

The Curated Interaction Model: A Client-Side Framework for Orchestrating Prosocial Real-World Conversation

Fernando Reyes
Independent Researcher
fer@beautifulandevil.com

Modern social technologies often paradoxically inhibit direct, meaningful human connection. We address this gap by proposing the Curated Interaction Model, a novel client-side framework designed to orchestrate engaging, face-to-face conversations among small groups of young adults (ages 21–29) in a controlled social setting. The model employs a lightweight, three-item psychometric classifier to assign participants to one of six empirically derived conversational profiles. Based on the specific compositional dynamics of a five- to six-person group, the framework’s orchestration algorithm selects and blends curated, versioned collections of discussion prompts (“decks”). This process adaptively guides the group toward balanced, high-quality dialogue. Critically, all orchestration logic executes on-device (edge-first), a design choice that minimizes cloud dependency, reduces interactional latency, and inherently preserves user privacy. The model’s architecture is grounded in established social-psychological principles, including optimal dissonance, conversational flow, and social identity theory, to foster emergent communion. In this paper, we formalize the orchestration algorithm, provide a rigorous terminological framework, and present extended system visualizations, including client-server data flow and a session state transition diagram. We offer implementation sketches in Swift/Objective-C and delineate a comprehensive evaluation plan, complete with statistical power considerations. We conclude by proposing empirical validation pathways (simulation and pilot studies) and discussing the model’s potential for generalization beyond the initial demographic cohort.

Keywords: conversational curation, small-group dynamics, human-computer interaction, edge computing, flow theory, optimal dissonance, social identity, privacy-by-design

Introduction

In an era of ubiquitous digital connectivity, the capacity for spontaneous, meaningful face-to-face conversation among strangers has perceptibly atrophied. Initial social encounters frequently fail to progress beyond superficial pleasantries, a phenomenon we term *social conversational friction*. Prevailing technological solutions often exacerbate this issue. Segmentation-driven matching platforms algorithmically corral individuals into homogenous silos, limiting exposure to diverse perspectives. Gamified social mechanics, while stimulating, often prioritize transient engagement over genuine connection. These systems typically operate on a paradigm of identity-based personalization, matching users on explicit attributes (e.g., interests, demographics) rather than on the implicit, stylistic dynamics of how they communicate.

This paper challenges that paradigm. We conceptualize conversation not as a byproduct of a successful match, but as the product itself: a structured, yet adaptive, experience. We introduce the Curated Interaction Model (CIM), a framework that orchestrates small-group dialogue by tuning a sequence of prompts to the group’s emergent conversational *style*—defined by its collective pace, content preferences, and stance during disagreement—rather than to the static identity attributes of its members. Our central thesis is that by modulating the thematic and tonal properties of conversational prompts in response to a group’s compositional dynamics, it is possible to predictably accelerate the formation of rapport, reduce social anxiety, and promote more equitable participation (i.e., balanced airtime) within a 60–75 minute session.

Primary Contributions

This work makes several contributions to the fields of Human-Computer Interaction (HCI), groupware design, and applied social psychology:

1. A Novel Taxonomy of Conversational Archetypes:

Correspondence concerning this article should be addressed to Fernando Reyes, Independent Researcher. Email: fer@beautifulandevil.com

We propose and detail a six-profile taxonomy derived from a minimal, three-item questionnaire designed to infer interactional style.

2. **A Group-Aware Orchestration Algorithm:** We formalize a dynamic blending algorithm that computes prompt-deck weights based on the group's profile composition and adapts in-session based on implicit signals of engagement or fatigue.
3. **A Privacy-Preserving, Edge-First Architecture:** We present a client-side system architecture that executes all core logic on-device, utilizing versioned, immutable content catalogs and requiring minimal server state, thereby enhancing privacy and system resilience.
4. **A Rigorous Methodological Framework:** We provide extended system visualizations, a comprehensive evaluation plan including statistical power analysis, and concrete pathways for empirical validation through both simulation and pilot studies.
5. **An Ethics-by-Design Approach:** We outline a framework for prompt curation that incorporates programmatic guardrails for psychological safety, ensuring inclusive and low-risk social interaction.

Related Work

Our research is situated at the intersection of three domains: the social psychology of group dynamics, computational dialogue management in HCI, and privacy-centric system design.

