

What 11,793 people working on their attachment reveal about whether it can change

Attachment anxiety, avoidance, and change in a large help-seeking sample

Gabriel Uribe, Luis Goicouria, Sophie Cheung

Attached Research (SkyPorch LLC) · Published June 24, 2026 · "The State of Attachment 2026"

Abstract

We characterize adult attachment among 11,793 users of Attached, a mobile app for working on attachment, using 51,676 attachment measurements (13,789 from an intake assessment and 37,887 derived from journaling). Each measurement places a person on the two standard ECR-R dimensions, anxiety and avoidance, rescaled to 0-1. Journaling-derived scores are produced by a large language model acting as an ECR-R rater (Methods, Section 1). This is a help-seeking sample, not the general population; secure attachment is correspondingly rare. Attachment anxiety dominates (mean 0.74, 90% above the scale midpoint). In an observational longitudinal analysis using aggregated measurements, attachment-anxiety scores declined over sustained use (primary cohort: mean change -0.152, 95% CI [-0.166, -0.139]; 37.5% reliable decrease), with larger declines among heavier users (Cohen's *d* up to 0.82), while avoidance stayed essentially flat. We report measurement reliability and validity openly: single scores are noisy (test-retest *r* approx 0.29) but aggregation restores usable reliability; the anxiety signal tracks self-reported mood distress (*r* approx 0.22-0.25) while avoidance does not. We make no causal or general-population claims, and same-construct convergent validity (a concurrent ECR-RS questionnaire) is future work.

1. Methods and definitions

Because every figure in this report depends on how attachment is measured, we state the methodology first.

1.1 What counts as a measurement

A "measurement" is a single estimate of a person's attachment on two dimensions — anxiety and avoidance — each on a 0-1 scale. Measurements come from two distinct sources, which we keep separate throughout:

Intake assessment. a structured self-report taken once at onboarding (13,789 measurements). It yields each user's initial attachment point and self-described style.

Journaling-derived score. an estimate produced from a single journal entry (37,887 measurements). This is the novel measure and the basis for all change-over-time analysis. It is described next.

1.2 How a journaling-derived score is calculated (LLM-as-rater)

Each journal entry is scored by a large language model (Google Gemini 3 Flash) instructed to act as a clinical rater and assign attachment anxiety and avoidance on the 1-7 scale of the Experiences in Close Relationships-Revised (ECR-R) framework. The prompt directs the model to score enacted behavior and emotional regulation in the entry rather than topic words or theory labels — self-soothing, repair attempts, and grounded reflection lower the scores; unresolved spiraling, protest, or shutdown raise them. To reduce single-sample variance, each entry is scored three times independently (temperature 0.2) and the three ratings are averaged (a self-consistency ensemble). The averaged 1-7 values are linearly rescaled to 0-1 via $(x - 1) / 6$, so an ECR-R rating of 1 maps to 0.0 and 7 maps to 1.0. The model also classifies the entry's relationship domain (romantic, friendship, family, work, or self). A "journaling score" therefore is this 0-1 anxiety/avoidance pair for one entry; an "anxiety mean" or "avoidance mean" is the average of these 0-1 scores across a defined set of entries or people.

1.3 Current pattern (the snapshot)

A person's "current pattern" is an exponentially-weighted moving average of their journaling scores with a 30-day half-life — recent entries count more, older entries decay. This smooths the noisy per-entry signal into a running estimate of present-day anxiety and avoidance.

1.4 Self-report measures (separate from the attachment score)

Three optional in-app self-reports are used only to triangulate outcomes, never to compute the attachment score:

Mood check-in. users place how they feel on a valence (pleasant / unpleasant) by arousal (high / low energy) grid. "Mood distress" denotes unpleasant-valence check-ins.

7-day check-in. a brief self-report asking whether, after a week, anxiety, relationships, and mood "Improved" or showed "No change" (binary per item).

30-day check-in. includes a 5-point agreement item, "Attached has helped me feel less anxious" (Strongly agree to Strongly disagree).

