Why This and Not That? A Collaborative Reflection Approach for Understanding Thought Coverage in Decision Making Support Dialog

Morita Tarvirdians **Hayley Hung** **Catharine Oertel**
TU Delft, Delft, The Netherlands
M.Tarvirdians@tudelft.nl h.hung@tudelft.nl c.r.m.m.oertel@tudelft.nl

Abstract

Conversational agents that support reflection for decision-making often rely on adaptive dialog policies that map observed user behavior to actions such as probing, deepening, or redirecting. Yet the same pattern can reflect a range of different reasons such as deliberate prioritisation or limited self-access. By modeling the observable pattern rather than the user's reason for it, current policies risk premature assumptions about the user state and inappropriate next actions. To address this gap, we introduce a human-centered method for surfacing this hidden inference step. In a user study with 62 users and 232 collaborative moments, we pause a reflection-support agent when it would normally redirect the conversation, surface its observation, and ask users to interpret the pattern and decide how to proceed. We derive a taxonomy of nine interpretation categories and show that similar reflective states can call for substantially different follow-up actions. Our findings challenge the assumption that adaptive dialog policies can rely on observable behavior alone, and suggest how user-provided interpretations can inform more appropriate conversational actions.

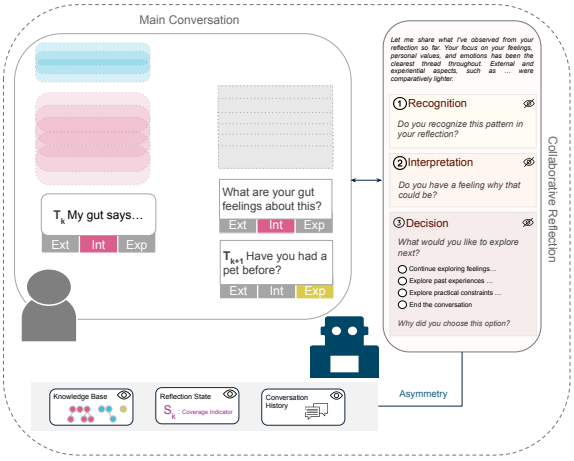


Figure 1: An illustration of a strategic decision moment in which the agent, based on observable signals (i.e., knowledge base, coverage indicator S_k, and conversational history after turn T_k), detects a coverage asymmetry across thought categories. Instead of silently acting on it, it triggers a side panel to surface its observation and collaboratively manage the decision, eliciting information not observable from dialog alone, such as whether the user recognizes the pattern, how they interpret it, and what direction they would choose for the next dialog action and why. The panel then closes, and the main conversation resumes at turn T_{k+1}, where the agent continues the dialog in the user-selected direction.

1 Introduction

Conversational agents are increasingly used to support people in thinking through complex situations, from health coaching (Kocielnik et al., 2018) to personal journaling (Kim et al., 2024) and decision-making (Park et al., 2026). These agents do more than generate contextually appropriate responses, they make *strategic decisions* about the trajectory of the conversation. At each turn, the agent decides which topic to explore, whether to probe deeper or move on, and how to frame the next question. These decisions determine which aspects of a person's situation receive attention and which do not.

Agents that manage dialog strategy face a fundamental tension, they must interpret the user's behavior to decide what to do next, but the reasons underlying these observable patterns are often not directly accessible to the agent. For example, a support agent that observes sustained focus on a single topic may interpret this as fixation and attempt to redirect the conversation elsewhere (e.g., in (Tarvirdians et al., 2026)). However, the same observable pattern may reflect avoidance of an uncomfortable topic, genuine relevance, or engagement with an issue that the user has not yet resolved. At the same time, agents that simply follow the user's stated preferences risk reinforcing blind spots; sycophantic behavior in LLMs reduces users' willingness to consider alternative perspectives and inflates their

Findings of the Association for Computational Linguistics: EMNLP 2026, pages 22495–22510
October 24-29, 2026 ©2026 Association for Computational Linguistics

confidence in existing beliefs (Sharma et al., 2024; Cheng et al., 2026). The gap between the behavioral pattern an agent observes and the user's actual reason for that pattern, what we term the *interpretive gap*, remains unstudied.

This gap is relevant for any dialog system that acts on inferred user state. Systems that select coaching strategies based on dialog state (Jörke et al., 2025), adapt questioning depth based on user responses (Kim et al., 2024), or learn tailored strategies for diverse users (Zhang et al., 2024; Zhao et al., 2025) all interpret user behavior as part of their decision process. These systems were not designed to address the interpretive gap but the gap is present whenever an agent acts on an observable pattern whose underlying cause is ambiguous. Without understanding the space of possible reasons behind such patterns, dialog strategy operates on untested assumptions. The challenge with interpreting a user's behavior accurately is that the set of possible reasons can be indeterminate due to ambiguity in the context of the situation. To account for ambiguity, we argue that an explicit dialog strategy is necessary to enable user and AI to collaboratively hone in on and disambiguate observed patterns of behavior in the user.

In this paper, we build on ReflectiMate (Tarvirdians et al., 2026), a reflection-support agent for decision making with a transparent and traceable policy. We make the agent's strategic decision moments *collaborative* — specifically, the moments where the agent would otherwise silently decide which aspect of the decision should be explored next. Rather than acting on its interpretation silently, the agent pauses and presents its observation to the user and gathers information on its recognizability and relevance to the user. Specifically, the user indicates whether they recognize the observed pattern, provides their interpretation of why it exists (their consciously aware reasons), and decides what should happen next in the dialog. Using data collected from 62 participants and 232 collaborative moments across diverse decision topics, we address the following research questions:

1. What interpretations do users provide when the agent surfaces its observations at strategic decision points, and what taxonomy of reasons emerges?

2. How do users' decisions align with or diverge from the agent's defaults, and what factors account for the differences?

2 Related Work

2.1 Dialog Strategy from Observable User Signals

Conversational agents that support coaching, reflection, or decision-making must continuously decide how to steer the conversation. These strategic decisions depend on inferring the user's state from observable patterns in the dialog. However, the same observable pattern can arise from fundamentally different user states, and current systems mainly do not account for this ambiguity.

This manifests differently across systems. GPT-Coach (Jörke et al., 2025) selects motivational interviewing strategies based on dialog state; when a user shows low motivation, the agent interprets this uniformly and applies an affirmation or open question. MindfulDiary (Kim et al., 2024) adapts follow-up questions based on what the user has shared; when a journal entry is brief, the agent interprets this as an opportunity to probe deeper. FutureMe (Dechant et al., 2025) generates contextual follow-ups to deepen reflective writing, interpreting shallow responses as warranting elaboration. ReflectiMate (Tarvirdians et al., 2026) monitors coverage across thought categories and redirects attention when it detects asymmetry, interpreting high coverage in one category as fixation.