Social Dynamics and Small Group Facilitation

The facilitation of small-group interaction has a rich history in social psychology. Foundational work by Tuckman (1965) on the stages of group development (forming, storming, norming, performing) highlights the natural, often challenging, progression groups undergo. Our model can be seen as a technological intervention designed to scaffold the "forming" and "norming" stages. Furthermore, the use of structured prompts is a well-established technique in manual group facilitation to accelerate bonding and ensure psychological safety (Kegan & Lahey, 2009). The CIM operationalizes this practice through a lightweight, adaptive, and scalable software layer. Unlike unstructured encounters, which can be dominated by a vocal few, our system aims to create a more equitable conversational floor, echoing Bales' (1950) early work on Interaction Process Analysis, which categorized contributions into task-oriented and socio-emotional functions.

Computational Approaches to Social Interaction

Within HCI, our work contrasts with mainstream approaches. Recommender systems, which are foundational to many social platforms, typically personalize content based

on user identity signals and past behavior (Ricci et al., 2011). The CIM, however, curates content based on inferred *interactional style*, a transient, context-dependent property of the group. While dialogue management systems have made significant strides in human-agent conversation (Jurafsky & Martin, 2009), their application to orchestrating multi-human, face-to-face dialogue is less explored. Our model functions not as a participant, but as a non-intrusive "digital facilitator," a concept with roots in the domain of Computer-Supported Cooperative Work (CSCW) and groupware (Grudin, 1994).

Edge Computing and Privacy-by-Design

The architectural decision to implement orchestration logic on the client-side aligns with the growing trend of edge computing (Shi et al., 2016). This approach offers three distinct advantages in our context: (1) **Reduced Latency:** On-device processing ensures near-instantaneous prompt selection, critical for maintaining conversational flow. (2) **Offline Functionality:** The system remains fully operational in environments with limited or unreliable connectivity. (3) **Privacy Preservation:** By minimizing data transmission and avoiding the server-side storage of sensitive interaction metrics (e.g., conversation duration, pass rates), the model adheres to the principle of privacy-by-design (Cavoukian, 2009), a crucial ethical consideration for technologies that mediate social behavior.

Theoretical Framework

The design of the Curated Interaction Model is explicitly grounded in three well-established psychological theories that collectively inform its core mechanics.

Optimal Dissonance. The theory of optimal dissonance, derived from Berlyne's work on aesthetic curiosity (1971), posits that engagement is maximized when an individual is exposed to stimuli that are moderately novel—neither too familiar (boring) nor too alien (threatening). Our blending algorithm directly operationalizes this principle. For a group with a dominant profile, the algorithm weights prompt selection heavily towards that profile's "home" deck (providing comfort and familiarity) while strategically injecting a smaller proportion of prompts from a contrasting secondary profile. This controlled introduction of novelty is designed to sustain curiosity and prevent conversational stagnation without inducing social anxiety.

Flow Theory. The concept of "flow," or a state of optimal experience, arises when an individual's perceived skills are well-matched to the perceived challenges of a task (Csikszentmihalyi, 1990). We model the 60–75 minute conversation as a flow-inducing activity. The structured progression of prompts from Level 1 (L1, low-risk icebreakers) to Level 2 (L2, personal experiences) and finally to Level 3 (L3, values and reflection) represents a calibrated scaling of challenge.

The immediate social feedback inherent in a face-to-face conversation (e.g., laughter, nods, reciprocal sharing) provides the continuous, unambiguous feedback loop necessary for achieving and maintaining this state. The fatigue-detection mechanism, which regresses the difficulty level in response to signals of disengagement, acts as a homeostatic control to keep the group within the "flow channel."

Social Identity Theory and Emergent Communion. According to Social Identity Theory (Tajfel & Turner, 1986), a significant part of an individual’s self-concept is derived from their perceived membership in social groups. The CIM is designed to foster the rapid, transient formation of a positive in-group identity at the dinner table. This is achieved through two mechanisms: first, the use of inclusive, non-evaluative prompts that minimize the risk of social judgment; and second, the creation of a shared ritual where all participants engage with the same prompt. This shared experience, distinct from the outside world, creates a temporary "us," a state of emergent communion that supports prosocial behaviors such as active listening, empathy, and mutual encouragement.

System Architecture and Method

Terminology: The *Deck* as a Curated Prompt Collection

Within the context of this framework, the term *deck* is defined with technical precision. A deck is a **versioned, cryptographically signed collection of curated prompts**, typically formatted as a JSON object. Each prompt (or "card") within a deck is tagged with metadata, including its associated conversational profile, difficulty level (L1, L2, L3), thematic content, and intended tempo (e.g., rapid, narrative). Decks are treated as immutable artifacts, delivered via a Content Delivery Network (CDN) and cached locally on the client device, ensuring consistency and offline functionality.

