Preprint · Not Peer Reviewed · Posted on PsyArXiv

Recurrence-Aware Conversation Decision Support in Emotionally Loaded Romantic Relationships: A Conceptual Framework

Mariia Yakovleva
Independent Researcher
mariia.creater@gmail.com

Author Note

The author is affiliated with a startup exploring product directions in AI-mediated support for difficult interpersonal conversations. This manuscript is a conceptual paper and does not report product efficacy data. The author declares no financial conflicts of interest with respect to the research presented here.

Correspondence concerning this article should be addressed to Mariia Yakovleva. Email: mariia.creater@gmail.com

March 31, 2026

Abstract

Many people do not struggle with a difficult romantic conversation only once. They face the same issue repeatedly: how to raise hurt, ask for clarity, set a boundary, or seek repair without making the situation worse. Current AI support is often framed as a prompt-response task in which a user gives context and a system returns candidate wording judged for fluency, empathy, or immediate helpfulness. That framing captures one layer of the interaction, but it misses a larger support problem. Hard romantic conversations are often repeated relational episodes shaped by prior attempts, recurring patterns, competing objectives, and uncertain reaction paths. Relationship research points to recurring interaction loops and multi-stage recovery processes, while human-computer interaction research treats hard conversations as problems of timing, pacing, emotional regulation, channel choice, and repair across stages. This paper does not claim that memory, branching, reflection, or supportive language are new. It argues that

when these elements are organized around repeated romantic-conversation episodes, the support problem changes from message generation to longitudinal relational decision support. The paper proposes six core functions: representing recurrence, identifying candidate patterns, clarifying objectives, detecting goal-strategy mismatch, preparing for plausible branches, and using post-conversation outcomes in later support. The contribution is conceptual rather than empirical: the paper introduces an analytical vocabulary for this class of systems, distinguishes it from adjacent categories, and outlines risks, limitations, and future directions for evaluation.

Keywords: conversation decision support; romantic relationships; AI-mediated support; recurrence; branching; conceptual framework

1. Introduction

People often enter difficult romantic conversations with much more at stake than finding the right sentence. They may be trying to raise a recurring hurt, ask for reassurance, define the relationship, set a boundary, repair after conflict, or decide whether the relationship still has a future. The challenge is not only what to say. It is also when to say it, how direct to be, what reaction to expect, how much prior history matters, and what consequences may follow after the opening move.

Yet AI support in this area is still often described as chatbot advice. In that framing, the user gives a prompt and the system returns wording evaluated for fluency, empathy, appropriateness, or immediate helpfulness. That can be useful when the user mainly wants phrasing help. But many hard romantic conversations are not simple wording problems. They are sequential, emotionally charged, and shaped by relational stakes that extend beyond a single reply.

Research across relationship science and human-computer interaction shows why this narrower framing misses the core difficulty. Difficult partner conversations often reflect recurring conflict structures, demand-withdraw cycles, and recovery processes that extend beyond one utterance (Baucom et al., 2015; Schrodt, Witt, & Shimkowski, 2014; Haydon, Jonestrask, Guhn-Knight, & Salvatore, 2017). Human-computer interaction work likewise treats hard conversations as design problems involving timing, emotional regulation, channel choice, pacing, and repair across stages (Baughan et al., 2021; Baughan et al., 2024; Chen et al., 2023). By contrast, reviews of empathic large language models (LLMs) and conversational support tools still tend to emphasize perceived empathy, fluency, supportive language, and immediate helpfulness (Sorin et al., 2024; Sharma et al., 2023). Even relationship-oriented chatbot studies may produce outputs that feel realistic or empathic while remaining thin on branching,

continuity, and strategic calibration across repeated episodes (Vowels, François-Walcott, & Darwiche, 2024a; Vowels, François-Walcott, & Darwiche, 2024b).

I argue that the dominant framing is too narrow in two ways. First, emotionally loaded romantic conversations are not well described as simple prompt-response tasks. Second, many of these conversations are not isolated one-time exchanges. They are repeated relational episodes shaped by prior attempts, recurring tensions, unfinished clarity, and familiar interaction loops. For that reason, AI support in this domain may be described more precisely as recurrence-aware conversation decision support rather than single-conversation wording assistance.

This claim is deliberately limited. I do not argue that memory, branching, reflection, or supportive language are new. The point is narrower: when these elements are organized around repeated romantic-conversation episodes, the main unit of support changes. The task is no longer just generating a locally helpful message. It becomes supporting judgment across time, under uncertainty, within an ongoing relationship.

The paper proceeds in five steps. It first explains why the chatbot framing is too narrow and why a single-conversation framing is also insufficient. It then defines the proposed category of recurrence-aware conversation decision support and identifies six core functions that organize it. Next, it distinguishes this category from adjacent systems, including therapy chatbots, journaling tools, generic LLM advice systems, and communication coaching tools. Finally, it outlines boundary conditions, risks, limitations, and a future empirical agenda for testing whether this framework adds value beyond strong one-shot alternatives.