The mapping from observable pattern to appropriate action is ambiguous, and the agent resolves this ambiguity with an implicit assumption rather than collaboratively with the user. How often these assumptions are wrong and what users' actual reasons are has not been empirically investigated.

2.2 Between Deference and Override

The problem of aligning agent behavior with user needs has been studied from two ends of a spectrum. Agents that defer to user preferences risk reinforcing blind spots: sycophantic AI reduces willingness to consider alternatives and inflates confidence in existing beliefs (Sharma et al., 2024; Cheng et al., 2026). Agents that act on their own objectives risk ignoring user diversity: one-size-fits-all strategies fail when users have different characteristics (Zhang et al., 2024), and generating strategically aligned responses remains challenging when the required strategy depends on the client's unobservable state (Başar et al., 2025; Sun et al., 2025).

Between these two ends, a small body of work explores collaborative approaches. Collaborative

management of the dialog has a long history. Plan-based models describe how participants interrupt an ongoing task to engage in a distinct *subdialogue*, negotiating a proposal or resolving a disagreement before returning to the task at hand (Litman and Allen, 1987; Lambert and Carberry, 1992). Our approach shares this move. It differs in what is exchanged. Earlier work negotiated beliefs about a shared plan, whereas we elicit the user's reason for a pattern in their own reflective behavior, which the agent has no plan-based route to infer.

Among recent systems, the closest to ours is Deliberative AI (Ma et al., 2025), which proposes a process where humans and AI deliberate together to resolve conflicting perspectives. However, their deliberation concerns the task-level decision itself (what to decide), whereas our work focuses on the agent's *strategic* decisions within the dialog (how to steer the conversation). We study what happens when the agent opens up the moments where it would normally act silently, producing empirical data about the gap between agent interpretations and user reasons.

2.3 User Modeling and the Interpretive Gap

Acting on an unobservable user state is a long-standing problem. Plan recognition infers the user's goal from observed actions (Charniak and Goldman, 1993), and POMDP-based dialog systems extend this to action selection, choosing what to say next under a distribution over user states (Williams and Young, 2007; Young et al., 2013). Recent work has modeled user emotions to improve dialog responses (Lin et al., 2023), learned user preferences across sessions (Mehri et al., 2026), and used reinforcement learning with world models that represent latent user states such as sentiment and belief (Zhao et al., 2025). These approaches go beyond simple behavioral rules, they build rich representations of user state and use them to select actions.

However, even rich user models capture *what* the user is doing or saying, not *why*. A user who disengages from a topic because it is irrelevant needs a different response than one who disengages because the topic is emotionally uncomfortable, yet both may produce the same latent representation. The "why" constitutes an interpretive layer that current models lack: the mapping from observable state to the user's reason for being in that state.

Our work addresses this gap with a collaborative approach at strategic decision moments in dialog, where the agent surfaces its observations and users provide interpretations. Rather than positing a state space and inferring over it, we characterize empirically how users interpret agent-observed patterns and select subsequent dialog actions.

3 Methodology

We describe our method in three subsections describing the reflection-support agent that provides the state-based policy, the determination of collaborative reflection moments that expose and query the policy's category-switching decisions, and the user study through which we collect users' interpretations and follow-up choices.

3.1 Baseline Reflection Support Agent

We build on ReflectiMate (Tarvirdians et al., 2026), a reflection-support agent with a transparent dialog policy. The agent maintains a structured knowledge base (KB) of the user's thoughts, organized into three categories grounded in the decision-making literature: *internal* (feelings, personal values, emotions), *external* (practical constraints, other people's opinions, real-world conditions), and *experiential* (past experiences, lessons from others).

The KB is constructed from the user's preliminary (unaided) reflection and updated after each conversational turn during assisted reflection phase. Each user utterance is parsed into one or more *thoughts*, each classified into one of the three categories based on content. The agent tracks two dimensions per category $k \in \{internal, external, experiential\}$:

- **Breadth** B_k: the number of distinct thoughts in category k.
- **Depth** D_k: the total number of follow-up expansions across thoughts in category k.

The agent computes a cumulative *coverage indicator* for each category:

$$S_k = B_k + D_k \tag{1}$$

where S_k is a heuristic measure that increases both with the number of thoughts in a category (breadth) and with the number of elaborations attached to them (depth), thus reflecting the overall focus of reflection within that category. It is updated after each turn and guides the agent's decisions.

A key objective of the ReflectiMate agent is to reduce reflection asymmetry across categories to

support more comprehensive decision-making, consistent with prior work highlighting the value of integrating affective, cognitive, and experiential perspectives (Dotlich et al., 2006).

Coverage asymmetry between categories is an observable signal for the agent, computed from the KB and the coverage indicator S_k. However, the underlying reason for this pattern is not observable. The agent therefore applies a uniform interpretation, assuming that higher-scoring categories are relatively overexplored and lower-scoring categories are underexplored. Based on this interpretation, the agent selects the category $\arg\min_k S_k$ and redirects the conversation accordingly, without consulting the user about the choice of direction. This is a strategic decision point, because the agent determines which aspect of the decision has received sufficient attention and which aspect should be further explored.

3.2 Collaborative Reflection Moments

To investigate how users interpret the reflection patterns observed by the agent (RQ1) and how these interpretations influence conversational direction choices (RQ2), we modified the ReflectiMate interaction pipeline by introducing *collaborative reflection moments*. These moments occur whenever the agent would normally make a silent strategic decision, i.e. redirecting the conversation toward another decision aspect.

At each collaborative reflection moment, the agent externalizes the reflective pattern it observes, asks the participant to interpret the pattern, and then asks how they would like to proceed. Each moment therefore produces: (1) the observable reflection state available to the agent, (2) the participant's interpretation of that state, and (3) the participant's preferred conversational direction. These moments constitute the primary unit of analysis for both RQ1 and RQ2.

Collaborative reflection moments can occur at the beginning of the assisted reflection phase, if asymmetry is already present in the user's preliminary reflection, or during the conversation when the agent's policy would normally switch categories.

Depending on the KB state, the agent presents one of two types of observations:

- **Simple Asymmetry:** one category has substantially higher coverage than the others.
- **Cross-category Mismatch:** despite prompting reflection within one category, the user's responses continue to predominantly produce thoughts from another category.

The observation text adapts to the KB state at each strategic decision point.