1.5 Why the attachment score is not based on self-report

The self-report items above are coarse (binary or single-item), measure general wellbeing rather than attachment specifically, and are subject to recall and demand effects. The primary measure is instead the continuous, per-entry journaling score, which is collected passively at higher frequency and maps onto an established attachment instrument (ECR-R). We use the self-reports only as external, cross-construct triangulation (Section 4).

1.6 Statistical definitions

0.5 midpoint. categories are formed by splitting each 0-1 dimension at 0.5, which corresponds to the ECR-R scale midpoint (a raw rating of 4). Above 0.5 = "high" on that dimension. We report category proportions across 0.4-0.6 because this cut is arbitrary.

Style quadrants. secure (low anxiety, low avoidance), anxious (high anxiety, low avoidance), avoidant (low anxiety, high avoidance), fearful-avoidant (high on both).

Insecure to secure. the share of users whose quadrant moved from any insecure pattern (anxious, avoidant, or fearful-avoidant) at first measurement to secure (both dimensions below 0.5) at last measurement.

Cohen's d. a standardized effect size expressing a change as a multiple of its standard deviation (about 0.2 small, 0.5 medium, 0.8 large).

Reliable Change Index (RCI). flags an individual change as "reliable" only if it exceeds measurement error, i.e. 1.96 times the standard error of the difference, derived from the measure's test-retest reliability.

Test-retest reliability (r). reported as the Pearson correlation between two measurements close in time, the conventional reliability coefficient (we report r , not r -squared, because r directly indexes rank-order consistency). We estimate it over a 12-hour-to-7-day window, where genuine trait change should be minimal, so the coefficient reflects measurement consistency rather than real change — the same stability that, at longer horizons, lets us read movement as signal.

Confidence intervals. all reported as 95%.

1.7 AI assessment disclosure and limitations

All journaling-derived attachment scores are generated by an automated large language model (Google Gemini 3 Flash), not by human clinicians. Following emerging publishing norms (e.g. ICMJE and COPE guidance), we disclose this AI use explicitly and do not list the model as an author: it is a measurement instrument, not a contributor. AI rating carries specific risks — sensitivity to writing style, verbosity, and language; possible conflation of general negativity with attachment anxiety; and the absence of validation against human ECR-R coding. These are mitigated by the ECR-R anchoring and the three-sample ensemble, but not eliminated, and are listed again in Limitations (Section 5).

1.8 Privacy

All outputs are aggregate and de-identified. No journal text, identifiers, or row-level records are used. Subgroup breakdowns are reported only for groups of 500+ users, with category cells under 30 individuals suppressed.

2. Sample

Participants are users who adopted Attached specifically to work on their attachment and relationship patterns. This self-selection is the most important interpretive caveat: a help-seeking group is far more insecure than the general public, and nothing here should be read as a population prevalence rate. Geography skews English-speaking (about 53% United States, then the United Kingdom, Canada, and Australia; these four account for roughly three-quarters of users), across 100+ countries.

Data asset	Count
Users with a current attachment estimate	11,793
Total attachment measurements	51,676
- from the intake assessment	13,789
- from journaling sessions	37,887
Users with at least one journaling-based score	5,947

Table 1. Data assets underlying the report.

2.1 The app and feature engagement

Attached delivers a daily curriculum of short activities. Beyond guided journaling, members complete interactive lessons, guided and AI-generated meditations, and structured emotion-regulation tools - a "loop log" (attachment-loop exercise), cognitive reframes, "secure scripts," "let it go," trigger cards, and mood check-ins - plus affirmations, a relationship "bonds" map, and an AI chat companion. Journaling is the substrate for the attachment score (it is what gets scored); the remaining activities constitute the practice whose volume defines the dose-response in Section 3.6. Active members keep a near-daily practice: the typical member active at least a month journals about twice a week and completes roughly nine guided activities a week. For the features with enough returning users (members who use it across 4+ weeks) to estimate reliably, the table shows how often those members use each one and its overall reach, sorted by weekly cadence:

Feature	Times/week (returning users)	Used by
Journaling (the scored input)	2.2	83%
Cognitive reframes	2.1	13%
Interactive lessons	1.9	92%
Affirmations	1.7	26%
Meditations	0.8	23%
'Let it go' (thought release)	0.6	28%
Secure scripts	0.6	16%

Table 2. Weekly engagement cadence for well-sampled features (200+ returning users), sorted by frequency. 'Times/week' = median uses per active week among members who use that feature across 4+ weeks (its returning users); 'Used by' = share of active members who used it since May 2025. Lower-reach features with too few returning users to estimate reliably (trigger cards, bonds, AI chat, loop log) are omitted. Relative figures, no absolute volumes.

These activities are the app's behavioral utility. The dose-response analysis (Section 3.6) shows the most engaged tier completed about 356 activities on average versus about 22 in the least-engaged tier, and that heavier engagement tracks larger declines in attachment anxiety - an association, not a proven cause.

3. Results

3.1 The two dimensions (primary)

On the 0-1 ECR-R scale (Section 1.2), attachment anxiety runs high across the community (mean 0.74, median 0.78), with 90% of users above the 0.5 midpoint. Avoidance is lower and more dispersed (mean

0.46). We report the continuous dimensions first because they are stable, whereas category labels depend on the cut point (Section 3.2).

Cohort	Users	Anxiety mean (SD)	Avoid. mean (SD)	% anx ≥ 0.5	% avd ≥ 0.5
All current snapshots	11,793	0.74 (0.19)	0.46 (0.22)	90.4%	43.1%
1+ journaling score	5,947	0.76 (0.17)	0.45 (0.21)	91.5%	38.0%
5+ journaling scores	1,312	0.73 (0.18)	0.43 (0.20)	90.7%	31.9%

Table 3. Continuous dimension summaries by cohort (0-1 scale; 0 = ECR-R 1, 1 = ECR-R 7).

3.2 The four styles, and why the cut point matters

Splitting each dimension at the 0.5 midpoint (ECR-R raw 4) yields the four quadrants defined in Section 1.6. Because the categories depend on where that line falls, we report a sensitivity analysis across 0.4-0.6: the secure share alone moves from 2.6% to 14.3%. Treat the categories as a rough map, not a measurement.

Threshold	Anxious	Fearful-avoidant	Avoidant	Secure
0.4	38.7%	55.7%	3.0%	2.6%
0.5 (primary)	51.3%	39.1%	4.0%	5.6%
0.6	57.9%	22.7%	5.1%	14.3%

Table 4. Category distribution at three thresholds (all current snapshots, $n=11,793$).

3.3 Self-described vs. journaling-derived pattern

Here we compare the two measurement sources from Section 1.1. At intake (self-report) users describe themselves as even more anxious-leaning than their journaling-derived signal suggests, and very few self-identify as secure. The two sources agree only moderately (anxiety correlation $r = 0.36$; same quadrant 55%), a reminder that a single measurement of either kind is noisy.

Pattern	Self-described (intake)	Journaling-derived (0.5)
Anxious	63.5%	51.3%
Fearful-avoidant	12.2%	39.1%
Avoidant	4.8%	4.0%
Secure	1.9%	5.6%
Not sure	17.6%	n/a

Table 5. Intake self-report vs. journaling-derived signal. Categories are approximate.

3.4 By relationship situation

Using the journaling-derived current pattern (Section 1.3), people in a breakup show the highest anxiety; single users are the only group where fearful-avoidant outnumbers anxious. Cross-sectional associations, not causes.

Situation	Anxious	Fearful-avoidant	Secure	Avoidant
Breakup	55.0%	34.9%	6.9%	3.2%
Long-term relationship	51.8%	36.9%	6.3%	5.0%
New relationship	53.8%	40.6%	3.2%	2.4%
Dating	47.9%	45.1%	4.4%	2.6%
Single	38.4%	48.3%	6.5%	6.8%

Table 6. Current pattern by relationship status (% within group).

