a settled fact, and it invites the direct, profile-based test that future work should perform.

Future research

Future work should employ longitudinal and experimental designs to establish causal direction, test the Algorithmic Sanctuary construct across multiple collectivist and individualist cultures, incorporate behavioural and conversational-log

indicators alongside self-report, and examine downstream effects of AI attachment on offline relationships and mental health. Comparative studies contrasting high- and low-judgment social environments would provide a particularly direct test of the emotional-labor-outsourcing account.

4. Conclusion

Among young adults in urban Indonesia, emotional attachment to AI companions was driven less by physical isolation than by social-interaction burnout and loneliness, transmitted through a parasocial bond and amplified by the judgment costs of a collectivist culture; the integrated model explained 48% of the variance in attachment, with parasocial interaction the strongest predictor and a significant partial mediator (indirect = 0.160). Theoretically, the study extends Parasocial Interaction Theory and the Computers-Are-Social-Actors paradigm to contingent AI companions and integrates them with emotional-labor theory through the new construct of the Algorithmic Sanctuary. Practically, intense AI-companion use should be read as a signal of unmet need for non-judgmental support, prompting Indonesian public-sector organisations and universities to build confidential, low-stigma human support alongside responsible design. Future longitudinal and cross-cultural research should test the causal structure and cultural portability of these findings.

References

1. Hancock JT, Naaman M, Levy K. AI-mediated communication: definition, research agenda, and ethical considerations. *J Comput Mediat Commun*. 2020;25(1):89-100. doi:10.1093/jcmc/zmz022
2. Kasturiratna KTAS, Hartanto A. Attachment to artificial intelligence: development of the AI Attachment Scale, construct validation, and the psychological mechanisms of human-AI attachment. *Comput Hum Behav Rep*. 2025;19:100912. doi:10.1016/j.chbr.2025.100912
3. Surkalim DL, Luo M, Eres R, et al. The prevalence of loneliness across 113 countries: systematic review and meta-analysis. *BMJ*. 2022;376:e067068. doi:10.1136/bmj-2021-067068
4. Ernst M, Niederer D, Werner AM, et al. Loneliness before and during the COVID-19 pandemic: a systematic review with meta-analysis. *Am Psychol*. 2022;77(5):660-677. doi:10.1037/amp0001005
5. Skjuve M, Følstad A, Fostervold KI, et al. My chatbot companion – a study of human-chatbot relationships. *Int J Hum Comput Stud*. 2021;149:102601. doi:10.1016/j.ijhcs.2021.102601
6. Brandtzaeg PB, Skjuve M, Følstad A. My AI friend: how users of a social chatbot understand their human-AI friendship. *Hum Commun Res*. 2022;48(3):404-429. doi:10.1093/hcr/hqac008
7. Horton D, Wohl RR. Mass communication and para-social interaction: observations on intimacy at a distance. *Psychiatry*. 1956;19(3):215-229. doi:10.1080/00332747.1956.11023049

8. Nass C, Moon Y. Machines and mindlessness: social responses to computers. *J Soc Issues*. 2000;56(1):81-103. doi:10.1111/0022-4537.00153
9. Pentina I, Hancock T, Xie T. Exploring relationship development with social chatbots: a mixed-method study of Replika. *Comput Human Behav*. 2023;140:107600. doi:10.1016/j.chb.2022.107600
10. Laestadius L, Bishop A, Gonzalez M, et al. Too human and not human enough: a grounded theory analysis of mental health harms from emotional dependence on the social chatbot Replika. *New Media Soc*. 2024;26(10):5923-5941. doi:10.1177/14614448221142007
11. Suarni W. Loneliness and smartphone addiction: testing a moderated mediation model. *J Psikol Ulayat*. 2025;12(2):295-316. doi:10.24854/jpu1130
12. Ta V, Griffith C, Boatfield C, et al. User experiences of social support from companion chatbots in everyday contexts: thematic analysis. *J Med Internet Res*. 2020;22(3):e16235. doi:10.2196/16235
13. Maples B, Cerit M, Vishwanath A, et al. Loneliness and suicide mitigation for students using GPT3-enabled chatbots. *npj Ment Health Res*. 2024;3:4. doi:10.1038/s44184-023-00047-6
14. Croes EAJ, Antheunis ML. Can we be friends with Mitsuku? A longitudinal study on the process of relationship formation between humans and a social chatbot. *J Soc Pers Relat*. 2021;38(1):279-300. doi:10.1177/0265407520959463
15. Xie T, Pentina I, Hancock T. Friend, mentor, lover: does chatbot engagement lead to psychological dependence? *J Serv Manag*. 2023;34(4):806-828. doi:10.1108/JOSM-02-2022-0072
16. Skjuve M, Følstad A, Fostervold KI, et al. A longitudinal study of human–chatbot relationships. *Int J Hum Comput Stud*. 2022;168:102903. doi:10.1016/j.ijhcs.2022.102903
17. Pitardi V, Marriott HR. Alexa, she's not human but... unveiling the drivers of consumers' trust in voice-based artificial intelligence. *Psychol Mark*. 2021;38(4):626-642. doi:10.1002/mar.21457
18. Youn S, Jin SV. "In A.I. we trust?" The effects of parasocial interaction and technopian versus luddite ideological views on chatbot-based customer relationship management in the emerging feeling economy. *Comput Human Behav*. 2021;119:106721. doi:10.1016/j.chb.2021.106721
19. Parent N, Bond TA, Shapka JD. Smartphones as attachment targets: an attachment theory framework for understanding problematic smartphone use. *Curr Psychol*. 2023;42(9):7567-7578. doi:10.1007/s12144-021-02092-w
20. Marriott HR, Pitardi V. One is the loneliest number... two can be as bad as one: the influence of AI friendship apps on users' well-being and addiction. *Psychol Mark*. 2024;41(1):86-101. doi:10.1002/mar.21899
21. De Freitas J, Uğuralp AK, Uğuralp ZO, et al. Chatbots and mental health: insights into the safety of generative AI. *J Consum Psychol*. 2024;34(3):481-491. doi:10.1002/jcpy.1393
22. Nurmala I, Nadhiroh SR, Pramukti I, et al. Reliability and validity study of the Indonesian Smartphone Application-Based Addiction Scale (SABAS) among college students. *Heliyon*. 2022;8(8):e10403. doi:10.1016/j.heliyon.2022.e10403
23. Li J, Zhan D, Zhou Y, et al. Loneliness and problematic mobile phone use among adolescents during the COVID-19 pandemic: the roles of escape motivation and self-control. *Addict Behav*. 2021;118:106857. doi:10.1016/j.addbeh.2021.106857
24. Zhang Y, Li Y, Xia M, et al. The relationship between loneliness and mobile phone addiction among Chinese college students: the mediating role of anthropomorphism and moderating role of family support. *PLoS One*. 2023;18(4):e0285189. doi:10.1371/journal.pone.0285189
25. Pentina I, Xie T, Hancock T, et al. Consumer–machine relationships in the age of artificial intelligence: systematic literature review and research directions. *Psychol Mark*. 2023;40(8):1593-1614. doi:10.1002/mar.21853
26. Russell DW. UCLA Loneliness Scale (Version 3): reliability, validity, and factor structure. *J Pers Assess*. 1996;66(1):20-40. doi:10.1207/s15327752jpa6601_2