Following the observation, participants complete a three-step protocol:

Step 1 — Recognition. *"Do you recognize this pattern in your reflection?"* If participants disagreed with the observation, they were asked to describe the pattern they perceived instead.

Step 2 — Interpretation. *"Do you have a feeling for why that could be?"* Participants provided an open-text explanation of the observed pattern.

Step 3 — Decision. *"What would you like to explore next?"* Participants selected the next conversational direction from dynamically generated options corresponding to the available thought categories and optionally explained their choice.

The decision options were generated using descriptive labels rather than category names (e.g., "Explore your feelings, personal values, and emotions" instead of "Explore internal factors"). Option order was randomized except for "End the conversation," which was always presented last. Importantly, the agent never revealed its own default interpretation or preferred conversational direction. The protocol was refined through expert review and a pilot study with five participants using cognitive interviewing (Willis, 2005) to identify ambiguities and improve question clarity.

Each collaborative moment produced a logged data tuple consisting of: (1) the KB state at trigger time (S_k, B_k, D_k for each category), (2) the agent's default decision (logged but hidden from the participant), (3) the observation text, and (4) the participant's recognition response, interpretation, decision, and rationale.

4 Experiment

User Study Procedure. A user study was conducted on a web-based platform. After reviewing and agreeing to an informed consent form outlining the study purpose, data usage, and participant rights, participants completed a pre-study questionnaire collecting demographic information as well as the short form of the Self-Reflection and Insight Scale (SRIS) (Silvia, 2022). Participants then selected a decision topic and entered the "unaided reflection" phase, during which they freely

wrote down their thoughts without any agent support. Once they indicated that they had no further thoughts to add, they proceeded to an "assisted reflection" phase, in which they interacted with the conversational agent. This interaction continued until participants chose to end the session.

Participants. We recruited 62 participants via the crowd-sourcing platform Prolific[1]. Inclusion criteria were fluency in English and currently being in the process of making one of the big life decisions (Camilleri, 2023) listed in Appendix F.

The sample was diverse in terms of age (M = 34.27, SD = 8.66, median = 30, range 21–60), gender (30 female, 32 male), education level (ranging from high school to doctoral degrees), self-insight scores (M = 26.55, SD = 6.22, median = 26), and selected decision topics.

4.1 Interpretation Taxonomy Construction

For RQ1, the unit of analysis was each participant's open-text interpretation response produced at a collaborative decision moment.

We developed the interpretation taxonomy using an inductive qualitative coding approach (Thomas, 2006), deriving categories iteratively from participant responses. All responses were manually coded by the primary researcher. To encourage diversity during early codebook development, coding order followed maximum variation sampling (Patton, 2014), a strategy that deliberately selects cases expected to differ on a theoretically relevant dimension. Here, we conditioned on SRIS-S insight scores because self-insight may shape how users interpret and explain their reflective patterns, selecting participants from the high and low ends of the distribution.

For each response, the coder reviewed the full interaction context (agent observation, recognition response, and interpretation), assigned a provisional code, and recorded a brief rationale. After each batch of six participants, codes were consolidated into candidate categories with definitions and examples and iteratively refined through comparison with new data. Following Hennink et al. (2017), saturation was reached when six consecutive participants introduced no new categories.

Inter-rater reliability. We evaluated coding reliability using local LLMs as independent coders, following work on using LLMs as qualitative coders (Than et al., 2025), and recent work on LLM-assisted reliability assessment (Jain et al., 2025). We used models from four families (Qwen 2.5 (Hui et al., 2024), Mistral (Jiang et al., 2023), Phi-4 (Abdin et al., 2024), and Gemma 2 (Team et al., 2024)) to reduce shared-bias effects (Than et al., 2025). Each model independently coded all responses using the final codebook without access to researcher labels, and all runs were executed locally via Ollama to ensure data privacy.

Agreement between the human coder and the individual models ranged from $\kappa = 0.42$ to $\kappa = 0.50$, with the highest inter-model agreement at $\kappa = 0.58$ (Appendix B).

4.2 Reflection Trajectory Analysis

To investigate alignment between user-selected and agent-selected reflection trajectories (RQ2), we compared the participant's chosen conversational direction against the agent's logged default decision (according to S_k) at each collaborative moment.

We report overall alignment rates as well as alignment rates conditioned on interpretation category. To examine whether interpretation categories mediate the relationship between observable reflection states and trajectory choices, we additionally analyze the distribution of user-selected directions within each dominant reflection state.

Finally, we characterize the nature of trajectory divergence by classifying each user choice as either: (1) moving toward an underexplored reflection category, or (2) remaining within the dominant category. We report these distributions overall and across interpretation categories.

5 Results

5.1 Taxonomy of User Interpretations (RQ1)

From 232 collaborative reflection moments, we derived a taxonomy of 9 interpretation categories (Table 1; full taxonomy in Appendix A). Inductive coding stabilized after 18 participants, with no new categories emerging from participants 19–62.

A central finding was that similar observable reflection states often corresponded to substantially different user interpretations. When grouping moments by shared observable KB states, 28 groups contained at least two different interpretation categories, and 18 groups contained three or more categories. This indicates that the interpretive gap was pervasive rather than limited to isolated cases.

[1] https://www.prolific.com

Category	Condensed Description	Example	Distribution
Deliberate Prioritisation	Conscious judgment about what to focus on	"My personal experience is most important"	n=64 (27.6%)
Self-Diagnosed Focus	Diagnoses a personal psychological tendency	"I think when I feel negative emotions I try very hard to rationalise them rather than confront them"	n=46 (19.8%)
Emotional Driver	Emotion or need named as reason	"Guilt"	n=32 (13.8%)
Decision Nature	Nature/circumstances of the decision	"Out of necessity, to pay bills"	n=26 (11.2%)
Thinking Style	Habitual cognitive approach	"That's just the way my brain works"	n=20 (8.6%)
No Accessible Reason	Cannot articulate why the pattern exists	"I really don't know why"	n=16 (6.9%)
Interaction Attribution	Reason attributed to structure or state of the interaction itself	"I don't feel the conversation has progressed to this stage yet"	n=15 (6.5%)
Pattern Disputed	Rejects the agent's characterization	"I was focusing on experiences, not feelings"	n=9 (3.9%)
Content-Level Response	No meta-cognitive explanation; response remains at surface content level	"I know deep down that it's the right long term decision"	n=4 (1.7%)

Table 1: Taxonomy of interpretation categories (N=232), ordered by frequency, with condensed definitions, illustrative examples, and distribution of occurrences.

