

# **A Pilot Rule-Based Operationalization of Conversational Objectives and Opening Strategies in Reddit Relationship-Help Posts**

Mariia Yakovleva

Independent Researcher

April 2026

*Author note.* We are affiliated with a startup exploring product directions in AI-mediated support for difficult interpersonal conversations. Our paper reports a pilot corpus analysis and does not claim product efficacy.

## **Abstract**

We ask a simple but important question: when people turn to online relationship-help forums before a hard romantic conversation, can we recover what they want from that conversation, how they plan to open it, and whether that opening seems to fit the goal? We build a rule-based extraction and pre-annotation pipeline for this task and test it on Reddit submissions. In our larger streamed corpus, we scanned 300,001 submission-level posts from relationship-oriented subreddits and extracted 332 candidate conversation-target narratives. After a second-pass screen, 91 posts were retained as likely on-target cases for pilot pre-annotation. In that subset, automated objective recovery remained partial: 47 posts stayed unclear, while 25 were labeled accountability, 11 closure, 6 clarity, 1 boundary-setting, and 1 repair. Strategy detection was easier but coarse, with 82 of 91 posts collapsing into direct discussion. Mismatch was recoverable in only a small number of cases under our initial rule set. We interpret these results

as evidence that goal-strategy mismatch is a meaningful but harder-to-operationalize construct than broad conversational objective alone. The paper therefore contributes a pilot workflow, a transparent account of what worked and what failed, and a concrete roadmap for stronger future annotation and modeling.

*Keywords:* relationship-help narratives, conversational objectives, goal-strategy mismatch, pilot operationalization, Reddit corpus

## **1. Introduction**

Hard romantic conversations are interesting because they look simple from the outside and messy from the inside. A person says, “I just need help figuring out what to say.” But once we look closer, the problem is usually bigger than wording. People may want clarity, reassurance, accountability, repair, a boundary, or a clean ending. They may also want several of those things at once. In real life, however, they still have to choose a first move. They have to decide whether to ask directly, soften the issue, wait, confess, confront, or walk away. That first move matters because it can pull the conversation toward clarity – or push it toward confusion.

That is why we treat difficult romantic conversations as decision problems rather than sentence problems. Prior relationship research shows that there is no single universally “best” way to communicate during conflict. Directness can help in some contexts and backfire in others, depending on the issue, the relationship, and the interaction pattern around it (Overall & McNulty, 2017). Research on attachment and conflict communication likewise shows that romantic conflict often unfolds through repeated patterns such as mutual constructive communication, demand-withdraw, mutual avoidance, and withholding (Domingue & Mollen,

2009). Demand-withdraw matters here because it shows how emotional pressure and relational expectations can shape the very form of the opening move (Leo et al., 2020).

For us, the most important implication is simple: a person can want one thing and still choose an opening that points in another direction. Someone who wants clarity may keep waiting. Someone who wants accountability may soften the issue so much that the core request disappears.

Someone who wants repair may open with accusation and trigger defensiveness. In other words, difficult romantic conversations are not just about sounding kind, calm, or articulate. They are about goal selection, strategic fit, and anticipated consequences.

Our paper grows directly out of our earlier work. In our conceptual paper on recurrence-aware conversation decision support, we argued that difficult romantic conversations should be understood not as one-shot wording tasks but as repeated relational decision environments. In that framework, goal-strategy mismatch was one of six core functions. In our later empirical bridge paper, we took a first quantitative step by showing that relationship-relevant text contains recoverable structure across dialogue turns, help-seeking narratives, and rough episode trajectories. At the same time, we were explicit that our earlier empirical work did not directly operationalize goals, intentions, or goal-strategy mismatch. Our present paper is our attempt to take that next step in a practical way.