1.1 Why Romantic Conversations Are a Strategically Concentrated Test Domain

Romantic conversations are used here as a strategically concentrated test domain, not as a claim of universal scope. This domain concentrates repeated episodes, emotional stakes, interdependence, ambiguity, delayed feedback, and nontrivial cost of miscalibration in unusually dense form. The same issue may return with the same partner across multiple attempts, yet each episode can be reshaped by prior disappointment, partial repair, hope, fear, or accumulated resentment. That concentration makes romantic conversations a useful setting for testing whether recurrence-aware decision support adds analytical value beyond one-shot wording help. The framework may later generalize to other close-relationship contexts, but the present paper remains bounded to romantic episodes in order to avoid premature generalization.

2. Why the Chatbot Framing Is Too Narrow

The chatbot framing remains attractive because it reduces a complex support problem to a familiar interaction model: the user asks, the model answers, and the exchange is judged by the apparent quality of the reply. For emotionally loaded romantic conversations, that model is often too thin. The support task is not only to produce fluent language; it is to support judgment under relational stakes, temporal continuity, and branching consequences.

First, the chatbot frame encourages output-quality confusion. A response can sound fluent, calm, empathic, and realistic while still being strategically wrong for the user's actual situation. A locally supportive answer may misidentify the governing objective, ignore recurrence, encourage premature disclosure, or underweight timing constraints. Warmth is therefore relevant but insufficient.

Second, the chatbot frame compresses a multi-stage process into a one-shot exchange. In practice, the opening move is followed by clarification, defensiveness, vagueness, withdrawal, repair, or renewed ambiguity. A system evaluated only on its first answer is being judged on the wrong slice of the problem.

Third, the chatbot frame underrepresents decision variables that are native to difficult romantic conversations. Timing, channel, recurrence, emotional readiness, perceived risk, objective conflict, and likely reaction paths are not peripheral details. They materially change what kind of support is appropriate.

Fourth, the chatbot frame encourages weak evaluation fit. Helpfulness, warmth, coherence, and empathy remain relevant, but they do not capture whether the system improved objective clarity, strategy alignment, branch preparedness, or post-conversation judgment. Those are decision-support outcomes, not merely response-quality outcomes.

3. Why the Single-Conversation Framing Is Also Too Narrow

Many emotionally loaded romantic conversations are not isolated events. They recur within a larger relational history, often involving the same unresolved issue, the same partner, and a similar pattern of escalation, avoidance, ambiguity, or disappointment. A user may present a conversation as a single decision point when it is actually one episode in a repeated interpersonal loop.

In recurrent situations, the relevant question is not only what wording or opening strategy fits the present exchange, but also what has already been tried, what pattern appears to organize the loop, whether the user's stated objective has remained stable across episodes, and whether previous conversations produced clarity, temporary relief, renewed ambiguity, or no behavioral change at all. A one-shot framing compresses this longitudinal structure into a single prompt.

Relationship research supports this shift because recurring conflict structures, demand-withdraw dynamics, and recovery sabotage are not edge cases in romantic interaction. They are central to how many conflicts unfold over time. That makes repeated episodes part of the support problem itself rather than incidental background detail.

This longitudinal view changes what support should aim to do. The question is no longer only, "What should I say now?" It becomes: "What keeps happening, what am I actually trying to obtain, what have prior attempts produced, and what is most likely to move the situation toward clarity rather than another loop?"

4. Hard Romantic Conversations as Repeated Relational Decision Environments

Hard romantic conversations differ from ordinary information exchange because they combine incomplete information, emotional pressure, relational memory, competing objectives, and uncertain consequences. The speaker rarely knows how the partner will interpret the issue, whether the conversation will lead to repair or rupture, or whether immediate honesty will reduce conflict or intensify it. The uncertainty is not only about wording. It is about consequences.

For conceptual purposes, these conversations can be treated as repeated relational decision environments. In decision-support research, a decision environment is defined not simply by the presence of a choice, but by the need to act under incomplete information, competing objectives, and uncertain consequences (Keen, 1980; Eierman, Niederman, & Adams, 1995; Power, 2001; Camacho et al., 2020). That description fits difficult romantic conversations closely, especially when the same issue has appeared before and the present episode is shaped by prior attempts.

The structure of outcomes is branching rather than singular. The opening move may be followed by receptiveness, defensiveness, vagueness, withdrawal, minimization, softening, or escalation. Consequences may also be delayed. A conversation that seems calm in the moment may later produce distance or renewed ambiguity, while an initially difficult exchange may eventually create clarity or repair. Gottman and Silver (1999) made this broader point influential by showing that stable and unstable couples differ not only in what they say, but in how negative cycles escalate and how repair attempts interrupt them. Finkel et al. (2013) provide a related precedent: even a lightweight structured conflict-reappraisal intervention was associated with preserved marital quality over time. That does not establish that AI support systems would produce equivalent effects, but it does suggest that structured cognitive scaffolding can matter in relational settings.