System Model and Operational Assumptions

The model is designed for a specific, controlled environment to minimize confounding variables. The primary assumptions are:

- A single table of 5–6 participants per session.
- A homogenous demographic cohort of young adults (ages 21–29).
- A premium, low-distraction venue conducive to conversation.
- Any participant’s personal mobile device can host the session, acting as the orchestrator.
- The system must be robust to intermittent or non-existent network connectivity post-setup.

- Sensitive PII (Personally Identifiable Information), such as photos, gender, or ethnicity, is never collected and therefore cannot influence content selection. The server’s persistence layer stores only a user ID and the derived, opaque profile token.

Questionnaire Design and Profile Classification

To classify participants without resorting to extensive personality inventories, we designed a minimalist, three-item forced-choice questionnaire. The items were developed to elicit preferences along three core dimensions of conversational style identified in preliminary research:

1. **Pace:** Preference for rapid, energetic exchange versus measured, thoughtful turns.
2. **Content:** Preference for abstract, analytical topics versus personal, narrative-driven stories.
3. **Stance:** Approach to disagreement, ranging from direct debate to harmony-seeking.

Each of the three possible answers (A/B/C) for a question contributes one point to two "allied" profiles (see Table 1). The final profile assigned is the one that achieves the maximum score (arg max) across the six possibilities. Ties are broken deterministically, first by prioritizing profiles associated with the answer to Q3 (stance on disagreement, considered the strongest signal), and second by selecting the profile currently least represented at the table, a heuristic designed to enhance group diversity.

Table 1

The Three-Item Classifier: Mapping of A/B/C Responses to Allied Conversational Profiles

| Question Dimension | A/B/C Response Mapping to Profiles |
|---------------------------|---|
| Q1: Conversation Pace | A → Vanguard, Connector; B → Storyteller, Aesthete; C → Classic, Analyst. |
| Q2: Content Preference | A → Vanguard, Analyst; B → Storyteller, Classic; C → Aesthete, Connector. |
| Q3: Handling Disagreement | A → Vanguard, Analyst; B → Storyteller, Classic; C → Connector, Aesthete. |

The Conversational Profile Taxonomy

The six profiles constitute a behavioral, not dispositional, taxonomy. They describe styles of interaction rather than fixed personality traits.

Vanguard — Characterized by energetic novelty-seeking, rapid idea exchange, and a comfort with micro-debates. *Behavioral Signals:* Low turn-taking latency, intellectual curiosity for contrarian viewpoints. *System Affordances:* Decks

with shorter L1/L2 prompts to maintain momentum, and a concise L3 prompt designed to crystallize core values.

Classic — Prefers calm, in-depth discussion with a measured pace and an appreciation for elegance in expression. *Behavioral Signals*: Longer, more deliberate turn-taking; high attention to linguistic phrasing and social etiquette. *System Affordances*: Highly reflective L2 prompts; a single, profound L3 prompt focused on practice and principles.

Storyteller — Engages through personal narratives, memory, and the building of empathy via lived experience. *Behavioral Signals*: Use of scene-setting, episodic recall, and emotional language. *System Affordances*: Prompts that anchor L1/L2 in concrete personal episodes; L3 prompts designed for meaning-making, explicitly avoiding therapeutic territory.

Analyst — Approaches topics with a desire for structure, principles, and causal reasoning. *Behavioral Signals*: Framing arguments in terms of trade-offs, curiosity about metrics and systems. *System Affordances*: L1 prompts based on factoids or "what-if" scenarios, L2 prompts on principles, and a single L3 focused on aligning values with logical frameworks.

Aesthete — Possesses a high sensitivity to sensory detail, ambiance, and design. *Behavioral Signals*: Frequent comments on light, sound, materials, or atmosphere. *System Affordances*: Vivid, imagistic L1/L2 prompts; L3 prompts that connect personal taste to identity without promoting social gatekeeping.

Connector — Acts as the social lubricant of the group, focusing on inclusive turn-taking and using gentle humor to build cohesion. *Behavioral Signals*: Actively inviting quieter participants to speak, summarizing conversational threads. *System Affordances*: "Bridge" prompts designed to appeal to multiple styles; periodic "reset" prompts to use after moments of silence or disconnection.

