

EmoMirror: Negotiable Affective Mirroring for Students with Elevated Alexithymia

Jinghe Yu*

ftyujinghe@mail.scut.edu.cn
South China University of Technology
Guangzhou, Guangdong, China

Dong She*

ftdshe@mail.scut.edu.cn
South China University of Technology
Guangzhou, Guangdong, China

Xinyao Li

202464870199@mail.scut.edu.cn
South China University of Technology
Guangzhou, Guangdong, China

Fanbo Meng

202464870243@mail.scut.edu.cn
South China University of Technology
Guangzhou, Guangdong, China

Yang Gao

gaoyang@cs.ecnu.edu.cn
South China University of Technology
Guangzhou, Guangdong, China

Zhanpeng Jin[†]

zjin@scut.edu.cn
South China University of Technology
Guangzhou, Guangdong, China

Abstract

Expressive writing can support mental health reflection, but it assumes that writers can recognize and name what they feel. Students with elevated alexithymia tendencies often face emotion vocabulary poverty: they experience distress yet struggle to distinguish, express, and revisit affective states. We examine how affective journaling systems can support such writers without turning uncertain model inference into an emotional verdict. We conducted a formative study with nine university students with elevated alexithymia tendencies, identifying three requirements: low-friction movement from capture to reflection, AI feedback as negotiable hypotheses, and user control over feedback timing and intensity. We introduce EmoMirror, a journaling prototype that renders model-inferred affect through kinetic typography, candidate emotion labels, valence-arousal fine-tuning, saved records, body-sensation logging, and user-triggered reflection. A within-session evaluation with eight students provides preliminary design insights. Participants used editable labels to move from vague affect toward more specific emotion language, engaged more readily in cause-focused reflection than hidden self-disclosure, and treated affective mismatch as useful only when they could revise, delay, or reject cues. EmoMirror shows how real-time affective typography can operate as negotiable scaffolding for emotional sense-making while preserving user interpretive authority.

CCS Concepts

• **Human-centered computing** → **Human computer interaction (HCI)**; *Empirical studies in HCI*; *Interaction design*; • **Applied computing** → *Health informatics*.

1 Introduction

Emotion journaling can support reflection only when writers have words for what they feel. Expressive writing and journaling help people organize emotional experiences and make sense of stressful events [8, 9]. Yet these practices assume that writers can recognize, name, and articulate their emotions. This assumption breaks down for students with elevated alexithymia tendencies, who often experience distress as an undifferentiated sense of being unwell.

Alexithymia describes difficulties identifying and describing feelings [1]; in everyday journaling, these difficulties appear as emotion vocabulary poverty, uncertainty about bodily and affective cues, and friction when turning felt experience into language. This paper examines how an affective journaling system can support these students without treating them as clinical patients or replacing their own interpretation.

Prior work gives mental health journaling systems several building blocks, but leaves open how affective feedback should enter the act of writing. Mental wellness self-tracking studies show that personal records can help users notice emotional patterns, while also raising concerns about burden, interpretation, and data sensitivity [2]. Recent AI journaling systems use large language models, contextual prompts, and multimodal records to support reflection, including systems such as MindfulDiary, DiaryMate, and MindScape [4, 5, 7]. Separately, kinetic and affective typography research shows that animated or adaptive text can communicate emotional qualities through motion, color, scale, and font form [6, 11, 12]. However, much of this work focuses on post-hoc summaries, conversational guidance, reader perception, or system-generated interpretations. These studies leave open how model-inferred emotion should become a real-time, visible, and revisable cue during first-person emotional writing.

The central design challenge is to make uncertain affective inference useful without making it appear certain. Mental health applications can cause harm when they present uncertain outputs as authoritative judgments [3]. This risk matters for alexithymic users because they may want help naming feelings while still needing authority over their own experience. We therefore frame affective feedback as a negotiable hypothesis: a cue the writer can inspect, question, revise, delay, or ignore. This stance aligns with HCI work on ambiguity as a resource, where open-ended representations can support reflection when users receive scaffolds for interpretation [10].