Online relationship-help forums are a strong place to start because they capture people in the middle of uncertainty. They do not just contain opinions about relationships in the abstract. They often contain moments of decision: should I tell him, should I ask her, should I bring this up, should I break up, should I confess, should I wait? Relationship help-seeking also happens far beyond formal therapy. It happens through books, friends, internet resources, and anonymous online communities (Stewart et al., 2016). Prior work shows that seeking relationship help from

strangers online can be studied systematically and at scale (Entwistle et al., 2021). That makes these narratives useful as a naturalistic corpus, even though they are noisy and incomplete.

At the same time, we do not want to oversell what a first-pass corpus study can do. Public Reddit data come with ethical and methodological issues around scraping, publicness, and researcher responsibility (Adams, 2024). More importantly, our target construct is difficult. It is one thing to detect that a post is about an upcoming conversation. It is another to infer the user's main objective. It is harder still to infer whether the planned strategy is aligned or mismatched with that objective. Our paper therefore has a deliberately narrow aim. We do not claim to validate an entire decision-support framework. We do not claim product efficacy. We ask whether we can build a transparent pilot pipeline that gets us closer to the construct in a disciplined way.

To address this question, we begin with a broad submission-level Reddit corpus, narrow it to posts that resemble conversation-target narratives, and then apply a rule-based pre-annotation layer to recover objective, strategy, mismatch, and confidence. The resulting pattern is asymmetric: conversational objectives are more recoverable than mismatch under the current rule set, while broad strategy is easier to detect than fine-grained strategic fit. This asymmetry is itself informative, as it highlights differential levels of recoverability across components of the construct.

Seen this way, our paper contributes three things. First, we provide a transparent extraction and screening workflow for isolating conversation-target relationship-help posts from a much larger Reddit stream. Second, we show what a first-pass operationalization can and cannot recover from those posts. Third, we treat the limitations of the pipeline as substantive findings about the difficulty of operationalizing the construct. For a pilot study, that is useful: it tells us where the

construct is already measurable and where stronger annotation, richer models, or human-coded audits will be needed next.

## **2. Research Question and Hypotheses**

Our central question is whether goal-strategy mismatch can be operationalized and quantified in online relationship-help narratives using a corpus-based pilot framework. We answer that question in stages, because our workflow itself has several layers: extraction, screening, pre-annotation, and interpretation.

We began with two hypotheses:

- H1. A meaningful subset of filtered online relationship-help posts will contain enough information to code both a conversational objective and an intended opening strategy.
- H2. Goal-strategy mismatch will appear at a non-trivial rate in the coded corpus.

Our final results support H1 most clearly.

## **3. Method**

### ***3.1 Corpus Development Logic***

We developed the corpus in two stages. First, for pipeline development, we used the Kaggle Reddit dataset janldeboer/reddit-relationships, a submission-level relationship-help corpus containing title and body fields. This pilot corpus comprised 1,885 submission rows. A strict conversation-target filter yielded 51 candidate posts, which were reduced through second-pass screening to a strict core of 23 posts and an expanded clean core of 28 posts. These smaller

subsets were used to debug the extraction logic and identify early weaknesses in the label system.

Second, after confirming that the pilot material consisted of original advice-seeking submissions rather than comments, we constructed the main workflow by applying a rule-based screening pipeline to a larger submission-level Reddit corpus. We streamed the first 300,001 submissions from relationship-oriented subreddits and restricted attention to `relationship_advice`, `relationships`, and `dating_advice`. Intermediate outputs from this workflow were saved as `streamed_screening_sheet_v2.csv`, `keep_only_91.csv`, and `preannotated_keep_91_v2.csv`. The final pilot pre-annotation analysis reported in this paper was conducted on the 91-post keep set retained after second-pass screening.

**Table 1. Corpus flow from broad stream to pilot analysis set.**

Stage	Count	Meaning
Smaller pilot submission corpus	1,885	Used to debug extraction logic and verify that we were working with original submissions rather than comments.
Strict candidates from smaller pilot corpus	51	Posts that matched the narrow conversation-target filter.
Strict core from smaller pilot corpus	23	Very clean pilot cases after a second-pass screen.
Expanded clean core from smaller pilot corpus	28	Slightly larger but somewhat noisier pilot set.
Large streamed submission corpus scanned	300,001	Submission-level Reddit corpus from relationship-oriented subreddits.
Strict candidates from large corpus	332	Posts matching romantic-context and conversation-target rules.
Second-pass screen: keep	91	Likely on-target posts used for main pilot pre-annotation.