The most frequent category was *Deliberate Prioritisation* ($n = 64$, 27.6%), where users described their reflective focus as an intentional judgment about which aspect of the decision deserved greater attention. *Self-Diagnosed Focus* ($n = 46$, 19.8%) captured cases in which users analytically explained the source of their reflective pattern, while *Emotional Driver* ($n = 32$, 13.8%) captured cases where users primarily attributed the pattern to an emotion without further analysis. Together, these three categories accounted for 61.2% of all collaborative moments.

To further examine interpretive diversity under similar observable conditions, we analyzed moments sharing the same dominant category, observation type, and comparable coverage-score gaps. The most diverse group — characterized by a dominant internal category and a score gap of approximately 10 — contained 19 moments spanning 8 of the 9 taxonomy categories. Interpretations ranged from *"it's nice to talk about my feelings"* (*Emotional Driver*) to *"just the way I am"* (*Thinking Style*) to *"I really don't know"* (*No Accessible Reason*).

Despite this diversity, the agent's policy would have treated all 19 moments identically, as they corresponded to the same observable reflection state from the agent's perspective.

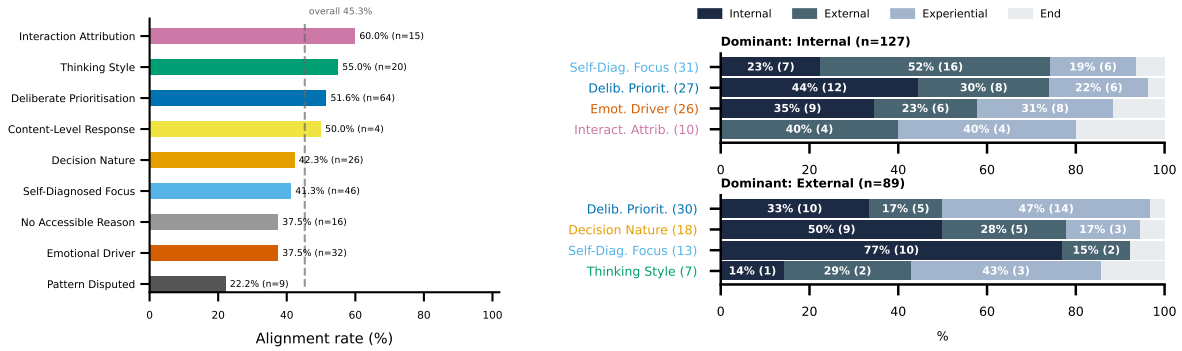
5.2 Alignment and User Choices (RQ2)

Users selected the same action as the agent's default in 105 of 232 moments (45.3%). In the remaining cases, users selected a different direction than the coverage-based policy.

Divergence does not imply resistance to exploration. Importantly, disagreement with the agent did not mean users stayed with what they had already explored. Among the 127 misaligned moments, 45.7% moved toward an underexplored dimension (albeit not the one selected by the agent), 41.7% remained within the dominant dimension, and 12.6% ended the conversation. Across moments in which the conversation continued, 74.4% (160/215) of user choices moved toward a dimension other than the current dominant one. The agent's observation therefore influenced *whether* users broadened their exploration, while the direction of broadening varied.

Interpretation is associated with divergence patterns. Alignment rates varied across taxonomy categories, ranging from 22.2% to 60.0% (Figure 2a). This indicates that the likelihood of following the agent's default decision depends in part on how users interpret the observed reflection pattern.

Figure 2b further shows that user choices differ systematically within the same observable state depending on interpretation category. For exam-



(a) Alignment rate by taxonomy category. Dashed line = overall rate (45.3%).

(b) User choice direction by taxonomy category, split by dominant KB state. Row labels give the number of moments in each category; segment labels give percentage and raw count. The four most frequent categories are shown in each panel.

Figure 2: Alignment and divergence behavior across taxonomy categories.

ple, when internal aspects dominated reflection ($n = 127$), users attributing the pattern to *Self-Diagnosed Focus* most frequently shifted toward external aspects (52%), whereas those attributing it to *Deliberate Prioritisation* more often continued exploring internal aspects (44%). When external aspects were dominant ($n = 89$), users attributing the pattern to *Self-Diagnosed Focus* never selected a shift toward experiential aspects, whereas other interpretation categories did.

Overall, these results suggest that the same observable reflection state can lead to different trajectory choices depending on the user's interpretation, which is not accessible to the agent's policy.

5.3 Exploratory Analyses

Recognition and observation type. Users recognized the agent's observation in 89.2% of collaborative moments. Recognition was universal for simple asymmetry cases, but all non-recognition cases (25 moments) occurred in cross-category mismatch observations.

Notably, lack of recognition did not prevent users from providing interpretations: 17 of 25 such cases still contained explicit explanations, including *Deliberate Prioritisation* (5/25) and *Self-Diagnosed Focus* (3/25). Only a minority explicitly attributed the mismatch to incorrect pattern detection.

Preliminary indicators of interpretation patterns. We explored whether observable pre-interaction signals may be associated with interpretation categories. Users with longer preliminary reflections were more likely to produce *Self-Diagnosed Focus* interpretations (33.8% in the highest tertile vs. 17.1% in the lowest) and less likely to produce *Emotional Driver* responses (5.2% vs. 17.1%; $\chi^2 = 24.8$, $df = 16$, $p = .073$, $V = .23$).

A similar but weaker trend was observed for SRIS insight scores: *Emotional Driver* interpretations were associated with lower insight scores ($M = 23.4$) than *Self-Diagnosed Focus* interpretations ($M = 26.1$, $p = .14$).

These results are exploratory and should be interpreted as hypothesis-generating. They suggest that pre-interaction signals such as reflection length and insight may be weakly associated with how users interpret reflection patterns.

6 Discussion

In this paper we investigated what interpretations users give when an agent surfaces the reflection patterns it has observed, and how those interpretations relate to the next dialog action they choose: whether to redirect, or stay with the current focus.

6.1 A Taxonomy of User Reasons Accounting for Coverage Asymmetries

To answer **RQ1**, we examined the reasons users give when an agent makes an observed asymmetry in their reflective behavior explicit. From these responses, we derive a taxonomy of nine interpretation categories (cf. Table 1), showing that the same behavioral pattern can reflect structured and recurring differences in user reasons. This opens a new modeling direction for adaptive dialog policy: agents can represent likely reason categories be-

hind observable patterns, enabling them to choose follow-up strategies that better fit the user's situation and the task. For instance, a follow-up strategy should be different if the reason for coverage asymmetry lies in a *deliberate prioritization* in contrast to an *emotional driver*.