The Orchestration Algorithm

Semantic Description

The core of the CIM is a heuristic-based weighting algorithm that modulates prompt selection based on the group's real-time compositional dynamics. Upon session initialization, the client application calculates the frequency distribution of the six profiles present at the table. This distribution dictates the initial weighting scheme for sampling from the available prompt decks. The logic follows a set of prioritized rules:

- **Dominant Profile**: If a clear majority profile exists (e.g., three or more participants), its corresponding deck is heavily weighted to establish a comfortable baseline, with a secondary weight assigned to the next most common profile to introduce optimal dissonance.
- **Balanced Subgroups**: In the case of two balanced subgroups (e.g., a 2-2-1 composition), the algorithm primarily alternates between the two dominant profile

decks, periodically injecting "Connector" prompts to bridge the two styles and maintain group cohesion.

- **High Diversity**: For a highly heterogeneous group (e.g., five different profiles), the algorithm defaults to a "common ground" strategy, prioritizing "Connector" and "Aesthete" prompts, which are designed to be broadly accessible. It then samples proportionally from the other present profiles.
- **Fatigue Adaptation**: The system monitors for implicit negative feedback, defined as a high rate of skipped prompts (passes) or prolonged silences between turns. If a predefined fatigue threshold is crossed, the algorithm enacts a state transition, temporarily downshifting the prompt difficulty level (e.g., from L2 to L1) or changing the thematic category.

Formal Definition

Let $\mathcal{P} = \{p_1, \dots, p_6\}$ be the set of six profiles. For a given table instance T with N participants, let the vector $\mathbf{c} = (c_1, \dots, c_6)$ represent the counts of participants corresponding to each profile, where $\sum_{i=1}^6 c_i = N$. We define a weight vector $\mathbf{w} = (w_1, \dots, w_6)$ used for weighted random sampling of the next prompt's profile category. The weights are assigned based on the following heuristic function:

$$w_p = \begin{cases} 0.70, & \text{if } p = \arg \max(\mathbf{c}) \text{ and } \max(\mathbf{c}) \geq 3 \\ 0.30, & \text{if } p = \text{second_max}(\mathbf{c}) \text{ in the above case} \\ 0.50, & \text{if } \mathbf{c} \text{ shows a two-pairs pattern and } p \in \{p_A, p_B\} \\ 0.40, & \text{if } \mathbf{c} \text{ has high diversity and } p = \text{Connector} \\ 0.20, & \text{if } \mathbf{c} \text{ has high diversity and } p = \text{Aesthete} \\ \text{proportional,} & \text{otherwise (distributed among present profiles)} \end{cases} \quad (1)$$

The weight vector \mathbf{w} is normalized such that $\sum_{p \in \mathcal{P}} w_p = 1$. The prompt selection loop then performs weighted sampling from \mathcal{P} to choose a profile deck, followed by uniform random sampling of an appropriate-level card from within that deck, subject to constraints against immediate repetition of topics or speakers. Fatigue is detected if the pass rate over a moving window of 5 prompts exceeds 30%.

Selection Loop (Pseudocode)

```
function run_session(profiles):
    weights = calculate_initial_weights(
        profiles)
    history = new HistoryQueue(size=3)
    level = L1

    while session_is_active():
        // Sample profile, avoiding immediate
        // repetition
        deck_profile = weighted_sample(weights,
```

```

        avoid=history.
            last_profile()
        )

// Sample card, avoiding recent topics
// and speakers
card = sample_card_from_deck(
    deck_profile, level,
        avoid_topics=
            history.topics
        ),
    avoid_users=
        history.users
    ())

present(card)

// Record interaction signals
signal = wait_for_user_signal() // e.g.,
    pass, laugh
history.add(deck_profile, card.topic,
    signal.user)

// Adapt based on signals
if fatigue_is_detected(history):
    level = regress_level(level)
    weights = adjust_weights_for_fatigue(
        weights)
else:
    level = advance_level(level,
        session_progress)
    
```


System Visualizations

Client-Server Communication Flow

The system’s architecture (Figure 1) is intentionally minimalist to support the core principles of privacy and on-device processing. The client application interacts with two server-side components: a CDN for efficient delivery of static, versioned prompt decks, and a minimal API for the sole purpose of user authentication and storage of the derived profile token. All dynamic, session-specific logic occurs client-side.

Figure 1

Client-server data flow, emphasizing the edge-first architecture. All real-time orchestration logic resides on the client device.