Once the task is framed in this way, the relevant support problem changes. The issue is not merely message generation. It is structured support under interpersonal uncertainty, where the user is making a sequence of judgments about timing, disclosure, objective prioritization, emotional regulation, conversational framing, and follow-up. Language generation remains important, but it should be placed inside a broader architecture of relational decision support.

5. Defining the Category: Recurrence-Aware Conversation Decision Support

For the present framework, the inputs to support include issue type, recurrence, remembered relational history, current emotional state, stated objective, feared outcome, communication channel, and urgency. These inputs matter because the same wording can function differently depending on whether the issue is new or repeated, whether the user seeks clarity or reassurance, and whether the setting makes escalation more or less likely.

The intermediate judgments of the system concern what pattern may be recurring, whether the user seeks clarity, reassurance, repair, accountability, or action; whether the chosen strategy fits the stated objective; and which branch risks are most plausible. The central

analytical claim is that these judgments belong to the support problem itself rather than to optional elaboration after an answer has already been generated.

This paper proposes the term recurrence-aware conversation decision support system for a subset of AI-mediated tools designed to help users navigate emotionally loaded romantic conversations across repeated episodes rather than only generate one-shot advice. This paper does not argue that memory, branching, reflection, or support language are new. It argues that when these are jointly organized around repeated romantic-conversation episodes, the relevant analytical object changes from message generation to longitudinal relational decision support. The defining feature is not simply that the system produces text. It is that support is organized around six recurring functions: representing recurrence, identifying candidate patterns, clarifying objectives, detecting strategy mismatch, preparing for plausible branches, and incorporating post-conversation outcomes into later support. The framework is offered as an analytical description of a design and research space rather than as a proprietary claim about a single implementation.

5.1 Why These Six Functions

The six functions are proposed because they jointly cover the minimum conceptual territory required for recurrence-aware conversational judgment support in this domain. Taken together, they capture temporal continuity, pattern-level interpretation, decision-objective specification, alignment checking, contingency preparation, and longitudinal updating. The claim is not that these six functions are the only possible components of a future system. The narrower claim is that they constitute a minimally sufficient conceptual set for describing support whose primary object is the repeated romantic-conversation episode rather than the isolated generated message.

The selection is therefore principled rather than decorative. Recurrence representation addresses whether the present issue is new or repeated. Candidate-pattern identification addresses whether the episode appears embedded in a recognizable loop. Objective clarification addresses what the user is actually trying to obtain. Goal-strategy mismatch detection addresses whether the chosen opening is likely to advance that objective. Branch preparation addresses the contingent structure of the partner's response. Outcome integration addresses whether the results of one episode become analytically relevant to the next. Omitting any one of these functions would narrow the framework back toward one-shot support, generic coaching, or reflection without structured decision logic.

5.2 Why This Is a Category and Not a Feature Bundle

A system should be classified as recurrence-aware conversation decision support only if it passes three tests. First, the default-structure test: recurrence representation, pattern inference, objective clarification, mismatch detection, branch preparation, and outcome integration must be built into the default support flow rather than appearing only when an unusually skilled user asks for them. Second, the evaluation test: the system must be judged on cross-episode judgment

support, not only on local response quality such as fluency, empathy, or immediate helpfulness. Third, the unit-of-analysis test: the repeated relational episode, rather than the single generated message, must be the primary object of support. If a system fails two of these three tests, it should not be treated as a clean member of the proposed category.

This framework should not be interpreted as claiming novelty for memory, branching, reflection, or supportive language as isolated capabilities. The category claim is structural. A feature bundle merely aggregates available capabilities. A category changes what the system is designed to do by default, how it is evaluated, and what counts as the core support object. On that logic, recurrence-aware conversation decision support is not defined by having more features than a chatbot. It is defined by reorganizing support around repeated relational episodes as the central analytical target.

5.3 Minimal Criteria and Falsifiability of the Proposed Category

The proposed category is not a loose feature checklist. It makes a stronger claim about default system structure. Category membership requires that all six functions operate as standard components of support rather than as ad hoc behaviors elicited only by unusually skilled prompting. A system that can occasionally simulate recurrence reasoning, branch planning, or reflection does not thereby instantiate a recurrence-aware conversation decision support architecture.

The claim is also falsifiable. It would be weakened if strong one-shot systems, under matched scenario constraints, performed comparably on objective clarity, goal-strategy alignment, branch preparedness, and cross-episode adaptation without requiring those functions to be built into default support flow. It would also be weakened if the six functions failed to cohere empirically as a meaningful design logic and instead behaved as loosely connected optional features. In that case, the proposal would be better treated as a useful checklist than as a defensible category.

5.4 Core Constructs

1. Recurrence representation. The construct is the system's representation of whether the issue has occurred before with the same partner, in what form, and with what remembered outcome. It is not mere storage of prior text or generic memory availability. The minimal observable sign is that support changes when the same issue is marked as repeated rather than new. The easiest failure mode is false reset: the system treats a repeated episode as if history does not matter.