We also provide empirical evidence for the existence of the interpretation gap between the agent and the user. For example, a single observable agent state, i.e. the same pattern in the knowledge-base representation that the agent uses to make strategic decisions, was interpreted in eight different ways by users. Therefore, dialog systems that select dialog trajectories solely based on observable signals, such as latent representations of inferred user beliefs (Zhao et al., 2025), latent representations of user coverage patterns (Tarvirdians et al., 2026), or LLMs that prematurely lock into assumptions about user intent (Laban et al., 2026), risk offering uniform solutions to fundamentally different underlying user states. While (Zhang et al., 2024) argue that one-size-fits-all strategies fail for diverse users, our findings suggest that one mechanism underlying such failures is that the same behavior can arise from different reasons, such as deliberate focus, emotional states, cognitive habits, situational constraints, and properties of the interaction itself.

6.2 Alignment and Divergence Between User Choices and State-Based Policy Actions

In **RQ2**, we investigated users' decisions regarding the reflection trajectory and the alignment between their choices and the trajectory the agent would have selected. A key finding was that, when users were informed of the pattern the agent observed in their reflection and asked to interpret it, 74.4% of choices moved toward an underexplored dimension of their decision, without requiring the agent to act unilaterally or explicitly reveal its default choice. Yet in 54.7% of moments users chose a different direction than the agent's default. Users broadened their exploration, but the direction was informed by reasons the agent could not have inferred from its state representation. In other words, although the agent could identify potentially underexplored regions of reflection, users' interpretations shaped how they chose to engage with those regions.

6.3 A Hybrid Intelligence Approach to Dialog Strategy

Our design and findings suggest a hybrid intelligence approach to dialog strategy: the agent does not treat its interpretation of user behavior as sufficient for action, but surfaces the observed pattern and asks the user to help interpret it before the next dialog move is selected. The agent contributes structural awareness of the reflection process; the user contributes self-knowledge about the reason behind the pattern. This is not merely a way to make the agent more agreeable: surfacing the pattern can also create meta-awareness for the user, helping them recognize how their own reflection is unfolding and giving them a stronger role in steering the conversation.

This perspective differs from approaches that emphasize collaboration at the level of task outcomes, such as joint decision-making (Ma et al., 2025), whereas we instead propose focusing on collaboration at the level of dialog strategy. In addition, our approach lies between two dialog strategy paradigms in prior work, systems that solely follow user preferences, the sycophantic agents (Sharma et al., 2024; Cheng et al., 2026), and systems that follow their own interaction objectives (Tarvirdians et al., 2026).

Together with the taxonomy, our proposed approach opens up opportunities for designing future agents that reason over multiple plausible interpretations of user states rather than collapsing them into a single assumption. Such agents could surface uncertainty when necessary and empower users in gaining further awareness in the interpretation process itself.

7 Future Work

The taxonomy maps the interpretive space that adaptive dialog policies would need to account for. This paper characterizes that space and does not evaluate whether acting on it improves dialog outcomes. Establishing that requires a controlled comparison against an agent that steers silently, together with an account of what counts as improvement. For big life decisions, for example, decision quality cannot be established within a single reflection session, since outcomes unfold over months or years, so evaluation would rely on process-level proxies. This line of experimentation is an important direction for future work.

Beyond this comparison, future work could condition action selection on inferred reason types rather than on observable state alone. The observed association between signals (e.g., reflection length, insight scores) and interpretation types further sug-

gests that partial adaptation without explicit exchange may be feasible, though this remains preliminary and requires validation on larger samples. Finally, extending the taxonomy to other support domains would further test its generality beyond the present setting.

8 Conclusion

We presented the first empirical characterization of the interpretive gap in support dialog, the gap between the behavioral patterns an agent observes and the reasons underlying them. By surfacing agent observations to 62 users across 232 strategic decision moments, we derive a 9-category taxonomy of user reasons, showing that identical observable states can reflect deliberate choices, emotional states, cognitive habits, situational constraints, or interactional structure. Users diverged from the agent's default action in 54.7% of cases, yet 74.4% of continuing choices shifted toward an underexplored dimension, suggesting disagreement concerned direction rather than integration. These findings highlight that dialog strategy in support agents operates over inherently ambiguous signals, and that structured agent–user exchange reveals information neither side can access alone.

Limitations

Several limitations constrain the generalizability of our findings. The taxonomy was derived from a single agent in a single domain (personal decision-making reflection), and while it captures general cognitive phenomena such as deliberate prioritisation, emotional influence, and habitual thinking, their relative distribution may differ in other contexts such as therapeutic dialog or educational tutoring.

Our double-coding approach used local LLMs rather than a second human coder. While inter-model agreement supports codebook clarity, moderate kappa values reflect the inherent difficulty of distinguishing nuanced categories such as Emotional Driver and Self-Diagnosed Focus, a boundary that also challenges automated classification.

User-reported reasons are self-reports and may reflect post-hoc rationalization rather than actual causes of behavior (Nisbett and Wilson, 1977). In addition, the study protocol may have influenced responses by prompting users to recognize patterns that might not have emerged spontaneously.

Ethical Considerations

This study was approved by the Human Research Ethics Committee of Delft University of Technology. All participants provided informed consent before participation and were compensated in accordance with Prolific's fair payment policies. Given the sensitivity of personal decision-making topics, we took several precautions. When participants requested opinions or advice, the agent explicitly stated it could not provide any. All language models used for response generation and qualitative coding were deployed locally to ensure that reflection data remained on secure institutional servers and was not transmitted to third-party APIs. Data were pseudonymised, and personally identifiable information was stored separately from reflection content; participants were informed of data handling and retention practices during consent.

Gen AI Declaration

Generative AI tools were used to assist with copy-editing and coding, but the authors are fully accountable for the content and ideas.

Acknowledgments

We thank Catholijn M. Jonker for her continued support of this line of work. We also gratefully acknowledge the support of the Delft AI Initiative and the Designing Intelligence (DI) lab.

This research was partially funded by the Hybrid Intelligence Center, a 10-year programme funded by the Dutch Ministry of Education, Culture and Science through the Netherlands Organisation for Scientific Research, Hybrid Intelligence Centre, grant number 024.004.022. It has also been partially supported by the project RISE: Relapse Intervention and Self-Management, file number KICH1.GZ03.21.002 of the research programme KIC - MISSIE Zorg in eigen leefomgeving 2021, partly financed by the Dutch Research Council (NWO).