Second-pass screen: drop	99	Likely off-target posts.
Second-pass screen: review	142	Ambiguous cases left unresolved at this stage.

Public data and ethical handling. The corpus used in this pilot study consisted of publicly accessible Reddit submissions drawn from relationship-help subreddits. The analysis was conducted at the submission level and did not rely on usernames, profile-level attributes, or attempts to identify individual posters. Because the posts concern emotionally sensitive interpersonal situations, we treated the dataset with caution and limited the reporting of example content. The paper does not reproduce verbatim sensitive posts from the corpus; where illustrative cases are provided, they are shortened, transformed, or summarized to minimize traceability. The study is exploratory and is intended to examine the operationalization of conversation-target constructs in public online text rather than to make claims about identifiable individuals.

All examples in Appendix A are transformed for illustration and are not reproduced verbatim from the source corpus.

### ***3.2 Second-Pass Screening***

After the strict extraction stage, we ran a second-pass screen. We did this because the first-pass rules were intentionally broad enough to catch likely conversation-target posts, which also meant they still admitted some off-target material. The second-pass screen separated posts into three groups: likely keep, likely drop, and review. We used lexical rules to identify obvious off-target cases such as sibling, parent, roommate, or boss problems, along with broad “need perspective” style posts that did not center on an upcoming conversation. We also used stronger keep markers

such as “how do I tell,” “should I ask,” “how do I break up,” and similar phrases that strongly imply a contemplated opening move.

The result was a workable separation: of 332 strict candidates, 91 were classified as likely on-target, 99 as likely off-target, and 142 as requiring more judgment. For the main pilot analysis reported here, we began with the 91 likely on-target cases rather than trying to resolve all 142 review cases immediately. This choice limited the pilot analysis to a cleaner subset before extending the workflow to more ambiguous cases. It let us test whether the construct was even recoverable on a cleaner subset before spending time on broader ambiguity.

### ***3.3 Coding Framework***

Our coding framework had four dimensions: objective, strategy, mismatch, and confidence. For objective, we used six categories: clarity, reassurance, repair, accountability, boundary-setting, and closure. For strategy, we used six broad categories: direct discussion, indirect or softened approach, avoidance or delay, boundary statement, reassurance-seeking approach, and confrontational approach. We defined mismatch as a case in which the planned opening appeared structurally unlikely to advance the stated objective. Confidence was coded on a three-point scale from low to high.

To clarify the operational boundaries of the coding system, Table 2 defines the objective, strategy, mismatch, and confidence categories used in the pilot workflow.