We first conducted a formative study with nine university students who reported elevated alexithymia tendencies. The study identified three design requirements for affective journaling. First, systems should lower the threshold for moving from immediate expression to later reflection. Second, AI feedback should remain a negotiable hypothesis rather than an emotional verdict. Third, the timing and intensity of affective feedback should remain under user control. These findings informed EmoMirror as a reflective

*These authors contributed equally to this work.

[†]Corresponding author. Email: zjin@scut.edu.cn

journaling prototype rather than an intervention that claims to improve emotion regulation.

We introduce EmoMirror, a negotiable affective mirror for journaling. EmoMirror lets writers capture typed or dictated fragments, renders model-inferred affect through real-time kinetic typography, presents candidate emotion labels, supports optional valence-arousal fine-tuning, saves records, and offers user-triggered reflection. The system turns affective inference into visible and revisable cues: the writer sees a possible interpretation, compares it with their own experience, and decides whether to keep, revise, hide, or revisit it. EmoMirror therefore treats dynamic typography as an interactional scaffold for emotional sense-making while leaving emotional certainty outside the system's scope.

We evaluated EmoMirror in a short within-session mixed-method prototype study with eight university students who reported TAS-20 scores above 52. We report descriptive statistics and interview findings, not efficacy claims. The preliminary evaluation suggests that EmoMirror helped some participants move from vague affect toward more specific emotion labels, prompted cause-focused reflection more consistently than hidden self-disclosure, and made mismatch useful only when participants could revise or reject the cue. We contribute three claims. First, students with elevated alexithymia tendencies need AI-mediated affective feedback they can negotiate. Second, EmoMirror operationalizes affective inference as visible and revisable scaffolding. Third, interpretive control, feedback pacing, and non-diagnostic language are first-order requirements for real-time affective journaling. We next describe the formative study, system design, prototype evaluation, and discussion.

2 Formative Study

We conducted a formative study to understand how university students with elevated alexithymia tendencies express and revisit emotions, and how they negotiate AI-mediated affective feedback. The study informed the design requirements of *EmoMirror*; it did not evaluate whether the system improves emotion awareness or regulation.

2.1 Participants and Procedure

We recruited nine university students through campus channels. As shown in Table 1, participants were 19–22 years old and reported Toronto Alexithymia Scale-20 (TAS-20) scores ranging from 53 to 68. We used the TAS-20 only to describe the recruited sample, rather than to compare participant groups or explain individual experiences [1].

Each participant completed a semi-structured interview lasting approximately 30 minutes, conducted either face-to-face or through video conferencing. The interviews examined participants' everyday emotion-expression practices, difficulties in identifying and describing emotions, journaling practices, reactions to a hypothetical AI-assisted journaling tool, and preferences for visual affective feedback. We introduced *EmoMirror* as a non-diagnostic journaling system that could reflect possible emotional cues through affective typography, emotion prompts, and optional post-writing reflection. Participants also completed a questionnaire that provided descriptive evidence regarding journaling barriers, feedback preferences, privacy concerns, and desired system features.

Table 1: Participants in the formative study. All participants were university students.

Participant	TAS-20	Gender	Age	Major
P01	59	Male	20	Electronic Information
P02	53	Female	20	Artificial Intelligence
P03	68	Female	19	Civil Engineering
P04	63	Female	19	Electrical Engineering
P05	61	Female	20	International Economics and Trade
P06	60	Female	19	Artificial Intelligence
P07	60	Female	21	Design
P08	65	Female	22	Intelligent Vehicles
P09	61	Female	21	Primary Education

We analyzed anonymized structured interview records using reflexive thematic analysis [?]. We first familiarized ourselves with each record and generated initial codes from participants' recorded responses. We then iteratively consolidated the codes, retained contradictory cases, and developed candidate themes concerning emotion expression, AI feedback, and design requirements. Because the available corpus consisted of structured interview records rather than verbatim transcripts, our analysis focuses on explicitly documented accounts and their design implications. All quotations below were translated from the original Chinese interview records by the authors.

2.2 Findings

2.2.1 Emotional Journaling Needs a Low-Friction Bridge from Expression to Reflection. Participants did not lack ways to express emotion. Instead, they often relied on socially situated and immediate channels, such as talking with friends, venting, or posting a brief status update. As one participant explained, "Chatting with friends feels the most natural" (P08). However, these channels did not always create a private and durable record that participants could revisit. At the same time, journaling was not perceived as a frictionless alternative. Participants described lacking time, motivation, or an established journaling practice. P04 stated that "laziness is the main reason" and identified assistance with starting a record as the most valuable form of support.