References

Marah Abdin, Jyoti Aneja, Harkirat Behl, Sébastien Bubeck, Ronen Eldan, Suriya Gunasekar, Michael Harrison, Russell J. Hewett, Mojan Javaheripi, Piero Kauffmann, James R. Lee, Yin Tat Lee, Yuanzhi Li, Weishung Liu, Caio C. T. Mendes, Anh Nguyen, Eric Price, Gustavo de Rosa, Olli Saarikivi, and 8 others. 2024. Phi-4 technical report. *Preprint*, arXiv:2412.08905.

Erkan Başar, Xin Sun, Iris Hendrickx, Jan de Wit, Tibor Bosse, Gert-Jan De Bruijn, Jos A Bosch, and Emiel Krahmer. 2025. How well can large language models reflect? a human evaluation of llm-generated reflections for motivational interviewing dialogues. In *Proceedings of the 31st international conference on computational linguistics*, pages 1964–1982.

Adrian R Camilleri. 2023. An investigation of big life decisions. *Judgment and Decision Making*, 18:e32.

Eugene Charniak and Robert P Goldman. 1993. A bayesian model of plan recognition. *Artificial Intelligence*, 64(1):53–79.

Myra Cheng, Cinoo Lee, Pranav Khadpe, Sunny Yu, Dyllan Han, and Dan Jurafsky. 2026. Sycophantic AI decreases prosocial intentions and promotes dependence. *Science*, 391(6792):eaec8352.

Martin Dechant, Eva Lash, Sarah Shokr, and Ciarán O'Driscoll. 2025. Future me, a prospection-based chatbot to promote mental well-being in youth: two exploratory user experience studies. *JMIR Formative Research*, 9(1):e74411.

David L Dotlich, Peter C Cairo, and Stephen H Rhinesmith. 2006. *Head, Heart and Guts: How the world's best companies develop complete leaders*. John Wiley & Sons.

Monique M Hennink, Bonnie N Kaiser, and Vincent C Marconi. 2017. Code saturation versus meaning saturation: how many interviews are enough? *Qualitative health research*, 27(4):591–608.

Binyuan Hui, Jian Yang, Zeyu Cui, Jiaxi Yang, Dayiheng Liu, Lei Zhang, Tianyu Liu, Jiajun Zhang, Bowen Yu, Keming Lu, Kai Dang, Yang Fan, Yichang Zhang, An Yang, Rui Men, Fei Huang, Bo Zheng, Yibo Miao, Shanghaoran Quan, and 5 others. 2024. Qwen2.5-coder technical report. *Preprint*, arXiv:2409.12186.

Nilesh Jain, Hyungil Suh, Seyi Adeyinka, Leor Roseman, and Aza Allsop. 2025. Multi-llm thematic analysis with dual reliability metrics: Combining cohen's kappa and semantic similarity for qualitative research validation. *arXiv preprint arXiv:2512.20352*.

Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Lélio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2023. Mistral 7b. *Preprint*, arXiv:2310.06825.

Matthew Jörke, Shardul Sapkota, Lyndsea Warkenthien, Niklas Vainio, Paul Schmiedmayer, Emma Brunskill, and James A Landay. 2025. Gptcoach: towards llm-based physical activity coaching. In *Proceedings of the 2025 CHI conference on human factors in computing systems*, pages 1–46.

Taewan Kim, Seolyeong Bae, Hyun Ah Kim, Su-woo Lee, Hwajung Hong, Chanmo Yang, and Young-Ho Kim. 2024. Mindfuldiary: harnessing large language model to support psychiatric patients' journaling. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, pages 1–20.

Rafal Kocielnik, Lillian Xiao, Daniel Avrahami, and Gary Hsieh. 2018. Reflection companion: a conversational system for engaging users in reflection on physical activity. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2(2):1–26.

Philippe Laban, Hiroaki Hayashi, Yingbo Zhou, and Jennifer Neville. 2026. Llms get lost in multi-turn conversation. In *International Conference on Learning Representations*.

Lynn Lambert and Sandra Carberry. 1992. Modeling negotiation subdialogues. In *30th Annual Meeting of the Association for Computational Linguistics*, pages 193–200.

Hsien-Chin Lin, Shutong Feng, Christian Geishauser, Nurul Lubis, Carel van Niekerk, Michael Heck, Benjamin Ruppik, Renato Vukovic, and Milica Gasić. 2023. Emous: Simulating user emotions in task-oriented dialogues. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2526–2531.

Diane J Litman and James F Allen. 1987. A plan recognition model for subdialogues in conversations. *Cognitive science*, 11(2):163–200.

Shuai Ma, Qiaoyi Chen, Xinru Wang, Chengbo Zheng, Zhenhui Peng, Ming Yin, and Xiaojuan Ma. 2025. Towards human-ai deliberation: Design and evaluation of llm-empowered deliberative ai for ai-assisted decision-making. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, pages 1–23.

Shuhaib Mehri, Priyanka Kargupta, Tal August, and Dilek Hakkani-Tür. 2026. Multisessioncollab: Learning user preferences with memory to improve long-term collaboration. *arXiv preprint arXiv:2601.02702*.

Richard E Nisbett and Timothy D Wilson. 1977. Telling more than we can know: Verbal reports on mental processes. *Psychological review*, 84(3):231.

Jeongeon Park, Bryan Min, Kihoon Son, Jean Y Song, Xiaojuan Ma, and Juho Kim. 2026. Choicemates: Supporting unfamiliar online decision-making with multi-agent conversational interactions. In *Proceedings of the 31st International Conference on Intelligent User Interfaces*, pages 1526–1550.

Michael Quinn Patton. 2014. *Qualitative research & evaluation methods: Integrating theory and practice*. Sage publications.

Mrinank Sharma, Meg Tong, Tomasz Korbak, David Duvenaud, Amanda Askell, Samuel R. Bowman, Newton Cheng, Esin Durmus, Zac Hatfield-Dodds, Scott R. Johnston, Shauna Kravec, Timothy Maxwell, Sam McCandlish, Kamal Ndousse, Oliver Rausch, Nicholas Schiefer, Da Yan, Miranda Zhang, and Ethan Perez. 2024. Towards understanding sycophancy in language models. In *International Conference on Learning Representations*.

Paul J Silvia. 2022. The self-reflection and insight scale: Applying item response theory to craft an efficient short form. *Current Psychology*, 41(12):8635–8645.

Xin Sun, Xiao Tang, Abdallah El Ali, Zhuying Li, Pengjie Ren, Jan de Wit, Jiahuan Pei, and Jos A Bosch. 2025. Rethinking the alignment of psychotherapy dialogue generation with motivational interviewing strategies. In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 1983–2002.

