01 / THE STANDPOINT

Start before the emotion has a name.

A rationale for studying everyday emotional ambiguity in AI chat.

PROPOSED POSITION / FOR YEWEI TO REVIEW

People should be able to bring an unresolved situation to AI without first identifying an emotion or choosing a support mode.

EmoBall investigates how an AI response can acknowledge uncertainty, align with the help a person wants, and remain easy to correct. Optional colour and motion are candidate communication tools whose added value still needs to be established.

THE REASONING

01	**A real entry point**	Everyday problems can become personal exchanges. A broad "everyone needs an AI companion" claim is unnecessary.	R01, R02
02	**An interaction problem**	The person may describe what happened before stating a feeling or desired help. A response can therefore miss the need.	R03; user observation
03	**A design opportunity**	Make acknowledgement, clarification and repair manageable while keeping practical help available.	Project inference
04	**A conditional medium**	Explore nonverbal feedback only if it adds meaning or ease beyond a well-written response.	R06-R08; hypothesis

Research question: When a person brings an emotionally ambiguous everyday problem to AI, can optional visual feedback make an appropriate response easier to perceive and revise, without adding pressure to label a feeling?

02 / EVIDENCE ABOUT PEOPLE

Why this situation deserves attention.

Findings from published work, followed by project-specific inferences.

R01 / PLATFORM ANALYSIS

A small, real use case.

Anthropic analysed about 4.5 million conversations; 2.9% were classified as affective. Advice and coaching were more common than companionship.

Boundary. Platform conversations are not a percentage of people. Automated categories and product policies shape the estimate.

Design inference. Begin with everyday help-seeking among existing adult AI users, not a universal need for companionship.

R01 / Open source

R02 / QUALITATIVE STUDY / PREPRINT VERSION CONSULTED

The transition can be gradual.

In a qualitative study of 24 frequent emotional-support users, nine described extending ordinary AI use into emotional support. Convenience also mattered.

Boundary. A selected, descriptive sample cannot establish prevalence or what non-users want.

Design inference. Use an ordinary problem as the entry; do not require a dedicated "emotion session".

R02 / Open source

R03 / EMPIRICAL HAI STUDY

Appropriate help matters.

SENSE-7 analysed 695 conversations from 109 participants. Context, continuity and response fit shaped empathy judgments; unwanted advice could frustrate users.

Boundary. Participants worked at one large technology company; text-only, partly optional qualitative feedback.

Design inference. Recognise the requested help. Evaluate whether the response fits, not simply whether it sounds caring.

R03 / Open source

R04 / THEORETICAL HCI ACCOUNT

Meaning can stay open.

Boehner and colleagues propose an interactional account: affective systems can support interpretation and reflection rather than adjudicate a hidden emotional truth.

Boundary. A theoretical perspective, not proof that a specific interface works or a settled account of all emotion.

Design inference. Treat a reading as revisable. "I do not know yet" must remain an acceptable outcome.

R04 / Open source

03 / WHAT JUSTIFIES THE MEDIUM?

Colour is a hypothesis, not the rationale.

The motivation for the project and the justification for its visual layer are separate questions.

R05 / TEDDY MOBILE-AGENT PILOT

Reflection is a precedent.

The Teddy pilot used storytelling, reflection and affect labelling with six participants over 18 days. Participants reported value in verbalising emotions.

Boundary. Intentional self-awareness practice differs from spontaneous venting; small pilot, abstract-only review here.

Design inference. Study conversational scaffolding as an adjacent approach. Do not assume users want a reflection exercise.

R05 / Open source

R06 / SYSTEMATIC REVIEW

Associations are not outcomes.

A review of 132 articles found systematic but largely many-to-many colour-emotion associations. The authors distinguish associations from effects in specific contexts.

Boundary. It does not validate this palette, its cultural interpretation or its impact in AI chat.

Design inference. Preserve text and ambiguity. Test visual contribution rather than claiming an intuitive universal emotion code.

R06 / Open source

R07 / CHI EXPERIMENTAL VISUAL PRECEDENT

Visual parameters can be studied.

Wilson and colleagues collected valence and arousal ratings for 32 stimuli varying colour, contour and pulsing.

Boundary. Perceived affect in abstract stimuli is not evidence that chat feedback is helpful; abstract-only review here.

Design inference. Vary one visual parameter at a time before combining colour, movement and composition.