2. Relational pattern inference. The construct is the system's inference that the episode may be embedded in a recognizable loop such as repeated avoidance, repeated reassurance-seeking, repeated ambiguity, inconsistency, or disappointment. It is not diagnosis, motive-reading, or certainty about the partner. The minimal observable sign is that the system surfaces a candidate pattern with explicit uncertainty language. The easiest failure mode is interpretive overreach: a thin pattern guess is presented as settled truth.

3. Objective clarification. The construct is explicit specification of what the user is actually trying to obtain, such as clarity, reassurance, repair, accountability, boundary-setting, behavioral change, or closure. It is not generic asking what the user wants in a casual sense. The minimal observable sign is forced selection or prioritization of a governing objective. The easiest failure mode is objective blur: multiple aims remain collapsed into undisciplined "better communication."

4. Goal-strategy mismatch detection. The construct is detection of cases in which the user's selected conversational strategy is structurally unlikely to advance the stated objective. It is not the system merely preferring different wording or style. The minimal observable sign is an explicit flag that the chosen opening optimizes for a different outcome than the one the user says they want. The easiest failure mode is cosmetic editing masquerading as judgment support.

5. Branch preparedness. The construct is preparation for plausible partner responses such as receptiveness, defensiveness, vagueness, withdrawal, minimization, or counterattack. It is not generic reassurance that "things may go different ways." The minimal observable sign is differentiated response planning for at least several credible branches. The easiest failure mode is linearity bias: the system assumes the partner will respond as hoped and leaves the user unprepared for divergence.

6. Post-conversation outcome integration. The construct is incorporation of what happened after the conversation into later support. It is not passive archiving or a dead reflection journal. The minimal observable sign is that later recommendations change in response to prior outcomes. The easiest failure mode is inert memory: the system stores post-conversation information but does not use it analytically.

5.5 Implementation-Facing Support Components

The six constructs above define the conceptual center of the framework. A second layer concerns implementation-facing components through which support may actually be delivered. These include message-entry design, tone calibration, reaction-path support, next-step support, and a structured reflection artifact. They matter for usability and delivery, but they are secondary to the category claim itself.

Message-entry design concerns how the conversation should begin, including whether the opening should be direct or consent-first, concise or explanatory, emotionally open or more contained, and adapted to text, call, voice note, or in-person settings. Tone calibration concerns how the message should feel, not only what it should say, including warmth, firmness, vulnerability, brevity, and emotional containment. Reaction-path support concerns explicit preparation for defensiveness, vagueness, avoidance, receptiveness, and repair. Next-step support concerns what to do after the initial exchange. A structured reflection artifact concerns how prior support is assembled into a usable plan and how post-conversation observations are captured for future sessions.

These implementation-facing components remain important, but they should not dominate the conceptual account. Otherwise the framework risks collapsing back into a better coaching tool. Their role is to operationalize a system whose central logic is recurrence-aware decision support rather than one-shot phrasing assistance. Figure 1 summarizes this support flow from inputs to intermediate judgments, decision-support outputs, post-conversation updating, and the next episode.

Figure 1. Conceptual support flow for recurrence-aware conversation decision support.

6. Operationalization and Future Measurement

One reason conceptual papers in this area remain weak is that they often stop at naming functions and do not specify what could later be measured. The present framework can be operationalized without implying efficacy evidence that has not yet been produced. Table 1 maps each construct to a candidate indicator and a conservative future measure. The point is not to imply that these variables have already been validated. The point is to show that the framework can be translated into observable, falsifiable components rather than remaining purely conceptual.

To keep the measurement logic disciplined, three outcomes should be treated as primary in future work: objective clarity, goal-strategy alignment, and branch preparedness. Other outcomes such as perceived usefulness, emotional validation, or satisfaction may remain relevant, but they should be treated as secondary or exploratory unless they are tied to a clearer decision-support construct.

6.1 Formal Propositions

P1. In repeated romantic-conversation scenarios, recurrence changes the support problem from local message generation to sequential relational judgment support.

P2. Objective ambiguity is not peripheral noise; it is a central variable in relational conversation design because different objectives require materially different openings, tones, and follow-up paths.

P3. Branch preparedness is analytically distinct from affective helpfulness; a response can feel validating while leaving the user strategically unprepared for plausible partner reactions.

P4. Systems that do not explicitly distinguish objective clarification from message generation will more often produce strategically misaligned support.

P5. Post-conversation outcome integration is necessary for longitudinal adaptation across repeated relational episodes.

P6. If the six proposed functions do not cohere empirically as a recognizable bundle, the category claim should be revised or abandoned.

7. Comparison with Adjacent Categories

A useful category must do more than sound plausible. It must distinguish itself from adjacent categories without overstating novelty or imposing artificial boundaries. The proposed framework overlaps with therapy chatbots, journaling tools, generic LLM advice systems, and communication coaching applications. The analytic question is therefore not whether those neighboring categories are unrelated. It is whether a subset of AI-mediated tools is better described by an organizing logic centered on repeated, emotionally loaded romantic-conversation decisions.