Session State Transition Diagram

A typical session progresses through a finite set of states, primarily defined by the difficulty level of the prompts (Figure 2). The session begins in a Setup state, moves sequentially from L1 to L3, and concludes with a Closing phase. Critically, the detection of a fatigue signal can trigger a regression to a lower-level state, ensuring the conversational challenge remains within the group’s comfort zone.

Figure 2

A simplified finite state diagram of a session’s progression. Forward transitions are driven by time and successful interaction, while backward transitions (regressions) are triggered by fatigue detection.



Security, Privacy, and Safety by Design

The design of the CIM incorporates a robust framework for ethical considerations.

Principles: Our approach is guided by data minimization, local-first processing, and user transparency.

Threats and Mitigations: The primary attack vectors are profile inference and content catalog tampering. We mitigate these through the use of opaque profile tokens (preventing reverse-engineering of user answers), mandatory TLS pinning for all API communication, and cryptographic signing of content catalogs to ensure their integrity.

Programmatic Safety Guardrails: To ensure psychological safety, the prompt content is governed by a strict safety policy. Politically charged topics, personal health, religion, and intimate romantic history are explicitly avoided. Furthermore, two programmatic guardrails are enforced by the orchestrator: a "comfort floor" prevents the introduction of high-level (L3) prompts too early in the session, and a "variance cap" limits the number of consecutive deep or challenging prompts to prevent emotional exhaustion. User agency is paramount; an explicit opt-out (pass) is available for any prompt without penalty.

Evaluation Plan

We propose a multi-stage empirical evaluation to validate the efficacy of the Curated Interaction Model.

Hypotheses

We will test the following primary hypotheses:

H1: Groups using the CIM will exhibit a more equitable distribution of speaking time compared to control groups using unstructured prompts.

H2: The latency to the first instance of shared laughter (a proxy for initial rapport) will be significantly shorter in CIM-facilitated groups.

H3: CIM groups will achieve a greater depth of conversation, measured by the percentage of Level 3 prompts successfully engaged with, without a corresponding increase in self-reported social anxiety.

H4: Participants in the CIM condition will report significantly higher levels of perceived connection (interpersonal closeness) and return intent than participants in the control condition.

Experimental Design and Instrumentation

We will employ a between-groups experimental design. The experimental group will use the full CIM, while the control group will use an identical application interface that delivers prompts of equivalent length and reading level, but drawn randomly from a single, thematically neutral pool. Assignment to conditions will be randomized by table and date.

Instrumentation will include:

- **Behavioral Metrics (logged by the app):** Per-prompt timestamps, speaker IDs (ephemeral), pass flags, and user-initiated laughter markers.
- **Perceptual Metrics (post-session survey):** Self-reported measures of connection (e.g., the Inclusion of Other in the Self Scale; Aron et al., 1992), comfort, memorability of the conversation, and return intent, captured via Likert scales.
- **Qualitative Data:** Open-ended responses to identify which conversational prompts or sequences had the most significant perceived impact.

All procedures will be subject to Institutional Review Board (IRB) approval.

Power Considerations

To detect a medium effect size (Cohen's $d \approx 0.5$), which is typical for interventions in HCI and social psychology, with a standard alpha of $\alpha = .05$ and statistical power of .80, a two-sample t-test requires approximately 64 participants per condition. Given the clustered nature of our data (participants nested within tables), we will either analyze at the table level or employ hierarchical linear models (HLM) to account for non-independence, adjusting our sample size accordingly. A sequential analysis plan will be considered to allow for staged rollout with robust error control.

Analysis Plan

For our primary hypotheses, we will use non-parametric Mann-Whitney U tests for between-condition comparisons on our key metrics, as we do not assume normality of the underlying distributions. Effect sizes will be reported using Cliff's delta. A Holm-Bonferroni correction will be applied across multiple comparisons to control the family-wise error rate. Qualitative data from open-ended responses will be analyzed using thematic analysis to identify emergent themes and archetypal prompts that demonstrate the strongest positive or negative valence.

Discussion, Limitations, and Future Work

This paper has presented the Curated Interaction Model, a comprehensive framework for technologically facilitating high-quality, face-to-face conversation. By prioritizing interactional style over user identity and adopting a privacy-preserving, on-device architecture, the CIM offers a new paradigm for social software.

Limitations

Our initial proposal has several limitations that provide avenues for future research. The six-profile taxonomy is empirically derived but requires large-scale validation. The algorithm's weighting heuristics are theory-driven but not yet empirically optimized. The initial study population (21-29 year olds in a specific cultural context) limits the generalizability of our findings. The model's efficacy may also be dependent on the quality and diversity of the curated prompt decks.