**Table 2. Operational definitions for objective, strategy, mismatch, and confidence categories.**

<b>Dimension</b>	<b>Category</b>	<b>Operational Definition</b>
<b>Objective</b>	<i>Clarity</i>	The user seeks information, explanation, or relational definition from the conversation.
<b>Objective</b>	<i>Reassurance</i>	The user seeks confirmation of safety, commitment, care, or emotional security.
<b>Objective</b>	<i>Repair</i>	The user seeks to restore connection, reduce damage, or improve the relationship after strain or conflict.
<b>Objective</b>	<i>Accountability</i>	The user seeks to disclose wrongdoing, confront wrongdoing, request responsibility, or address a violation.
<b>Objective</b>	<i>Boundary-setting</i>	The user seeks to establish limits, distance, space, or conditions for continued interaction.
<b>Objective</b>	<i>Closure</i>	The user seeks to end uncertainty, conclude the relationship, or bring the situation to a definite endpoint.
<b>Strategy</b>	<i>Direct discussion</i>	The user contemplates raising the issue explicitly through a direct question, disclosure, request, or conversation opener.
<b>Strategy</b>	<i>Indirect or softened approach</i>	The user contemplates approaching the issue cautiously, partially, or indirectly in order to reduce intensity or avoid abrupt confrontation.
<b>Strategy</b>	<i>Avoidance or delay</i>	The user contemplates waiting, withholding, postponing, or not initiating the conversation despite recognizing the issue.
<b>Strategy</b>	<i>Boundary statement</i>	The user contemplates opening with a limit, condition, refusal, or statement of personal constraint.
<b>Strategy</b>	<i>Reassurance-seeking approach</i>	The user contemplates opening in a way primarily aimed at eliciting comfort, validation, or confirmation from the partner.
<b>Strategy</b>	<i>Confrontational approach</i>	The user contemplates opening with accusation, blame, challenge, or a hostile framing likely to increase defensiveness.
<b>Mismatch</b>	<i>Aligned</i>	The contemplated opening appears broadly consistent with the user's stated or inferred primary objective.
<b>Mismatch</b>	<i>Mismatched</i>	The contemplated opening appears structurally unlikely to advance the user's stated or inferred primary objective.
<b>Mismatch</b>	<i>Unresolved</i>	Objective, strategy, or their relation cannot be assigned with enough confidence to judge alignment.
<b>Confidence</b>	<i>Low</i>	The available text provides limited evidence for assigning the label, and substantial ambiguity remains.
<b>Confidence</b>	<i>Moderate</i>	The available text provides some supporting evidence for the label, but ambiguity remains.

<b>Confidence</b>	<i>High</i>	The available text provides strong and relatively explicit evidence for the label.
-------------------	-------------	--

*Note.* Operational definitions reflect the coding framework applied during rule-based pre-annotation. Labels are pilot pre-annotations rather than validated human judgments.

These definitions were intended to support transparent first-pass pre-annotation rather than fully validated human coding. In cases where posts contained multiple possible motives or underspecified openings, the coding logic prioritized the most clearly stated objective and the most explicitly contemplated opening move; where such prioritization was not possible, the case remained unresolved.

Although the coding framework appears compact, its application is complicated by overlapping motives and incomplete self-description within individual posts. The same post can contain several motives at once. Someone may want clarity and reassurance. Someone may want accountability but also want to avoid conflict. Someone may want to repair the relationship while simultaneously preparing to leave it. That is why we treated our paper as a pilot operationalization rather than a finished annotation standard. Our goal was to see whether the categories could be recovered at all under a transparent and reproducible set of rules.

We also made an important methodological change during the project. Our earliest version of the paper assumed a manual annotation design. Once we saw the realities of the pipeline, we revised the method to match what we had actually built: a rule-based pre-annotation workflow. This distinction is important because the resulting labels should be interpreted as pilot pre-annotations rather than validated human judgments. The purpose of the workflow was to determine whether a first-pass rule system could generate informative outputs and reveal where the operationalization succeeds and fails. Illustrative transformed examples of the coding logic are provided in Appendix A.

### ***3.4 Rule-Based Pre-Annotation***

We pre-annotated the 91 likely on-target posts using weighted lexical rules that drew from both title and body. We gave title matches more weight than body matches because titles often contain the clearest summary of the user’s concern. For example, titles such as “How do I tell my girlfriend I need some space?” or “Should I tell my partner...?” already compress objective and strategy into a short form. Body text then adds detail, context, and ambiguity. The weighting therefore helped us use the most informative part of the post rather than treating every phrase as equally diagnostic.

The initial version of the rule set showed limited performance. Objective labels were frequently unresolved, strategy labels collapsed into an overly broad direct-discussion category, and mismatch was rarely computable. We then revised the objective rule set to better capture plain-language formulations such as “need some space,” “break up,” “should I tell,” “what next,” “when do I tell,” “should I ask where this is going,” “cheated,” and related expressions. We reran the pipeline after this revision and used those updated outputs as the main pilot results reported here.