Therapy chatbots overlap because some already provide structured reflection and guided dialogue rather than free-form conversation alone (Yuksel & Kocaballi, 2022). Their failure pattern in this context appears when emotional support substitutes for relational decision analysis. A system may validate feelings effectively yet leave the user without explicit objective clarification, mismatch detection, or branch preparation under relational stakes.

Generic LLM advice systems are a closer comparator because they can simulate multi-stage support when users explicitly ask for it. Their failure pattern appears when the user receives locally empathic text that is strategically misaligned with the actual objective. The distinction is therefore not technical impossibility. It is default architecture and evaluative target: recurrence-aware support makes recurrence representation, objective clarification, mismatch detection, and branch planning standard parts of support rather than optional add-ons invoked only when the user knows to ask.

Communication coaching applications also overlap substantially. Some already include scripts, rules, branching examples, or rehearsal features. Their failure pattern appears when scripts remain static across repeated episodes with changing outcomes. The proposed category is narrower and more situational: it focuses on repeated romantic-conversation episodes under relational stakes, where recurrence, candidate-pattern surfacing, and outcome-linked adaptation are central rather than peripheral.

Journaling tools remain an important adjacent case because they support private reflection and self-observation. Their failure pattern appears when insight is generated but no branch-ready action structure is produced. Even when journaling improves self-understanding, it usually does not infer candidate loops, detect mismatch between objective and strategy, prepare branch responses, or integrate outcomes into a structured next session. The value of the proposed category lies in making that recurring bundle visible enough to classify, compare, and evaluate directly. See Table 3 for a systematic comparison.

8. Boundary Conditions and Non-Claims

A stronger framework also requires a precise account of what this category does not claim to do. A recurrence-aware conversation decision support system should not be described as diagnosing attachment style, determining a partner's motives as fact, adjudicating which participant is objectively correct, or identifying abuse solely from ordinary conversational input.

It may generate candidate interpretations, but those interpretations remain defeasible and should be presented as hypotheses rather than truths.

It should also not be framed as a therapist, mediator, legal advisor, or safety assessor. In emotionally loaded but non-dangerous scenarios, the system may help organize judgment and preparation. In coercive, abusive, retaliatory, or clinically complex situations, the appropriate output may instead be caution, delay, referral, or non-engagement. These non-claims are not cosmetic disclaimers. They are part of the scientific boundary of the proposed category because they define where the framework stops being decision support and starts drifting into unjustified authority.

9. Normative Risks and Objections

A stronger scientific account requires treating the risks of such systems as central rather than peripheral. Seven objections warrant careful attention.

One-sided narrative reinforcement. A conversation decision support system usually receives only one participant's account. That creates structural asymmetry: the system can generate coherent support while lacking access to the partner's interpretation, omitted context, or contradictory evidence. Alon-Barkat and Busuioc (2023) show that users may selectively adhere to algorithmic advice when it confirms a prior narrative. In romantic conflicts, that risk is especially serious because a system may feel clarifying while actually hardening a one-sided interpretation.

Automation bias. Research on automation has long shown that people may over-defer to machine recommendations when verification is costly or emotionally taxing (Skitka, Mosier, & Burdick, 1999; Parasuraman & Riley, 1997; Lyell & Coiera, 2017). Hard romantic conversations are already cognitively and emotionally demanding. That makes over-reliance a realistic possibility rather than a remote edge case.

Pseudo-manipulation risk. Branch preparation and strategic framing can be used for constructive preparation, but they can also become instrumentalized. A system that helps a user anticipate likely reactions may drift from preparation into tactical steering of the other person's emotions, concessions, or disclosures. Design must therefore distinguish preparation for clarity and repair from covert leverage.

Unsafe relationship context. Some emotionally loaded relationships involve abuse, coercion, intimidation, or retaliation risk. In those contexts, a conventional planning system may be inappropriate or unsafe. The right response may be caution, delay, referral, safety-oriented support, or non-engagement rather than better wording (Khawaja & Belisle-Pipon, 2023).

False interpretive authority. Because these systems can produce organized analyses, they may appear to possess privileged access to motives or hidden patterns. That appearance can mislead. The system is still reasoning from limited user-provided material. Ehsan et al. (2021) argue that socially situated AI systems require forms of transparency that make their reasoning legible in human terms. In this domain, that means surfacing uncertainty, distinguishing

inference from fact, and showing why one interpretation was suggested without implying that alternatives have been ruled out.

Avoidance reinforcement. A system may appear useful because it helps the user reflect, articulate, and plan. Yet in some cases that structure may become a substitute for action rather than preparation for action. A user may repeatedly engage the system to refine wording, revisit interpretation, or seek reassurance without actually moving toward clarity, boundary-setting, or decision. Future evaluation should therefore distinguish support that improves strategic readiness from support that merely extends avoidance in a more organized form.