Future Work

We envision several extensions to this work.

1. **Algorithmic Refinement:** Future iterations could replace the heuristic-based weighting system with a machine learning model, potentially using reinforcement learning to dynamically optimize prompt selection based on real-time group feedback signals.
2. **Longitudinal Studies:** A longitudinal study could track groups over multiple sessions to investigate how technologically-scaffolded rapport develops and persists over time.
3. **Generalization and Adaptation:** A key research pathway involves adapting and validating the model for different demographic cohorts (e.g., older adults, adolescents) and contexts (e.g., professional team-building, educational settings). This would likely require recalibrating both the classifier and the prompt content.

- 4. **Privacy-Preserving Learning:** We plan to explore federated learning or other privacy-preserving aggregation techniques to allow the model to learn and improve from usage across a large population without centralizing raw, sensitive interaction data.

References

- [1] Aron, A., Aron, E. N., & Smollan, D. (1992). Inclusion of Other in the Self Scale and the structure of interpersonal closeness. *Journal of Personality and Social Psychology*, 63(4), 596–612.
- [2] Bales, R. F. (1950). *Interaction process analysis: A method for the study of small groups*. Addison-Wesley.
- [3] Berlyne, D. E. (1971). *Aesthetics and psychobiology*. Appleton-Century-Crofts.
- [4] Cavoukian, A. (2009). *Privacy by Design: The 7 Foundational Principles*. Information and Privacy Commissioner of Ontario.
- [5] Csikszentmihalyi, M. (1990). *Flow: The psychology of optimal experience*. Harper & Row.
- [6] Grudin, J. (1994). Groupware and social dynamics: eight challenges for developers. *Communications of the ACM*, 37(1), 92–105.
- [7] Jurafsky, D., & Martin, J. H. (2009). *Speech and language processing: An introduction to natural language processing, computational linguistics, and speech recognition* (2nd ed.). Prentice Hall.
- [8] Kegan, R., & Lahey, L. L. (2009). *Immunity to change: How to overcome it and unlock the potential in yourself and your organization*. Harvard Business Press.
- [9] Ricci, F., Rokach, L., & Shapira, B. (2011). Introduction to recommender systems handbook. In *Recommender systems handbook* (pp. 1–35). Springer.
- [10] Shi, W., Cao, J., Zhang, Q., Li, Y., & Xu, L. (2016). Edge computing: Vision and challenges. *IEEE Internet of Things Journal*, 3(5), 637–646.
- [11] Tajfel, H., & Turner, J. C. (1986). The social identity theory of intergroup behavior. In S. Worchel & W. G. Austin (Eds.), *Psychology of intergroup relations* (pp. 7–24). Nelson-Hall.
- [12] Tuckman, B. W. (1965). Developmental sequence in small groups. *Psychological Bulletin*, 63(6), 384–399.

Appendix A

*

Appendix A: Deck Catalog (Schema and Examples)

```
{
  "deck_version": "2025.10",
  "profiles": ["vanguard", "classic", "storyteller", "analyst", "aesthete", "connector"],
  "cards": [
    {
      "id": "v1", "profile": "vanguard", "level": 1, "tempo": "rapid", "theme": "city",
      "text": "If you could rewrite one unwritten rule in this city, what would it be?"
    },
    {
      "id": "c2", "profile": "classic", "level": 2, "tempo": "story", "theme": "habits",
      "text": "Describe a timeless habit that keeps your week in order."
    },
    {
      "id": "s3", "profile": "storyteller", "level": 2, "tempo": "story", "theme": "turning_points",
      "text": "Tell a story about a small decision that led to a significant outcome."
    }
  ]
}
```

Appendix B

*

Appendix B: Full Mapping Matrix

Table B1

Complete Mapping Matrix for the Three-Item Classifier

| Item | Option | Profiles (Points Awarded) |
|------|--------|---------------------------|
| Q1 | A | Vanguard, Connector |
| Q1 | B | Storyteller, Aesthete |
| Q1 | C | Classic, Analyst |
| Q2 | A | Vanguard, Analyst |
| Q2 | B | Storyteller, Classic |
| Q2 | C | Aesthete, Connector |
| Q3 | A | Vanguard, Analyst |
| Q3 | B | Storyteller, Classic |
| Q3 | C | Connector, Aesthete |

Appendix C

*

Appendix C: Minimal On-Device Flow (TikZ)

Figure C1

On-device orchestration and data boundaries, illustrating the flow from user classification to real-time prompt selection.