Writing also introduced a tension between immediate capture and subsequent reflection. Some participants felt that formal writing was too slow or overly organized to preserve fleeting feelings. P08 noted that focusing on making a record coherent could make it feel "less real." Yet, once emotions had been recorded, revisiting them could create useful psychological distance. P05 described rereading a journal entry "from a third-person perspective" as helping them view their feelings more objectively.

These accounts suggest that *EmoMirror* should not assume that users are already committed journal writers. Instead, it should bridge fleeting expression and later reflection through lightweight entry points, support for informal or multimodal expression, and mechanisms that preserve records for subsequent review.

2.2.2 AI Feedback Should Remain a Negotiable Hypothesis, Not an Emotional Verdict. Participants frequently described a gap between feeling distressed and understanding what that distress meant. P08

recalled feeling low but being unable to determine “what type of emotion” it was. Participants therefore wanted AI feedback to provide more than a positive-versus-negative classification. They requested cues regarding emotion type, intensity, and possible causes. P02, for example, wanted to know “how distressed” the system inferred they were and “why” it had made that inference.

However, participants did not want the system to convert complex experiences into definitive labels. P01 cautioned that emotions are inherently complex and questioned why they should be reduced to a small number of fixed words. Similarly, participants resisted systems that would immediately replace a vague expression with a supposedly more accurate emotion term. Instead, they wanted feedback to remain open to interpretation. P07 asked the system to explain “which sentence” had led it to infer a particular emotion.

Thus, the central value of AI feedback lies not in claiming emotional certainty, but in opening an inquiry. *EmoMirror* should present affective outputs as provisional cues, expose the basis of its inferences where possible, and allow users to accept, ignore, question, or revise them. This framing preserves users’ authority over their own emotional interpretations while supporting reflection on feelings that may be difficult to articulate.

2.2.3 Affective Feedback Becomes Supportive Only When Users Can Negotiate Its Pace and Intensity. Participants responded to dynamic visual feedback in sharply different ways. Some anticipated that visible affective changes could interrupt writing or intensify negative emotions. P02 explained that real-time feedback might cause them to stop writing because it could create a sense of “losing control.” Others preferred post-writing feedback because immediate prompts might redirect their thoughts before they had fully expressed themselves.

In contrast, several participants viewed visual affective feedback as an opportunity for emotional release and deeper expression. P07 preferred receiving feedback after completing a sentence or short passage, while P08 described dynamic expression as resembling a friend who “helps speak out” their emotions. For this participant, visual amplification could make a feeling more manageable rather than more intense.

These contrasting responses indicate that affective visualization cannot function as a uniform intervention. Its usefulness depends on the writer’s momentary readiness for feedback, the intensity of that feedback, and whether the user can control when it appears. *EmoMirror* should therefore support multiple feedback rhythms, including on-demand, segment-level, post-writing, and real-time modes. It should also provide low-arousal visual feedback, clear pause and disable controls, and supportive language during negative emotional states.

Across these findings, trust emerged as a cross-cutting requirement. Participants could describe AI as a companion or reflective mirror, but they also worried about being monitored or judged. P02 stated that the idea of “an AI knowing all of my thoughts” felt frightening, whereas P09 described a supportive system as “more like companionship.” *EmoMirror* should therefore prioritize private and, where feasible, local interaction; non-diagnostic language; user-controlled feedback; and transparent boundaries around emotional inference.

2.3 Design Implications for EmoMirror

Our findings translate into three design requirements. First, *EmoMirror* should lower the threshold for moving from everyday expression to reflective journaling through lightweight and flexible input, together with support for later review. Second, it should frame emotion recognition as a negotiable hypothesis by exposing relevant cues, communicating uncertainty, and preserving opportunities for user correction. Third, it should treat feedback timing and intensity as configurable interaction parameters rather than fixed system behaviors. These requirements motivate *EmoMirror* as a private, user-paced affective journaling system whose effectiveness will be examined in a subsequent user study.