Clinical overreach. A recurrence-aware conversation decision support system is not a therapist, mediator, diagnosis tool, or adjudicator of truth. Even when the interaction feels empathic or insightful, it should not be framed as delivering therapy or authoritative relationship judgment (Khawaja & Belisle-Pipon, 2023).

10. Future Empirical Directions

Because the present paper is conceptual rather than empirical, a credible next step is not to infer effectiveness but to specify what future studies should test. The framework supports a research agenda organized around falsifiable questions about structure, judgment, and risk.

The framework would be weakened if future comparative studies showed that recurrence-aware structure provides no measurable advantage over strong one-shot support on objective clarity, strategy alignment, branch readiness, or adaptation across repeated episodes. It would also be weakened if the six proposed functions did not cluster coherently as a distinct design logic and instead behaved as loosely connected optional features. Four specific hypotheses follow:

H1. In repeated romantic-conversation scenarios, recurrence-aware support will be associated with higher blinded ratings of strategic usefulness than one-shot advice generation.

H2. Explicit objective clarification will be associated with better user discrimination between reassurance-seeking, clarity-seeking, repair-seeking, and boundary-setting in matched scenarios.

H3. Goal-strategy mismatch detection will be associated with lower preference for openings judged by blinded coders to be poorly aligned with the user's stated objective.

H4. Branch-preparedness support will be associated with higher branch-readiness scores, operationalized as the user's ability to generate differentiated responses to defensive, vague, avoidant, or receptive partner reactions.

A particularly important next step is comparative evaluation against strong baselines rather than weak strawmen. The relevant comparison is not between a recurrence-aware system and obviously poor advice. It is between recurrence-aware support and the best available

one-shot prompting, coaching, or journaling alternatives under the same scenario constraints. Table 2 summarizes example future study designs aligned to the framework.

11. Limitations

Several limitations should be stated clearly. First, this paper is conceptual rather than empirical. The proposed framework has not been validated through controlled testing, field deployment, or outcome comparison. Second, the six-function decomposition may be incomplete, incorrectly bundled, or theoretically overdrawn. Third, the category may collapse under comparative testing if strong one-shot systems perform similarly on objective clarity, goal-strategy alignment, branch preparedness, and cross-episode adaptation without requiring this architecture. Fourth, scenario-based gains, if later observed, may not transfer to real dyadic outcomes in live relationships. Fifth, user-reported post-conversation outcomes may be biased, selectively remembered, or strategically self-serving. Sixth, category boundaries may overlap in practice because some therapy chatbots, generic language-model systems, or coaching applications may already contain parts of the proposed logic. Seventh, the framework remains intentionally high level and should not be mistaken for a validated engineering blueprint. Different systems could instantiate the same conceptual architecture using prompts, rules, retrieval pipelines, memory structures, user-defined templates, or hybrid methods. Eighth, the framework focuses on user-facing support and does not resolve partner consent, dyadic symmetry, or the ethics of asymmetric AI assistance in couple-level contexts.

12. Conclusion

AI support for emotionally loaded romantic conversations is too often described only as chatbot-style advice generation. That description captures a real layer of the interaction, but it misses a more important analytical distinction: many hard romantic conversations are not isolated wording events, but repeated relational episodes shaped by recurrence, memory, ambiguity, and prior failed attempts at clarity.

The relevant design problem is therefore not only how a system generates locally helpful language. It is whether the system can represent recurrence, surface candidate patterns, clarify the user's governing objective, flag strategy mismatch, prepare for plausible branches, and incorporate what happened after the conversation into later support.

The contribution of this paper is deliberately narrower than an efficacy claim. This paper does not argue that memory, branching, reflection, or support language are new. It argues that when these are jointly organized around repeated romantic-conversation episodes, the relevant analytical object changes from message generation to longitudinal relational decision support. The paper therefore proposes a conceptual category and a set of constructs that can later be tested against strong baselines. If future work in this area is to be rigorous, the central question should not be whether the system sounds empathic. It should be whether the proposed architecture can be evaluated, and potentially falsified, on objective clarity, goal-strategy alignment, branch

preparedness, and cross-episode adaptation. Those questions are open. The framework provides one analytical vocabulary for pursuing them.

If repeated romantic conversations are structurally longitudinal and strategically contingent, then systems designed for this domain should be evaluated not only as generators of supportive language, but as supports for sequential relational judgment.

Data Availability Statement

No empirical dataset was generated or analyzed for this manuscript. No code repository is required for the present conceptual paper.

Disclosure and Conflict of Interest

The author is affiliated with a startup exploring product directions related to AI support for difficult interpersonal conversations. This manuscript is a conceptual paper and does not report product efficacy data. No funding was received for this work.