3.5 Change over time

OBSERVATIONAL FINDING

All changes below are in 0-1 ECR-R units (Section 1.2); for reference, 0.15 is about one ECR-R scale point ($0.15 \times 6 \approx 0.9$). To distinguish real change from day-to-day noise, the primary analysis averages each person's first six and last six journaling scores (single scores are too noisy; see Section 4). Anxiety declines consistently; avoidance barely moves.

Aggregation	Users	Anxiety change [95% CI]	Reliable decrease	Insecure to secure
First/last 3	684	-0.162 [-0.176, -0.147]	28.9%	21.1%
First/last 6 (primary)	549	-0.152 [-0.166, -0.139]	37.5%	20.2%
First/last 10	411	-0.149 [-0.163, -0.136]	47.4%	21.4%

Table 7. Aggregated within-person change (0-1 units). "Reliable decrease" = RCI-flagged anxiety drop (Section 1.6); "insecure to secure" defined in Section 1.6.

Longitudinal cohort	Users	Anxiety change [95% CI]	Style changed	Ins. to sec.
2+ scores	2,814	-0.136 [-0.144, -0.127]	42.6%	12.9%
5+ scores	1,312	-0.185 [-0.198, -0.171]	50.2%	17.1%
5+ scores, 28+ days	705	-0.228 [-0.247, -0.209]	52.3%	21.4%
10+ scores, 56+ days	446	-0.265 [-0.289, -0.241]	54.5%	24.4%

Table 8. Change by longitudinal cohort (first vs. last measurement, 0-1 units).

3.6 Dose-response

More practice is associated with larger anxiety declines, with effect sizes reaching Cohen's $d = 0.82$ (large; defined in Section 1.6). "Ins. to sec." is the insecure-to-secure share (Section 1.6). Heavier users also started more anxious, so this pattern partly reflects who chooses to engage and is not a clean causal contrast.

Journaling sessions	Users	Baseline anxiety	Anxiety change	Cohen's d	Ins. to sec.
2-4	50	0.80	-0.07	—	12.0%
5-9	111	0.84	-0.144	-0.28	14.4%
10-19	183	0.85	-0.189	-0.44	17.5%

Journaling sessions	Users	Baseline anxiety	Anxiety change	Cohen's d	Ins. to sec.
20-39	218	0.86	-0.255	-0.74	22.9%
40+	193	0.87	-0.281	-0.82	27.5%

Table 9. Dose-response in anxiety change (0-1 units). Cohen's d vs. the least-active (2-4) group.

3.7 Reported mood by attachment pattern

As an independent corroboration, we compared each current pattern against users' in-app mood check-ins (valence x arousal; Section 1.4). The result matches theory: users whose current pattern is secure report pleasant moods about 69% of the time, versus roughly 55% among anxious and fearful-avoidant users, and they report high-arousal distress ("anxious / agitated") least often. This is a concurrent, cross-sectional association in a self-reported, help-seeking sample - supporting evidence that the journaling-derived patterns track lived emotional states, not a causal claim. The avoidant subgroup is small (about 100 users), so its estimates are noisier.

Current pattern	Calm / content	Energized	Down / depleted	Anxious / agitated	Pleasant total
Secure	51.3%	17.3%	20.2%	11.2%	68.6%
Avoidant	42.5%	16.9%	26.0%	14.6%	59.4%
Fearful-avoidant	43.1%	12.4%	30.9%	13.6%	55.5%
Anxious	41.2%	13.9%	30.0%	15.0%	55.1%

Table 10. Mood check-in distribution by current attachment pattern (% of that pattern's check-ins; valence x arousal quadrants).

4. Measurement quality

4.1 Reliability and the role of aggregation

A single journaling score is noisy: two scores roughly two days apart correlate at only $r = 0.29$ (anxiety) and 0.33 (avoidance), consistent with each entry reflecting that day's context. Averaging restores reliability: by the Spearman-Brown projection, averaging six scores reaches r approx 0.71 and ten reaches r approx 0.80 . Every change figure in Section 3.5 is therefore built on averaged scores, not single entries.