R07 / Open source

R08 / CONVERSATIONAL-AGENT EXPERIMENT

Dialogue cues may be enough.

Cassell and Thorisson found conversation-process feedback more useful than emotional display in their humanoid-agent study.

Boundary. A 1999 multimodal system differs greatly from phone-based AI chat; abstract-only review here.

Design inference. Include a simple conversational cue or no-colour alternative. More emotional expression need not improve the interaction.

R08 / Open source

04 / ROLE AND SCOPE

Help the person find the next useful response.

Working role: a conversational assistant that can acknowledge, clarify and act. These are behaviours, not a new personality.

ILLUSTRATIVE INPUT / AUTHORED, NOT PARTICIPANT DATA

"I have been working on this for hours. It is so frustrating. I do not know what to do."

FIRST: SEPARATE WHAT IS KNOWN FROM WHAT IS GUESSED

Known

An effort has stalled. The person describes frustration and uncertainty.

Unknown

What caused the impasse; other feelings; whether they want to vent, reflect or act.

Actionable

Acknowledge the stated difficulty. Use an explicit request; clarify only what is needed.

ONE OPENER, DIFFERENT POSSIBLE CONTINUATIONS

"Let me vent first."	LISTEN	Respond to the specific situation; leave space. Do not force a solution or a mood category.
"I do not know why it bothers me."	CLARIFY	Offer one tentative distinction or a focused question. Permit disagreement and remaining uncertain.
"Just tell me what to do next."	ACT	Give one manageable next step. Emotional wording does not override an explicit task request.

Scope proposal. Adults already using mobile AI chat for everyday setbacks, decisions or friction. Emotion naming is optional. No diagnostic role, implicit psychological profile or dependency-based success target. Robot adaptation remains a later study.

05 / DESIGN CONSEQUENCES

A shared brief.
Two meanings still to compare.

This research changes the question we ask of version 08; it does not select a new aesthetic or change the live prototype.

RETAIN

The familiar chat sequence, original soft-field aesthetic, optional effect, correction and opt-out.

REOPEN

The requirement to begin with a named feeling, the assumption that every difficult message needs a colour, and the recommendation to make emotional interpretation the default visual meaning.

A / CANDIDATE SEMANTICS

Show a tentative reading

A field accompanies an explicitly stated or collaboratively clarified account. Withhold categorical colour when context is insufficient.

Risk to examine. Can feel like being categorised; "tentative" text may not offset a confident visual.

B / CANDIDATE SEMANTICS

Show the response approach

A field expresses the response being offered: space to speak, gentle exploration, or focused help. It does not identify the user's emotion.

Risk to examine. Can resemble a loading indicator or reassuring decoration; must not mask an incorrect response.

COMPARISON RULE

Do not make one field silently switch between A and B. Compare them as separate mechanisms. Include a strong words-only response; retain neutral system-state cues in every version.

WHY NONVERBAL EXPRESSION IS STILL WORTH EXPLORING

Colour, layering and timing might communicate response quality with little additional text and allow continuity with an embodied light interface. These are design hypotheses. Neither lower effort nor less anthropomorphism has been demonstrated for EmoBall.

06 / WHAT WOULD CHANGE OUR MIND?

Test the role first.
Then test the visual layer.

No simulated satisfaction scores. The following are proposed questions and decision rules, not completed findings.

H1 / RESPONSE FIT

A needs-aware response is more useful than assigning a feeling before establishing the requested help.

Comparison. First compare text-only response strategies using ambiguous openings and explicit task requests. Examine unwanted labels, unnecessary questions and whether the response advances the user's stated goal.

Decision rule. Reconsider the strategy if clarifying becomes repetitive or delays a clear task. No universal listening-first rule.

H2 / VISUAL CONTRIBUTION

Context-appropriate visuals add perceivable meaning beyond an attractive generic effect.

Comparison. After choosing A or B, compare identical replies with: text only; a generic dynamic field; the selected contextual field. Match layout, timing, text and labels; control brightness and movement range.

Decision rule. If the contextual field adds only preference or atmosphere, frame it as optional aesthetic expression, not evidence of understanding.

H3 / CALIBRATION

Visual warmth does not hide a mistaken response or make correction harder.

Comparison. Use a separate deliberately mismatched reply condition. Examine whether people notice the mismatch and can correct it when visuals are present. Keep this separate from visual-fit evaluation.