Morita Tarvirdians, Senthil Chandrasegaran, Hayley Hung, Catholijn M Jonker, and Catharine Oertel. 2026. Reflecti-mate: A conversational agent for adaptive decision-making support through system 1 and system 2 thinking. In *Proceedings of the 34th ACM Conference on User Modeling, Adaptation and Personalization*, pages 213–222.

Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, Johan Ferret, Peter Liu, Pouya Tafti, Abe Friesen, Michelle Casbon, Sabela Ramos, Ravin Kumar, Charline Le Lan, Sammy Jerome, and 179 others. 2024. Gemma 2: Improving open language models at a practical size. *Preprint*, arXiv:2408.00118.

Nga Than, Leanne Fan, Tina Law, Laura K Nelson, and Leslie McCall. 2025. Updating "the future of coding": Qualitative coding with generative large language models. *Sociological Methods & Research*, 54(3):849–888.

David R Thomas. 2006. A general inductive approach for analyzing qualitative evaluation data. *American journal of evaluation*, 27(2):237–246.

Jason D Williams and Steve Young. 2007. Partially observable markov decision processes for spoken dialog systems. *Computer Speech & Language*, 21(2):393–422.

G.B. Willis. 2005. *Cognitive Interviewing: A Tool for Improving Questionnaire Design*. SAGE Publications.

Steve Young, Milica Gašić, Blaise Thomson, and Jason D Williams. 2013. Pomdp-based statistical spoken dialog systems: A review. *Proceedings of the IEEE*, 101(5):1160–1179.

Tong Zhang, Chen Huang, Yang Deng, Hongru Liang, Jia Liu, Zujie Wen, Wenqiang Lei, and Tat-Seng Chua. 2024. Strength lies in differences! improving strategy planning for non-collaborative dialogues via diversified user simulation. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 424–444.

Yue Zhao, Xiaoyu Wang, Dan Wang, Zhonglin Jiang, Qingqing Gu, Teng Chen, Ningyuan Xi, Jinxian Qu, Yong Chen, and Luo Ji. 2025. Dream to chat: Model-based reinforcement learning on dialogues with user belief modeling. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 4764–4781, Suzhou, China. Association for Computational Linguistics.

A Taxonomy Codebook

Table 2 shows the full taxonomy codebook with detailed definitions and illustrative examples.

B LLM Double-Coding Setup

B.1 Prompt Template

The prompt 3 was used for double-coding. Placeholders `{observation_text}` and `{interpretation}` are filled with the agent observation and the participant's interpretation for each moment.

B.2 Evaluation

Four models were evaluated:

- Qwen 2.5 14B
- Mistral Small
- Phi-4
- Gemma 2

Table 3 reports inter-rater reliability summary, Cohen's κ.

C Per-Category IRR

Table 4 shows agreement rates between the human coder and Mistral Small (the best-performing model, $\kappa = 0.50$, $n = 199$ double-coded moments). n is the number of moments the human coder assigned to each category. Agreement % is computed as the proportion of those moments where Mistral agreed. *Most common confusion* reports the category Mistral most often assigned instead, with count in parentheses.

D Human–Model Confusion Matrix

Table 5 illustrates full 9×9 confusion matrix for human coder vs. Mistral Small on the $n = 199$ double-coded sample. Rows represent the human (true) label; columns represent the Mistral (predicted) label.

E Taxonomy Distribution

Table 6 shows the distribution of the 232 coded pause moments across the nine taxonomy categories, with sub-breakdowns by moment position (opening vs. mid-session) and agent alignment.

F Decision Topic Distribution

Distribution of the 13 decision topics represented in the sample, with counts of participants and pause moments, and the most frequent taxonomy category within each topic is shown in Table 7.

```
Classify this user's interpretation into exactly one category.

Categories:
1. Deliberate Prioritisation --- User made a conscious judgment about where to
    focus. Example: "My feelings shouldn't outweigh what I can provide for the
    bird."
2. Self-Diagnosed Focus --- User analyses HOW a personal difficulty drives the
    pattern. Must contain causal reasoning (because, makes me, stems from, leads
    to). Example: "My anxiety has a massive role and it makes me focus on feelings
    ."
3. Emotional Driver --- User names an emotion or emotional need as the reason,
    WITHOUT explaining how it operates. Example: "Guilt." / "I want to be happier
    ."
4. Decision Nature --- The decision's situation or circumstances explain the focus
    , not the user's psychology. Example: "It's a financial decision." / "Social
    pressures are the driving factors."
5. Thinking Style --- A stable personal trait or habit explains the focus. Example
    : "That's how my brain works." / "I am an empathetic individual."
6. Interaction Attribution --- The agent's questions or the study design caused
    the pattern. Example: "The chatbot guided me there."
7. No Accessible Reason --- User tried to explain but couldn't. Example: "I'm not
    sure." / "Not really."
8. Content-Level Response --- User gave decision content instead of interpreting
    the pattern. Example: "I have started a business before."
9. Pattern Disputed --- User rejects the agent's observation as inaccurate.
    Example: "You misinterpreted my thoughts."

Key distinction: Self-Diagnosed Focus requires causal/mechanistic reasoning about
    oneself. Emotional Driver just names the emotion. If unsure between these two,
     check: does the interpretation explain HOW the emotion shapes the pattern, or
     just NAME the emotion?

Agent's observation: {observation_text}
User's interpretation: {interpretation}

Respond in exactly this format:
CATEGORY: [category name]
REASONING: [one sentence]
```

Figure 3: LLM classification prompt template for taxonomy categorization.