3 EmoMirror

3.1 Overview of EmoMirror

EmoMirror operationalizes the formative findings as a negotiable affective mirror for journaling. Our formative study identified three requirements: writers need a low-friction bridge from immediate expression to later reflection, AI feedback should remain a negotiable hypothesis, and affective feedback should follow the writer’s pace. *EmoMirror* responds through a four-step loop: writers capture a typed or dictated fragment, encounter provisional affective cues, negotiate the interpretation, and revisit saved entries when they are ready for deeper reflection. Figure 1 shows the user-facing surfaces that instantiate this loop, and Figure 2 shows the implementation boundary that translates user input into revisable cues.

EmoMirror makes the writer’s interpretation the authority for each affective cue. The prototype includes a Journal interface for lightweight capture, kinetic affective typography for visible cues, candidate emotion labels, an optional valence–arousal fine-tuning surface, saved records, and user-triggered reflection views. These features form one workflow across capture, feedback, revision, and review. The design therefore treats affective feedback as material for sense-making: the system offers a cue, and the writer decides how to revise, save, ignore, or revisit it.

3.2 Key Features of EmoMirror

3.2.1 Low-Friction Capture Preserves Fleeting Expression for Later Reflection. Journal separates emotional capture from formal diary writing. Participants described formal journaling as valuable after the fact, but difficult to start during fleeting or messy emotional moments. *EmoMirror* therefore lets writers enter short text fragments or use browser-based speech dictation before they organize an experience into a coherent account. The browser transcribes speech, and *EmoMirror* processes the resulting text through the same affective cue pipeline used for typed entries.

Each Journal fragment becomes a saved record that writers can revisit later. This design preserves the immediacy of unfinished expression while still creating a durable path toward reflection. Diary functions as a later organizing space, where writers can return to saved material and request structured prompts about events, feelings, possible needs, and reflection questions. *EmoMirror* therefore turns journaling into a staged workflow: capture first, organize later.




Figure 1: User-facing surfaces of EmoMirror. (1) Journal pairs expressive writing with a live kinetic mirror, candidate labels, and V-A adjustment so writers can inspect and revise model cues during capture. (2) Body Sense records bodily sensations with recent journal context as optional reflection material. (3) Emo Echo provides user-triggered conversational reflection within a non-diagnostic framing.




Figure 2: Technical implementation of EmoMirror. The frontend handles journal views, client state, media input, and character-level typography rendering. A stable JSON API contract carries emotion segments, V-A mappings, design maps, and status metadata between the browser and backend services. The backend orchestrates FastAPI routing, authentication, storage, model clients, safety controls, and fallback behavior.

3.2.2 Kinetic Affective Typography Externalizes Cues While Preserving the Writer’s Words. Kinetic affective typography makes model-inferred affective cues visible while preserving the writer’s original text. As writers type or dictate an entry, EmoMirror analyzes the current text and renders provisional cues through motion, scale, color, weight, and local emphasis. Motion and scale express energy; color and weight suggest affective valence or tension. The text remains the writer’s own sentence, while the rendering creates a visible cue that the writer can inspect. At the implementation level, the backend returns a character-indexed design map, and the browser applies visual styles to each character (Figure 2).

This feature responds to participants who wanted more than positive-negative sentiment but resisted definitive emotion labels. EmoMirror uses typography to make a possible interpretation perceptible before turning it into a fixed word. Prior work shows that kinetic typography and adaptive font rendering can communicate affective qualities through animated text [6, 11, 12]. EmoMirror

redirects that design space toward first-person reflection by asking the writer to compare the rendered cue with their own account.

3.2.3 Candidate Labels and Optional V-A Fine-Tuning Keep Interpretation Negotiable. Candidate labels turn affect labeling into a revision process. After generating a provisional cue, EmoMirror surfaces a primary emotion label and nearby candidate labels. These candidates give writers a small vocabulary for moving between related feelings, such as low-energy negative states and high-energy negative states. Writers can select a candidate label, keep the current label, or revise the cue through the optional valence-arousal surface.

The V-A surface serves as an optional fine-tuning control that remains secondary to journaling. The interface presents the horizontal axis as negative-to-positive and the vertical axis as low-to-high energy, avoiding technical psychological terminology in the writer-facing view. Dragging the point updates the selected

label and candidate labels; selecting a label updates the point. The saved entry records the writer's negotiated interpretation as the authoritative account for that entry.