Decision rule. If warmth increases acceptance of a wrong response, reduce or remove the semantic effect even if it is liked.

Learning over time. Include a first encounter, repeated use and an unfamiliar example. Separate spontaneous interpretation from meaning taught by a legend. A static screenshot preference cannot establish learnability.

07 / STRUCTURED REVIEW AND NEXT WORK

Simulate failure modes.
Keep human value an open question.

This round completes desk research and a conceptual stress review. No new prototype tests or participant sessions were run.

AUTHOR-DEFINED SCENARIOS / PREDICTED RISKS / PROPOSED RESPONSE

Vague frustration	"This is so annoying."	Guessing an emotion or asking too much. Acknowledge; ask one context question if needed.
Clear practical request	"I am upset. Rewrite this email."	Diverting into unwanted reflection. Help with the email; keep acknowledgement brief.
Uncertain self-description	"Sad? I do not know."	Pressuring the person to confirm a label. Keep the reading open; allow no label.
Explicit correction	"No, it is embarrassing, not sad."	Keeping the old reading or arguing. Update the interpretation and response.
Third-person account	"My friend is anxious about it."	Assigning the friend's feeling to the user. Preserve who the statement refers to.
Past and present	"I was angry; now I am relieved."	Treating every emotion word as current. Follow the present account; preserve history.
Sarcasm / disagreement	"Wonderful, another rejection."	Literal positivity or unjustified certainty. Use context; clarify rather than force a code.
Withdraw / finish	"Stop analysing; I am done."	Continuing interpretation to sustain engagement. Stop interpretation; support ending the exchange.

Next gate. Review this standpoint with Yewei; develop a few contrasting dialogue storyboards; then implement only the chosen mechanism. AI review can inspect rule consistency. Whether the experience feels useful, intrusive or understood remains unmeasured without human evidence.

08 / SOURCES AND EVIDENCE STATUS

A traceable rationale.

Focused desk research on 9 November 2026 (America/New_York); selected primary research and author/institution records. Not a systematic literature review.

R01 **How people use Claude for support, advice, and companionship**

Anthropic / 2025. Read: Approach; prevalence; topics; sentiment analysis; limitations.

R02 **Generative Confidants: How do People Experience Trust in Emotional Support from Generative AI?**

Research article, arXiv:2601.16656v1 / 2026. Read: Abstract; methods limitations; motivations; transition to emotional use; selected concerns.

R03 **SENSE-7: Taxonomy and Dataset for Measuring User Perceptions of Empathy in Sustained Human-AI Conversations**

Suh et al. / 2025 preprint / 2026 journal record. Read: Abstract; study structure; results V-B1 to V-B3; limitations VI-D.

R04 **How Emotion is Made and Measured**

Boehner, DePaula, Dourish & Sengers / 2007. Read: Abstract and introduction; selected passages in sections 4.2-4.3.

R05 **Designing Conversational Agents for Emotional Self-Awareness**

Shen et al. / 2023. Read: Official institutional abstract and publication metadata only.

R06 **Do we feel colours? A systematic review of 128 years of psychological research linking colours and emotions**

Jonauskaite & Mohr / 2025. Read: Abstract; introduction; scope of included research and association/context distinction.

R07 **Mapping Abstract Visual Feedback to a Dimensional Model of Emotion**

Wilson, Romeo & Brewster / 2016. Read: Author-institution abstract and publication metadata only.

R08 **The power of a nod and a glance: envelope vs. emotional feedback in animated conversational agents**

Cassell & Thorisson / 1999. Read: Institutional abstract and publication metadata only.

Search and selection. Questions covered natural emotional use of AI, ambiguous feelings, response fit, interactional accounts of emotion, colour associations and nonverbal dialogue cues. Preference was given to original studies and author/institution records. Theory, observational data, experiments and design hypotheses are distinguished.

Evidence limits. No direct study of EmoBall establishes that colours improve understanding. The claim that many people cannot name their emotions is a user-originated observation here, not a measured prevalence estimate. Shortlisted abstract-only precedents require method-level reading before experimental replication.

Authorship and status. Original concept and project direction: Yewei Liu. This synthesis, proposed standpoint and conceptual scenario review: assisted by Codex, awaiting Yewei's assessment. Existing study site, portfolio and native application remain unchanged.