Table 1. Core Constructs, Example Indicators, and Possible Future Measures

Construct	Definition	Example indicator	Possible future measure
Recurrence representation	Whether the issue has occurred before with the same partner and with what remembered outcome	User reports that the same issue has been discussed repeatedly	Count of prior episodes or coded repeat-topic frequency
Relational pattern inference	Candidate repeated loop organizing the episode	Avoidance, reassurance-seeking, ambiguity, inconsistency, disappointment	Coded pattern-category frequency with uncertainty labels
Objective clarification	What the user most wants from the conversation	Clarity, reassurance, repair, accountability, closure	Categorical objective selection and objective-stability score across sessions
Goal-strategy mismatch	Inconsistency between stated objective and selected opening strategy	User says they want clarity but selects an opening optimized for temporary reassurance	Mismatch rate or blinded coder agreement on mismatch

Construct	Definition	Example indicator	Possible future measure
Branch preparedness	Preparation for plausible response paths	Defensiveness, vagueness, withdrawal, minimization, counterattack	Scenario-based branch-readiness score or branch-coverage score
Outcome integration	Use of post-conversation information in later support	System incorporates what happened after the last conversation	Presence of reflection artifact; change in later recommendations based on prior outcome

Note. The six constructs are proposed as minimally sufficient for recurrence-aware conversation decision support. "Possible future measure" entries are not validated; they indicate operationalization directions for later empirical work.

Table 2. Example Future Study Designs Aligned to the Framework

Study design	Comparison	Primary outcome	Key failure mode to watch
Scenario-based experiment	Recurrence-aware system vs. one-shot advice baseline	Blinded strategic-usefulness ratings, objective clarity, branch-readiness score	Users may rate polish or empathy higher without actual strategic gain
Blind coder assessment of opening choices	User-selected openings before vs. after mismatch detection	Reduction in choices poorly aligned with the user's stated objective	Coders may disagree on what counts as mismatch
Repeated-session diary study	Repeated-issue users across multiple sessions	Objective stability, recurrence detection, adaptation of later support to prior outcomes	Reflection may become a reassurance ritual rather than action support
Risk-focused safety evaluation	High-conflict or potentially unsafe scenarios with expert review	Appropriateness of caution, referral, delay, or non-engagement guidance	System may overconfidently encourage engagement where caution is warranted

Note. Designs are illustrative. Baseline comparisons should use the best available one-shot alternative, not a weak strawman.

Table 3. Comparison With Adjacent Categories

Category	Primary function	Interaction style	Typical limitation
Therapy chatbots	Emotional support / mental-health-style interaction	Affect-oriented dialogue; some structured reflection	Fails when emotional support substitutes for relational decision

Category	Primary function	Interaction style	Typical limitation
			analysis under romantic stakes
Journaling tools	Self-reflection	User-led writing and review	Fails when insight is generated but no branch-ready action structure is produced
Generic LLM advice systems	Generate broad advice from prompts	Reactive, prompt-response interaction; can simulate structure when explicitly asked	Fails when locally empathic text is strategically misaligned with the user's actual objective
Communication coaching apps	Skill practice / scripts / tips	Templates, scenario coaching, rules; some branching examples	Fails when scripts remain static across repeated episodes with changing outcomes
Conversation decision support systems	Structured support for navigating difficult interpersonal decisions	Default support centered on goal clarification, strategy support, branch preparation, and post-conversation integration	Still vulnerable to one-sided context, overreach, and misuse if poorly designed

Note. Category boundaries are porous in practice. Hybrid systems occupying more than one column are expected. The rightmost row (conversation decision support systems) describes the proposed category.

References

Note. References follow APA 7th edition format.

Alon-Barkat, S., & Busuioc, M. (2023). Human-AI interactions in public sector decision making: "Automation bias" and "selective adherence" to algorithmic advice. Journal of Public Administration Research and Theory, 33(1), 153–169. https://doi.org/10.1093/jopart/muac033

Baucom, B. R., Dickenson, J. A., Atkins, D. C., Baucom, D. H., Fischer, M. S., Weusthoff, S., & Hahlweg, K. (2015). The interpersonal process model of demand/withdraw behavior. Journal of Family Psychology, 29(1), 80–90. https://doi.org/10.1037/fam0000044

Baughan, A., Petelka, J., Yoo, C. J., Lo, J., Wang, S., Paramasivam, A., Zhou, A., & Hiniker, A. (2021). Someone is wrong on the internet: Having hard conversations in online spaces. Proceedings of the ACM on Human-Computer Interaction, 5(CSCW1), Article 156, 1–22. https://doi.org/10.1145/3449181

Baughan, A., Tian, L., Shekar, P., Zhang, A., & Hiniker, A. (2024). Supporting hard conversations in close relationships through design. Proceedings of the ACM on Human-Computer Interaction, 8(CSCW2), Article 379, 1–22. https://doi.org/10.1145/3686906

Camacho, J., Zanoletti-Mannello, M., Landis-Lewis, Z., Kane-Gill, S. L., & Boyce, R. D. (2020). A conceptual framework to study the implementation of clinical decision support systems (BEAR): Literature review and concept mapping. Journal of Medical Internet Research, 22(8), e18388. https://doi.org/10.2196/18388

Caughlin, J. P., & Vangelisti, A. L. (2000). An individual difference explanation of why married couples argue about the topics they do. Journal of Social and Personal Relationships, 17(4–5), 523–551. https://doi.org/10.1177/0265407500174003