This bidirectional revision loop preserves interpretive authority. Participants wanted AI feedback to open inquiry while leaving room to disagree, especially when emotions felt mixed or difficult to name. Candidate labels provide natural-language scaffolding, while V-A fine-tuning supports writers who want more control over nuance. This stance aligns with HCI work that treats ambiguity as productive when users receive scaffolds for interpretation [10].

3.2.4 User-Triggered Reflection Lets Writers Revisit Cues Across Time. EmoMirror distributes reflection across immediate, entry-level, and longitudinal time scales. During writing, Journal offers lightweight cues through dynamic typography, candidate labels, and optional V-A fine-tuning. After an entry, Diary reflection lets writers request a structured account of events, feelings, possible triggers, and reflection questions. Across multiple saved entries, Emotion Review summarizes recent Journal, Diary, and body-related records so writers can inspect patterns over time.

This temporal structure responds to participants' conflicting preferences about feedback timing. Some participants expected immediate visual feedback to support expression; others worried that it would interrupt writing or amplify negative affect. EmoMirror therefore keeps deeper reflection user-triggered. Body sensation recording and Records/history support later review as contextual resources, while the core interaction remains the writer-paced loop from capture to revisiting.

3.3 Safety and Data Boundaries

EmoMirror constrains affective mirroring through non-diagnostic language, user-triggered reflection, and scoped access to personal records. The interface presents affective outputs as possible cues for writer interpretation. When body-related or risk-related content appears, the system uses supportive language and directs users toward professional help when appropriate. These safeguards keep the prototype within reflective journaling and outside diagnosis, treatment, or crisis intervention.

EmoMirror also limits how personal records become visible across roles. Ordinary views return records associated with the current participant code, while research-facing access requires separate permission. Deeper AI-generated reflection appears only after the writer requests it. These boundaries answer the formative study's trust concern: affective journaling becomes threatening when AI appears to monitor, judge, or overclaim. EmoMirror instead positions the system as a revisable mirror whose authority remains subordinate to the writer's interpretation, responding to documented risks in mental health applications that present uncertain outputs as authoritative judgments [3].

4 User Study and Results

We conducted a short, within-session, mixed-method prototype evaluation to understand how participants experienced EmoMirror as a negotiable affective mirror. This evaluation connects the formative study and system design: the formative study identified the need for low-friction capture, provisional affective cues, and

user-controlled feedback; EmoMirror operationalized these requirements through kinetic affective typography, candidate labels, V-A fine-tuning, and user-triggered reflection. The study examined participants' perceived support for emotion naming, deeper writing, and interpretive control during affective journaling.

4.1 Participants and Procedure

We recruited eight university students with elevated alexithymia tendencies. We used TAS-20 scores above 52 as a screening descriptor for recruitment and sample characterization. Participants were 19–20 years old, and their TAS-20 scores ranged from 53 to 80 ($M = 62.50$, median = 60.50).

Each participant received a short introduction to EmoMirror as a reflective journaling prototype. Participants wrote about emotionally meaningful experiences, observed kinetic affective typography and provisional affective cues, and revised the system-generated cues through candidate labels or V-A fine-tuning when the output did not match their own interpretation. They also compared the experience with ordinary memo-style journaling. After using the prototype, participants completed a 1–5 Likert questionnaire and a brief semi-structured interview. The study focused on perceived short-term experience with the prototype.

4.2 Measures and Analysis

The questionnaire included items on three constructs: Emotion Naming Support, Self-Disclosure Support, and Mirror Mechanism Experience. We report descriptive results only, including means, medians, variation, and the number of participants selecting 4 or 5. We also report feedback timing preferences. We coded interview responses with a lightweight deductive codebook aligned with the three evaluation questions, including codes such as MoreSpecificLabel, MoreCause, UsefulMismatch, MisreadFeeling, Interruption, NeedControl, PrivacyConcern, and AccuracyLimit. We retained negative cases to identify design boundaries. All quotes were translated from Chinese by the authors.