### ***3.5 Analytic Strategy***

Our analysis was descriptive and exploratory. We used frequency counts to examine how many posts could be assigned to each objective, strategy, mismatch, and confidence category. We then interpreted the structure of those distributions in light of our hypotheses. Because the pilot results contained many unclear or unresolved cases, we did not apply inferential statistics beyond what the data could support. In earlier drafts, we planned chi-square tests and logistic regression. In

the final pilot analysis, however, we report the descriptive distributions and use them to evaluate the strengths and limits of the rule-based operationalization.

#### **4. Results**

The first clear result is that the target construct can be isolated, but only after several layers of narrowing. In the larger submission-level corpus, 332 posts were extracted from the first 300,001 scanned submissions. That is a small fraction of the stream. This fraction is substantively informative because it indicates that goal-strategy mismatch is not a general property of relationship-help text as a whole. It belongs to a narrower family of posts in which the author is explicitly preparing for, contemplating, or rehearsing a difficult conversation.

The second clear result is that screening and labeling are not equally difficult. Screening worked better than labeling. We were able to build a likely on-target set of 91 posts. Once we tried to assign objective, strategy, mismatch, and confidence labels, the difficulty of the construct became more visible. Our revised rule set improved objective recovery substantially compared with the earlier pass, but objective remained unresolved in many cases. Strategy was easier to recover, but only at a very broad level. Mismatch was the hardest layer of all.

On the 91 likely on-target posts, our revised pre-annotation system assigned the following objective labels: 47 unclear, 25 accountability, 11 closure, 6 clarity, 1 boundary-setting, and 1 repair. These counts are informative for two reasons. First, they show that objective is at least partially recoverable, because almost half of the corpus received a non-unclear objective label. Second, they show that our current pipeline over-recovers some objectives more easily than others. Accountability and closure were relatively visible because titles often contain confession,

wrongdoing, break-up, or disclosure language. By contrast, reassurance and repair remained rare under the current rules, and reassurance was not recovered at all in the final pass.

Strategy was easier to label, but this recoverability was achieved at a highly coarse level of categorization. Of the 91 posts, 82 were labeled direct discussion, 5 unclear, 3 confrontational, and 1 avoidance or delay. This means our broad strategy taxonomy is currently too coarse for the corpus. Many posts ask some version of “how do I tell,” “should I ask,” or “what should I say,” so a simple rule engine naturally maps them into direct discussion. That is not wrong, but it compresses too much variation. Asking for clarity, confessing wrongdoing, ending a relationship, and raising a painful issue are all different kinds of direct moves, yet our current scheme still treats them as one class.

Mismatch remained sparse. Fifty cases stayed unresolved because either objective or strategy remained too unclear for the mismatch rule to fire. Forty cases were labeled aligned, and only one case was labeled mismatched. That distribution is important because it shows the current rule system is much more comfortable declaring alignment than declaring mismatch. In other words, our mismatch rules are conservative. They under-call mismatches rather than over-call them. For a pilot, that is safer than the reverse, but it also means our current implementation is not yet strong enough to support confident mismatch-rate claims across categories.

Confidence scores tell the same story from another angle. Fifty cases received low confidence, 33 received moderate confidence, and 8 received high confidence. This means 41 of 91 cases had at least moderate support under the revised rule set, while most of the remaining cases still carried substantial ambiguity. That pattern matches what we saw qualitatively while inspecting the corpus: users often clearly signal that a conversation is coming, but they do not always

clearly state what they most want from it. They may describe pain, confusion, guilt, fear, or urgency without naming a single governing objective.