Table 2: Full Taxonomy Codebook

Category	Definition	Representative Examples
Deliberate Prioritisation	The user has made a conscious judgment about which dimension of their decision deserves more or less reflective attention. The pattern reflects a reasoned assessment of relevance, importance, or usefulness — the user explains what they chose to focus on or not focus on, and why.	*"Practicality is important when considering the welfare of a pet"* (ID 280) *"I don't have past experiences with wanting a child, this is hypotheical and only about my feelings"* (ID 67)
Self-Diagnosed Focus	The user has analytically identified a personal difficulty, tendency, or internal state that drives their reflective focus. They go beyond naming an emotion to explaining the mechanism — how a personal characteristic shapes their pattern. There is a diagnostic quality: the user is reasoning about their own psychology (System 2 reasoning about System 1 content).	*"I focus a lot on feelings and how other people think because I worry about their perception of me"* (ID 212) *"I feel like I have so much going on in my own head I sometimes can't take other peoples opinions on board."* (ID 128)
Emotional Driver	An emotion or unresolved emotional need explains the user's reflective focus. The emotion either pulls the user toward a dimension (e.g., sadness pulls toward feelings) or pushes them toward one (e.g., guilt pushes toward external obligations). The user names the emotion or emotional need as the reason without extensive analysis of how it operates (System 1 speaking).	*"insecurity"* (ID 167) *"Childhood trauma"* (ID 191)
Decision Nature	The decision's inherent nature, the user's real-world circumstances, or their stage in the decision process explains the reflective focus. The pattern reflects what the decision IS or WHERE the user is in the decision process — not the user's psychology, emotions, or cognitive style.	*"out of necessity"* (ID 135) *"You have to be practical in being able to live on the money you earn"* (ID 246)
Thinking Style	The user's habitual way of approaching problems or decisions explains the reflective focus. The pattern reflects a stable personal characteristic — how they think generally — not something specific to this decision or their current emotional state.	*"just the way I am"* (ID 295) *"I am fact / risk driven in decision making"* (ID 301)
Interaction Attribution	The user attributes the observed coverage pattern to the agent's questions, the conversational structure, the study design, or their expectations about the interaction format — rather than to their own reflective tendencies or the nature of their decision.	*"The bot encouraged me to open up and think about my feelings in a constructive way."* (ID 282) *"it's just the way the conversation has flowed"* (ID 206)
No Accessible Reason	The user recognises the observed pattern but cannot articulate why it exists. They express genuine uncertainty, possibly offering vague speculation without conviction. The reason for the pattern is not consciously available to the user. The user attempted to interpret but could not.	*"not sure"* (ID 126) *"Not really no"* (ID 169)
Content-Level Response	The user did not attempt to interpret the pattern. Instead of explaining why their reflection has a particular distribution, they responded with decision content — narrating their situation, adding new information about their decision, or continuing to reflect on the decision topic. The metacognitive step of interpreting one's own reflective process did not occur.	*"I have started a business before"* (ID 308) *"I didn't mention that its a lot harder for me to conceive because of my health conditions, I have…"* (ID 227)
Pattern Disputed	The user rejects the agent's characterisation of their reflection. They either see a different pattern than what the agent described, believe the agent miscategorised their responses, or feel the observation doesn't accurately represent their thinking.	*"I was focusing on experiences, not feelings"* (ID 210) *"I do feel like I've reflected deeply from an emotional standpoint as well as a practical one"* (ID 139)

Table 3: IRR κ across all evaluated configurations.

Model A	**Model B**	**Human–A κ**	**Human–B κ**	**A–B κ**
Qwen 2.5 14B	Mistral Small	0.46	0.50	0.54
Phi-4	Gemma 2	0.46	0.42	0.58

Table 4: Per-category human–Mistral agreement

Category	***n***	**Agree %**	**Most common confusion**
Content-Level Response	1	0%	Decision Nature (n=1)
Decision Nature	22	18%	Deliberate Prioritisation (n=12)
Deliberate Prioritisation	61	74%	Decision Nature (n=5)
Emotional Driver	28	79%	Deliberate Prioritisation / Self-Diagnosed Focus (n=3)
Interaction Attribution	11	82%	Decision Nature / Deliberate Prioritisation (n=1)
No Accessible Reason	11	36%	Pattern Disputed (n=4)
Pattern Disputed	7	71%	Deliberate Prioritisation / Interaction Attribution (n=1)
Self-Diagnosed Focus	42	55%	Deliberate Prioritisation (n=10)
Thinking Style	16	44%	Deliberate Prioritisation / Emotional Driver / Self-Diagnosed Focus (n=3)

Table 5: Confusion matrix: human coder (rows) vs. Mistral Small (columns). Diagonal cells (shaded) indicate agreement.

Human \ Mistral	**CL**	**DN**	**DP**	**ED**	**IA**	**NA**	**PD**	**SF**	**TS**
CL	0	1	0	0	0	0	0	0	0
DN	0	**4**	12	3	0	1	1	1	0
DP	0	5	**45**	4	4	1	2	0	0
ED	0	0	3	**22**	0	0	0	3	0
IA	0	1	1	0	**9**	0	0	0	0
NA	0	1	2	0	0	**4**	4	0	0
PD	0	0	1	0	1	0	**5**	0	0
SF	0	0	10	6	0	2	0	**23**	1
TS	0	0	3	3	0	0	0	3	**7**

CL = Content-Level Response. DN = Decision Nature. DP = Deliberate Prioritisation. ED = Emotional Driver. IA = Interaction Attribution. NA = No Accessible Reason. PD = Pattern Disputed. SF = Self-Diagnosed Focus. TS = Thinking Style.

Table 6: Full taxonomy distribution with position and alignment breakdown.

Category	***n***	**%**	**Opening**	**Mid-session**	**Aligned**	**Misaligned**
Deliberate Prioritisation	64	27.6%	14	50	33 (52%)	31 (48%)
Self-Diagnosed Focus	46	19.8%	18	28	19 (41%)	27 (59%)
Emotional Driver	32	13.8%	7	25	12 (38%)	20 (62%)
Decision Nature	26	11.2%	10	16	11 (42%)	15 (58%)
Thinking Style	20	8.6%	3	17	11 (55%)	9 (45%)
Interaction Attribution	15	6.5%	2	13	9 (60%)	6 (40%)
No Accessible Reason	16	6.9%	2	14	6 (38%)	10 (62%)
Content-Level Response	4	1.7%	3	1	2 (50%)	2 (50%)
Pattern Disputed	9	3.9%	3	6	2 (22%)	7 (78%)
Total	232	100.0%	62	170	105 (45%)	127 (55%)

Table 7: Decision topics by participant count, moment count, and dominant taxonomy category.

Decision Topic	**Participants**	**Moments**	**Most common category**
Start a new job/position	16	61	Deliberate Prioritisation
Quit a job/position	8	30	Self-Diagnosed Focus
Start a new business	6	25	Deliberate Prioritisation
Get a pet	5	20	Deliberate Prioritisation
Move to a new city	5	18	Decision Nature
Have/adopt a child	4	13	Decision Nature
Pursue a degree	4	16	Deliberate Prioritisation
End a romantic relationship	3	9	Deliberate Prioritisation
Buy a home	3	11	Deliberate Prioritisation
Get married	2	6	Pattern Disputed
Move to a new country	2	6	Self-Diagnosed Focus
Care for a family member	2	8	Thinking Style
Begin a romantic relationship	2	9	Emotional Driver