4.3 Findings

Editable labels helped some participants move from vague affect to more specific emotion words. Questionnaire responses suggest moderate and uneven perceived support for emotion naming. The strongest item in this construct concerned differentiating vague or mixed feelings through candidate labels and the coordinate panel (Q3: $M = 3.38$, median = 4.00; 5/8 agreement). Interviews clarify why this support worked for some participants: P07 described the candidate labels as a way to recover from an imperfect first output, saying, "After I finished writing, it gave many related emotion words; even if the first one was not what I wanted, I could usually find a fitting one nearby." At the same time, several participants described the label space as too broad, too sparse, or too rigid for mixed emotions. **Takeaway.** EmoMirror's perceived value for emotion naming came from editable affective scaffolds, especially candidate labels and V-A fine-tuning, rather than from a single model-generated label.

EmoMirror prompted reflection on causes more than disclosure of hidden needs. Questionnaire responses were weaker for self-disclosure than for emotion naming. The item about causes or triggers had the clearest trend (Q6: $M = 3.50$, median = 4.00; 5/8 agreement), while the item about hidden needs, grievances, fears, expectations, or unspoken thoughts had weaker support (Q7: $M = 2.75$, median = 3.00; 2/8 agreement). P04 described the mechanism as a prompt to continue writing: “I would keep writing to analyze why I had this emotion.” Other participants treated novelty, privacy, and negative-affect amplification as boundaries. One participant described the system as playful but said it did not necessarily elicit content they would otherwise leave out; another described emotional records as private and reported holding back on a stranger platform. **Takeaway.** EmoMirror can invite cause-focused reflection in a short prototype session, while sensitive self-disclosure requires privacy, optional prompts, and user-paced reflection.

The affective mirror created both productive mismatch and perceived misreading. Participants described mismatched affective cues as useful friction for some moments and a risk in others. Four participants agreed that mismatch could help them reinterpret themselves (Q10), while three agreed that mismatch could make them feel misread or resistant (Q11). The clearest questionnaire signal concerned interpretive authority: seven of eight participants agreed that revising labels or coordinates was important (Q14), and seven agreed that EmoMirror felt more like a reference cue than a judgment of their inner state (Q15). Feedback timing preferences also varied: 6/8 preferred real-time feedback, while 2/8 preferred receiving feedback after completing the whole entry. P05 explained the risk of premature feedback: “If I have not finished the context, and the system quickly and arbitrarily labels an emotion, I would feel uncomfortable.” **Takeaway.** Mismatch becomes supportive when users can inspect, reject, revise, delay, or soften the cue; EmoMirror should keep affective feedback provisional so the mirror remains a reflective partner rather than an interruption.

Overall, the preliminary evaluation supports a bounded design claim. Participants valued EmoMirror when model-inferred affective cues became visible, editable, and open to disagreement. These responses reinforce the design stance from the formative study: affective journaling systems should present emotional inferences as revisable reflection cues that preserve the writer’s interpretive authority.

5 Discussion

This paper argues that affective journaling for students with elevated alexithymia should not make emotion recognition more authoritative; it should make uncertain affective cues easier to inspect, revise, and revisit. The formative study identified three requirements: low-friction capture, AI feedback as a negotiable hypothesis, and user control over timing and intensity. EmoMirror operationalizes them through lightweight text/speech capture, kinetic affective typography, candidate labels, optional V-A tuning, saved records, and user-triggered reflection. The prototype evaluation provides preliminary support: participants valued EmoMirror when it helped them move from vague affect toward specific language, reflect on causes, and disagree with mismatched cues without surrendering interpretive authority.

The workshop-level contribution is therefore a bounded design claim: negotiable affective mirroring can turn model-inferred emotion into reflective material, not emotional verdicts. For mental health journaling, this distinction matters because deeper disclosure, diagnostic language, or always-on feedback can threaten privacy and trust. EmoMirror suggests that interpretive control, pacing, revision, and non-diagnostic boundaries are core design requirements.

This study has clear limitations. The evaluation involved eight participants, a short within-session interaction, descriptive statistics, and perceived support rather than longitudinal outcomes. Future work should evaluate repeated use, compare real-time and post-writing feedback rhythms, support richer multi-label and segment-level revision, and examine whether negotiable mirroring shapes emotion awareness over time.