**Table 3. Objective labels in the 91-post likely on-target subset.**

Objective label	Count	Share of 91
Unclear	47	51.6%
Accountability	25	27.5%
Closure	11	12.1%
Clarity	6	6.6%
Boundary-setting	1	1.1%
Repair	1	1.1%

**Table 4. Strategy, mismatch, and confidence distributions in the 91-post likely on-target subset.**

Variable	Category	Count	Share of 91
Strategy	Direct discussion	82	90.1%
Strategy	Unclear	5	5.5%
Strategy	Confrontational	3	3.3%
Strategy	Avoidance or delay	1	1.1%
Mismatch	Unresolved / blank	50	54.9%
Mismatch	Aligned (0)	40	44.0%
Mismatch	Mismatched (1)	1	1.1%
Confidence	Low (1)	50	54.9%
Confidence	Moderate (2)	33	36.3%
Confidence	High (3)	8	8.8%

#### ***4.1 Reading the Hypotheses Against the Results***

H1 is the hypothesis with the strongest support. We were able to isolate a meaningful subset of conversation-target posts and recover at least partial objective and strategy information from a substantial share of them. H1 is therefore supported at the pilot level.

H2 receives weak support at best. Under the current mismatch rule set, mismatch did not occur at a non-trivial measurable rate. More precisely, our pilot did not recover mismatch often enough to support a strong claim. We therefore treat H2 as not yet supported under the current automated operationalization.

### **5. Discussion**

The central finding of the pilot is that the construct appears empirically meaningful but difficult to recover reliably from raw text. Our workflow can identify a narrower family of relationship-help posts that are clearly about a contemplated conversation. That is already meaningful because it shows the construct does not belong to all relationship-help narratives equally. Once we move into this narrower zone, however, a second challenge appears: users often say enough for us to know that a conversation is coming, but not enough for us to know with confidence what the governing objective is.

In that sense, our pilot confirms part of the original theory while also disciplining it. The theory behind goal-strategy mismatch assumes that people have objectives and choose openings that either fit or misfit those objectives. Our corpus suggests that this is often true in substance, but not always easy to recover from raw text. Online posters often narrate their distress in a way that mixes motives together. One person may be asking whether to confess wrongdoing, seek forgiveness, keep the relationship, and protect self-image all in the same post. Another may ask

“what should I say?” without ever naming whether they want clarity, reassurance, or closure.

This pattern suggests that the main difficulty lies less in the conceptual plausibility of the construct than in its operationalization under noisy natural language conditions. To our knowledge, this study is an early pilot attempt to jointly operationalize conversational objective, opening strategy, and their structural fit in submission-level online relationship-help narratives. Our results also show a useful asymmetry between objective and strategy. Strategy is easier to detect than mismatch because users often write in overt speech-act language: should I tell, should I ask, how do I say, how do I break up. Objective is harder because goals are often psychologically embedded rather than grammatically explicit. Mismatch is hardest because it requires both pieces at once, plus a judgment about structural fit. This layered difficulty is one of our main findings. It helps explain why generic help-seeking text can feel easy to read as a human but still remain difficult to label consistently in code.

These findings have implications for AI-mediated support systems. If our system can detect only that a user is about to have a hard conversation, then it is still at the level of “conversation awareness.” If it can also detect the governing objective, then it moves closer to decision support. If it can reliably detect mismatch, then it begins to intervene at the exact point where the user’s opening may work against the outcome they say they want. Our pilot suggests that the move from awareness to objective recovery is possible but incomplete, while the move from objective recovery to robust mismatch detection remains an open problem.

That is why we think the negative space in the results is useful. For a pilot study, these negative results are informative rather than merely residual. The fact that 47 of 91 objectives remained unclear, that 82 of 91 strategies collapsed into direct discussion, and that mismatch was

positively recovered only once indicates that stronger annotation, a finer strategy taxonomy, or a human-audited subset will be needed in future work.

## **6. Limitations**

Our paper has several important limitations. First, the main labels are rule-based pre-annotations rather than validated human annotations. They should therefore be understood as a pilot operationalization rather than as a finished gold standard. Second, the study relies on a single rule-based approach and does not include a human-coded benchmark, dual annotation, or direct comparison with alternative modeling approaches. The reported labels should therefore be interpreted as outputs of one operationalization strategy rather than as definitive measurements of the construct.