Chen, T., Lee, C., Mindel, J. R., Elhaouij, N., & Picard, R. W. (2023). Closer Worlds: Using generative AI to facilitate intimate conversations. Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems. https://doi.org/10.1145/3544549.3585613

Ehsan, U., Liao, Q. V., Muller, M. J., Riedl, M. O., & Weisz, J. D. (2021). Expanding explainability: Towards social transparency in AI systems. Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems. https://doi.org/10.1145/3411764.3445188

Eierman, M. A., Niederman, F., & Adams, C. (1995). DSS theory: A model of constructs and relationships. Decision Support Systems, 14(1), 1–26. https://doi.org/10.1016/0167-9236(94)00014-8

Finkel, E. J., Slotter, E. B., Luchies, L. B., Walton, G. M., & Gross, J. J. (2013). A brief intervention to promote conflict reappraisal preserves marital quality over time. Psychological Science, 24(8), 1595–1601. https://doi.org/10.1177/0956797613488120

Gottman, J. M., & Silver, N. (1999). The seven principles for making marriage work. Harmony Books.

Haydon, K. C., Jonestrask, C., Guhn-Knight, H., & Salvatore, J. E. (2017). The dyadic construction of romantic conflict recovery sabotage. Journal of Social and Personal Relationships, 34(6), 915–935. https://doi.org/10.1177/0265407516657096

Keen, P. G. W. (1980). Decision support systems: A research perspective. MIT Center for Information Systems Research.

Khawaja, Z., & Belisle-Pipon, J.-C. (2023). Your robot therapist is not your therapist: Understanding the role of AI-powered mental health chatbots. Frontiers in Digital Health, 5, 1278186. https://doi.org/10.3389/fdgth.2023.1278186

Lyell, D., & Coiera, E. (2017). Automation bias and verification complexity: A systematic review. Journal of the American Medical Informatics Association, 24(2), 423–431. https://doi.org/10.1093/jamia/ocw105

Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. Human Factors, 39(2), 230–253. https://doi.org/10.1518/001872097778543886

Power, D. J. (2001). Supporting decision-makers: An expanded framework. Informing Science, 4(3), 113–123. https://doi.org/10.28945/579

Scholich, T., Barr, M., Stirman, S. W., & Raj, S. (2025). A comparison of responses from human therapists and large language model-based chatbots to assess therapeutic communication: Mixed methods study. JMIR Mental Health, 12, e69709. https://doi.org/10.2196/69709

Schrodt, P., Witt, P. L., & Shimkowski, J. R. (2014). A meta-analytical review of demand/withdraw patterns of interaction and their associations with relational, individual, and communicative outcomes. Communication Monographs, 81(1), 28–58. https://doi.org/10.1080/03637751.2013.813632

Shaikh, O., Chai, V. E., Gelfand, M. J., Yang, D., & Bernstein, M. S. (2024). Rehearsal: Simulating conflict to teach conflict resolution. Proceedings of the CHI Conference on Human Factors in Computing Systems. https://doi.org/10.1145/3613904.3642229

Sharma, A., Lin, I. W., Miner, A. S., Atkins, D. C., & Althoff, T. (2023). Human-AI collaboration enables more empathic conversations in text-based peer mental health support. Nature Machine Intelligence, 5, 46–57. https://doi.org/10.1038/s42256-022-00583-4

Skitka, L. J., Mosier, K., & Burdick, M. (1999). Does automation bias decision-making? International Journal of Human-Computer Studies, 51(5), 991–1006. https://doi.org/10.1006/ijhc.1999.0252

Sorin, V., Brin, D., Barash, Y., Konen, E., Charney, A., Nadkarni, G., & Klang, E. (2024). Large language models and empathy: Systematic review. Journal of Medical Internet Research, 26, e52597. https://doi.org/10.2196/52597

Vowels, L. M., François-Walcott, R. R. R., & Darwiche, J. (2024a). Are chatbots the new relationship experts? Insights from three studies. Computers in Human Behavior: Artificial Humans, 2, 100077. https://doi.org/10.1016/j.chbah.2024.100077

Vowels, L. M., François-Walcott, R. R. R., & Darwiche, J. (2024b). AI in relationship counselling: Evaluating ChatGPT's therapeutic capabilities. Computers in Human Behavior: Artificial Humans, 2(2), 100078. https://doi.org/10.1016/j.chbah.2024.100078

Vowels, L. M., François-Walcott, R. R. R., Grandjean, M., & Darwiche, J. (2024c). The efficacy, feasibility, and technical outcomes of a GPT-4o-based chatbot Amanda for relationship support: A randomized controlled trial. PLOS Mental Health. https://doi.org/10.1371/journal.pmen.0000136

Yuksel, B., & Kocaballi, A. B. (2022). Conversational agents to support couple therapy. In Proceedings of the 12th Nordic Conference on Human-Computer Interaction (NordiCHI 2022). https://doi.org/10.1145/3546155.3546700