References

- [1] R. Michael Bagby, James D. A. Parker, and Graeme J. Taylor. 1994. The Twenty-Item Toronto Alexithymia Scale—I. Item Selection and Cross-Validation of the Factor Structure. *Journal of Psychosomatic Research* 38, 1 (1994), 23–32. doi:10.1016/0022-3999(94)90005-1
- [2] Eun Kyoung Choe, Bongshin Lee, Matthew Kay, Wanda Pratt, and Julie A. Kientz. 2017. Self-Tracking for Mental Wellness: Understanding Expert Perspectives and Student Experiences. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. 629–641. doi:10.1145/3025453.3025750
- [3] Rachael M. Kang and Tera L. Reynolds. 2024. “This App Said I Had Severe Depression, and Now I Don’t Know What to Do”: The Unintentional Harms of Mental Health Applications. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI ’24)*. Association for Computing Machinery, New York, NY, USA, Article 1003, 17 pages. doi:10.1145/3613904.3642178
- [4] Taewan Kim, Seolyeong Bae, Hyun Ah Kim, Su-Woo Lee, Hwajung Hong, Chanmo Yang, and Young-Ho Kim. 2024. MindfulDiary: Harnessing Large Language Model to Support Psychiatric Patients’ Journaling. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 1–20. doi:10.1145/3613904.3642937
- [5] Taewan Kim, Donghoon Shin, Young-Ho Kim, and Hwajung Hong. 2024. Diary-Mate: Understanding User Perceptions and Experience in Human-AI Collaboration for Personal Journaling. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 1–15. doi:10.1145/3613904.3642693
- [6] Joohwan Lee, Soojin Jun, Jodi Forlizzi, and Scott E. Hudson. 2006. Using Kinetic Typography to Convey Emotion in Text-Based Interpersonal Communication. In *Proceedings of the 6th Conference on Designing Interactive Systems (DIS ’06)*. Association for Computing Machinery, New York, NY, USA, 41–49. doi:10.1145/1142405.1142414
- [7] Subigya Nepal, Arvind Pillai, William Campbell, Talie Massachi, Eunsoo Soul Choi, Orson Xu, Joanna Kuc, Jeremy Huckins, Jason Holden, Colin Depp, Nicholas Jacobson, Mary Czerwinski, Eric Granholm, and Andrew T. Campbell. 2024. Contextual AI Journaling: Integrating LLM and Time Series Behavioral Sensing Technology to Promote Self-Reflection and Well-Being Using the MindScape App. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*. doi:10.1145/3613905.3650767
- [8] James W. Pennebaker. 1997. Writing about Emotional Experiences as a Therapeutic Process. *Psychological Science* 8, 3 (1997), 162–166. doi:10.1111/j.1467-9280.1997.tb00403.x
- [9] James W. Pennebaker and Cindy K. Chung. 2011. Expressive Writing: Connections to Physical and Mental Health. In *The Oxford Handbook of Health Psychology*. Oxford University Press. doi:10.1093/oxfordhb/9780195342819.013.0018
- [10] Pedro Sanches, Kristina Höök, Corina Sas, and Anna Ståhl. 2019. Ambiguity as a Resource to Inform Proto-Practices: The Case of Skin Conductance. *ACM Transactions on Computer-Human Interaction* 26, 4, Article 21 (2019), 32 pages. doi:10.1145/3318143
- [11] Dong She, Jinghe Yu, Nuo Chen, Yuting Bai, Yang Gao, and Zhanpeng Jin. 2026. EmoType-Zero: Adaptive Fonts to Render the Speech Emotion. In *Extended Abstracts of the 2026 CHI Conference on Human Factors in Computing Systems (CHI EA ’26)*. Association for Computing Machinery, New York, NY, USA, Article 254, 7 pages. doi:10.1145/3772363.3798656
- [12] Jocelyn J. Shen, Kathryn Jin, Ann Zhang, Cynthia Breazeal, and Hae Won Park. 2023. Affective Typography: The Effect of AI-Driven Font Design on Empathetic Story Reading. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems (CHI EA ’23)*. Association for Computing Machinery, New York, NY, USA, Article 26, 7 pages. doi:10.1145/3544549.3585625