Third, the likely on-target subset was created through screening logic that still left 142 cases in unresolved review. We intentionally did not force all of those cases into the final analysis, but this also means that the reported results come from a cleaner and narrower subset rather than the entire candidate pool. As a result, the final 91-post analysis set may not fully reflect the ambiguity distribution present in the broader screened corpus.

Fourth, the strategy taxonomy remains too coarse. The dominance of direct discussion indicates that the current framework does not yet separate direct question, direct disclosure, direct request, and direct exit as clearly as it should. Fifth, the mismatch rule is conservative and underpowered. It appears more likely to under-call mismatch than to over-call it, which means the present workflow may under-capture one of the paper's central constructs.

Sixth, the weighting of title matches more heavily than body matches may introduce bias into label recovery. Titles often compress, simplify, or foreground one aspect of the problem,

whereas the body may contain a more mixed or less explicit objective structure. This may have improved recoverability for some categories while obscuring others.

Finally, the data are drawn from public relationship-help subreddits and cannot be treated as representative of all romantic conversations or all relationship-help behavior. The people who post on Reddit are already a selective population, and the posts themselves are public narratives shaped for strangers rather than complete private records of the relationship.

## **7. Conclusion**

We started with a simple question: can we detect when a person's planned opening for a hard romantic conversation does not fit what they actually want from that conversation? Our answer is cautious but useful. We can get part of the way there. We can isolate conversation-target relationship-help posts from a much larger Reddit stream. We can recover broad conversational objectives better than chance-like guessing. We can recover very coarse strategy forms even more easily. But robust mismatch detection remains harder than our first-pass rules can handle. These limitations do not negate the contribution of the pilot; rather, they clarify the current boundary of what the rule-based operationalization can recover. Our paper shows that goal-strategy mismatch is not merely a theoretical abstraction. It leaves observable traces in posts written immediately before difficult conversations. At the same time, our results show that the construct resists oversimplification. If future work wants to measure mismatch well, it will need richer annotation, finer strategy distinctions, and probably hybrid human-machine labeling rather than rules alone.

For now, our paper gives a transparent pilot workflow and a grounded empirical lesson: in online relationship-help narratives, broad conversational intention is easier to recover than strategic

misfit. For an early-stage quantitative paper, a narrower but well-supported claim is preferable to a broader claim that exceeds the evidence. We would rather end with a map of where the hard part really begins than with a claim our own results cannot support.

*Data and workflow note.* The pilot workflow was implemented in Google Colab using the Kaggle Reddit dataset `janldeboer/reddit-relationships`. Intermediate screening and pre-annotation files were saved as `streamed_screening_sheet_v2.csv`, `keep_only_91.csv`, and `preannotated_keep_91_v2.csv`.

## References

Adams, N. N. (2024). ‘Scraping’ Reddit posts for academic research? Addressing some blurred lines of consent in growing internet-based research trend during the time of Covid-19.

*International Journal of Social Research Methodology*, 27(1), 47–62.

<https://doi.org/10.1080/13645579.2022.2111816>

Artstein, R., & Poesio, M. (2008). Inter-coder agreement for computational linguistics.

*Computational Linguistics*, 34(4), 555–596. <https://doi.org/10.1162/coli.07-034-R2>

Domingue, R., & Mollen, D. (2009). Attachment and conflict communication in adult romantic relationships. *Journal of Social and Personal Relationships*, 26(5), 678–696.