Scores averaged	1	3	6	10	12
Projected anxiety reliability	0.29	0.55	0.71	0.80	0.83
Projected avoidance reliability	0.33	0.60	0.75	0.83	0.86

Table 11. Spearman-Brown reliability projection by number of averaged scores.

4.2 Snapshot stability (not reliability)

The running current-state estimate (the EMA, Section 1.3) is very smooth between adjacent readings. We do not report this smoothness as a reliability coefficient: adjacent snapshots are computed from overlapping entries, so their near-perfect correlation reflects autocorrelation by construction, not independent-error precision. The honest reliability ceiling for current-state estimates is the aggregation projection in Table 11.

4.3 Validity

We treat validity cautiously. At the user-month level, the journaling anxiety signal correlates with self-reported mood distress — the share of unpleasant-valence mood check-ins (Section 1.4) — at approx 0.22-0.25, strengthening as more data accrues; here r is a Pearson correlation, where roughly 0.1 is small, 0.3 moderate. The avoidance signal does not correlate with mood distress (r approx 0), the divergence attachment theory predicts. This is encouraging but cross-construct: mood is not attachment, so it is supporting evidence, not convergent validity. A same-construct test — a concurrent ECR-RS questionnaire — is the key item on our roadmap; until then we do not claim this is a validated attachment instrument.

User-month cohort	User-months	Anxiety vs. mood distress	Avoidance vs. mood distress
1+ journaling score / month	5,467	$r = 0.22$	$r = -0.03$
3+ journaling scores / month	2,457	$r = 0.25$	$r = 0.01$
6+ journaling scores / month	1,426	$r = 0.25$	$r = 0.08$

Table 12. Concurrent (level-on-level) correlations with mood-distress check-ins.

5. Limitations

Selection bias. Participants chose an attachment app; the sample is help-seeking and not population-representative.

AI assessment. Attachment scores are produced by a large language model, not human raters, and have not been validated against human ECR-R coding; they may be sensitive to writing style, verbosity, and language.

Self-report limits. The triangulation self-reports are coarse and subject to recall and demand effects; they corroborate but cannot validate the score.

Context bias. Journaling concentrates around difficult moments, which can inflate apparent insecurity and overstate later improvement (regression to the mean).

No control group. Change is observational; we cannot separate the app from time, life events, or natural recovery.

Engagement confound. Heavier users started more anxious, so dose-response partly reflects who engages.

Category fragility. Four-style proportions depend on threshold choice; continuous dimensions are the primary unit.

6. What we claim, and what we do not

Supported: group-level continuous distributions and group-level change over time within this help-seeking community; the divergent behavior of the two dimensions. **Not claimed:** general-population prevalence rates; clinical diagnoses; any causal effect of the app on attachment; individual-level verdicts about a

person's "true" style; convergent validity with an established attachment scale.

7. Author contributions

G.U. (Gabriel Uribe): conceptualization, data curation and analysis, and writing of the original draft. L.G. (Luis Goicouria): proofreading, and suggestions on the enumeration and clarification of the methods (writing - review and editing). S.C. (Sophie Cheung): resources and supervision - as co-founder of Attached, established and maintained the product and dataset on which this research depends. A large language model was used as an analysis instrument (Section 1.7) and is not listed as an author.

8. Data availability and citation

Aggregate tables are available on request. No raw, row-level, or identifiable data are released. This report is for educational purposes and is not medical advice.

Cite as:

Uribe, G., Goicouria, L., and Cheung, S. (2026). *The State of Attachment 2026: what 11,793 people working on their attachment reveal about whether it can change*. Attached Research.
<https://www.attachedapp.com/research/state-of-attachment>

Scores generated with Google Gemini 3 Flash as an ECR-R rater (disclosed per Section 1.7). (c) 2026 SkyPorch LLC.