Entwistle, C., Horn, A. B., Meier, T., & Boyd, R. L. (2021). Dirty laundry: The nature and substance of seeking relationship help from strangers online. *Journal of Social and Personal Relationships*, 38(12), 3472–3496. <https://doi.org/10.1177/02654075211046635>

Krippendorff, K. (2011). *Computing Krippendorff's alpha-reliability*. Annenberg School for Communication, University of Pennsylvania. Retrieved from [https://repository.upenn.edu/asc\\_papers/43](https://repository.upenn.edu/asc_papers/43)

Leo, K., Crenshaw, A. O., Hogan, J. N., Bourne, S. V., Baucom, K. J. W., & Baucom, B. R. W. (2020). A replication and extension of the interpersonal process model of demand/withdraw behavior: Incorporating subjective emotional experience. *Journal of Family Psychology*, 35(4), 534–545. <https://doi.org/10.1037/fam0000802>

Overall, N. C., & McNulty, J. K. (2017). What type of communication during conflict is beneficial for intimate relationships? *Current Opinion in Psychology*, 13, 1–5. <https://doi.org/10.1016/j.copsyc.2016.03.002>

Stewart, J. W., Bradford, K., Higginbotham, B. J., & Skogrand, L. (2016). Relationship help-seeking: A review of the efficacy and reach. *Marriage & Family Review*, 52(8), 781–803. <https://doi.org/10.1080/01494929.2016.1177631>

Tausczik, Y. R., & Pennebaker, J. W. (2010). The psychological meaning of words: LIWC and computerized text analysis methods. *Journal of Language and Social Psychology*, 29(1), 24–54.

<https://doi.org/10.1177/0261927X09351676>

Yakovleva, M. (2026). *Recurrence-aware conversation decision support in emotionally loaded romantic relationships: A conceptual framework* [Preprint].

Yakovleva, M. (2026). *An empirical bridge for recurrence-aware conversation decision support* [Preprint].

## Appendix A. Illustrative Transformed Examples of Objective, Strategy, and Mismatch Coding

The examples below are transformed and shortened for illustration. They are not reproduced verbatim from the source corpus. Their purpose is to clarify how the coding framework maps likely conversational objective, opening strategy, mismatch status, and confidence in typical conversation-target cases.

**Table A1. Illustrative transformed examples of coding logic.**

ID	Transformed Post Excerpt	Likely Objective	Likely Opening Strategy	Mismatch Status	Confidence	Rationale
A1	<i>"My boyfriend has been distant for weeks. I want to ask whether he still sees a future with me, but I do not know how to bring it up."</i>	Clarity	Direct discussion	Aligned	Moderate	The user seeks relational definition and explicitly contemplates asking directly.
A2	<i>"I cheated and need to tell my girlfriend what happened. I know I have to say it directly."</i>	Accountability	Direct disclosure / direct discussion	Aligned	High	The post centers on disclosure of wrongdoing and acceptance of responsibility.
A3	<i>"I think I need to end this relationship, but I keep hoping things will fix themselves if I wait a little longer."</i>	Closure	Avoidance or delay	Mismatched	Moderate	The stated goal is ending the relationship, but the contemplated approach is postponement.
A4	<i>"My partner crossed a line again and I need to tell him I cannot keep doing this."</i>	Boundary-setting	Boundary statement / direct discussion	Aligned	High	The post clearly indicates a limit-setting objective and a direct boundary move.
A5	<i>"I want us to repair things after our last fight, but every draft sounds like an accusation."</i>	Repair	Confrontational approach	Mismatched	Moderate	The stated goal is repair, but the contemplated opening risks defensiveness.

A6	<i>"I keep wanting to ask if she still loves me, but I am scared of sounding needy."</i>	Reassurance	Indirect or softened approach	Unresolved	Low	The user appears to seek reassurance, but the intended opening remains underspecified.
A7	<i>"I need some space and want to explain that without making him think I am ending it."</i>	Boundary-setting	Direct discussion	Aligned	Moderate	The user seeks distance rather than breakup and plans an explicit but bounded opening.
A8	<i>"Something feels wrong, but I do not know whether I want clarity, reassurance, or to leave. I just know I need to say something."</i>	Unclear	Direct discussion	Unresolved	Low	The post signals an upcoming conversation, but the governing objective is too mixed to assign confidently.

*Note.* All post excerpts are transformed to preserve construct illustration while removing identifying detail. Objective, strategy, mismatch, and confidence labels reflect rule-based pre-annotation rather than validated human coding.