

InnerPond: Fostering Inter-Self Dialogue with a Multi-Agent Approach for Introspection

Hayeon Jeon
hci+d lab., IPAI
Seoul National University
Seoul, Republic of Korea
jhy94520@snu.ac.kr

Dakyeom Ahn*
hci+d lab., Department of
Communication
Seoul National University
Seoul, Republic of Korea
adklys@snu.ac.kr

Sunyu Pang*
hci+d lab., IPAI
Seoul National University
Seoul, Republic of Korea
sunyu.pang@snu.ac.kr

Yunseo Choi*
hci+d lab., IPAI
Seoul National University
Seoul, Republic of Korea
dbstj0531@snu.ac.kr

Suhwoo Yoon
hci+d lab., Department of
Communication
Seoul National University
Seoul, Republic of Korea
yeopil@snu.ac.kr

Joonhwan Lee
hci+d lab., Department of
Communication
Seoul National University
Seocho-gu, Seoul, Republic of Korea
joonhwan@snu.ac.kr

Eun-mee Kim[†]
Department of Communication
Seoul National University
Seoul, Republic of Korea
eunmee@snu.ac.kr

Hajin Lim[†]
hci+d lab., Department of
Communication
Seoul National University
Seoul, Republic of Korea
hajin@snu.ac.kr

Abstract

Introspection is central to identity construction and future planning, yet most digital tools approach the self as a unified entity. In contrast, Dialogical Self Theory (DST) views the self as composed of multiple internal perspectives, such as values, concerns, and aspirations, that can come into tension or dialogue with one another. Building on this view, we designed InnerPond, a research probe in the form of a multi-agent system that represents these internal perspectives as distinct LLM-based agents for introspection. Its design was shaped through iterative explorations of spatial metaphors, interaction scaffolding, and conversational orchestration, culminating in a shared spatial environment for organizing and relating multiple inner perspectives. In a user study with 17 young adults navigating career choices, participants engaged with the probe by co-creating inner voices with AI, composing relational inner landscapes, and orchestrating dialogue as observers and mediators, offering insight into how such systems could support introspection. Overall, this work offers design implications for AI-supported introspection tools that enable exploration of the self's multiplicity.

*These authors contributed equally to this research.

[†]Co-corresponding authors



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License.

CHI '26, Barcelona, Spain

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2278-3/26/04

<https://doi.org/10.1145/3772318.3791248>

CCS Concepts

• Human-centered computing → Natural language interfaces.

Keywords

Introspection, inner dialogue, self-reflection, LLM, multi-agent, Dialogical Self Theory, DST

ACM Reference Format:

Hayeon Jeon, Dakyeom Ahn, Sunyu Pang, Yunseo Choi, Suhwoo Yoon, Joonhwan Lee, Eun-mee Kim, and Hajin Lim. 2026. InnerPond: Fostering Inter-Self Dialogue with a Multi-Agent Approach for Introspection. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26)*, April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 25 pages. <https://doi.org/10.1145/3772318.3791248>

1 Introduction

Individuals constantly face complex life choices [101], where introspection—the critical examination of one's own thoughts, values, and emotions [115]—plays a crucial role in navigating these decisions [13, 19]. A key mechanism that facilitates introspection is inner dialogue, in which different internal perspectives engage in a mental conversation [31, 85, 92]. To better understand the complexities of this inner dialogue, the **Dialogical Self Theory (DST)** [41, 45] provides a theoretical framework. Rather than viewing the self as an isolated, unified entity, DST conceives it as a relational process—a “dynamic multiplicity of I-positions” [42]. Put simply, our mind hosts various ‘selves’—like ‘the creative self’ or ‘the fearful self’—each possessing a distinct perspective and narrative [44]. These I-positions are not isolated; they engage in a constant dialogue, influencing and repositioning one another




Figure 1: InnerPond fosters AI-mediated introspection through multi-agent dialogue. The system transforms users’ I-positions into conversing LLM agents (lotus leaves), supporting dialogue among multiple I-positions within the self.

within a dialogical space [40, 43]. Through such interaction, a more reflective and emergent self-understanding could emerge [44, 45].

While inner dialogue is a vital mechanism for introspection, in its spontaneous form, it may remain constrained within familiar perspectives or spiral into unproductive rumination [14, 34]. To scaffold this process more constructively, various aids for introspection have been explored—from artifacts such as journals, photos, and music [21, 91, 110] to more recent LLM-based systems [54, 108]. These approaches have provided valuable ways to foster introspection. However, by framing users primarily as singular, unified selves, they have given less attention to the plural and dynamic nature of the self—particularly the way multiple internal voices coexist, compete, and negotiate with one another during complex life choices [44, 45].

We aim to address this gap by translating DST’s concept of a dialogical self into an interactive system and, through this instantiation, generating design knowledge about supporting dialogue among one’s multiple selves. In particular, LLMs can now embody distinct personas with consistent perspectives [23, 103], enabling new ways to externalize this plurality. Leveraging this capability, we investigate a design exploration that makes inner multiplicity explicit—representing multiple I-positions as distinct AI agents that can dialogue with one another and with the user. Through this approach, we explore how people experience and make sense of such multi-agent inner dialogue. Building on this, we address the following research questions (RQs):

- **RQ1:** What are the key design considerations for supporting dialogue among multiple selves in exploring the plural and dynamic self?
- **RQ2:** How do people experience and make sense of engaging with their multiple selves through multi-agent AI-mediated introspection?

In addressing the RQs, we designed and developed **InnerPond** as a research probe [12, 20, 49]—not to evaluate system effectiveness, but to inspire design exploration of how people experience dialogical engagement with their multiple I-positions. Informed by design considerations around spatial metaphors, scaffolding structures, and multi-agent conversational dynamics, InnerPond represents an individual’s various I-positions as independent LLM-based agents that can be visualized, elaborated, and brought into dialogue within a shared space. Anchored in the lotus leaf and pond metaphor (Figure 1), InnerPond aims to enact coexistence and interconnection among multiple selves, helping users explore their inner landscape and move toward more emergent forms of self-understanding.

To understand how people experience such engagement with their multiple selves, we conducted a user study (N=17) with young adults deliberating between two career paths—a context where competing internal voices often surface [3, 113] and structured support can help navigate their complexity [11, 64]. Our findings showed that through externalizing and dialoguing with their inner voices, participants could develop a meta-cognitive perspective—recognizing distinct selves as interconnected parts of a larger self. For some, this shift reframed how they approached career decisions, moving from reacting to external pressures toward reflecting on internal alignment.

Through this exploration, we use the term “*inter-self communication*” to describe dialogue among one’s own inner voices, as conceptualized in DST. Unlike solitary reflection, it gives distinct voices space to interact, making the multiplicity of the self more tangible; unlike dialogue with others, it keeps exchanges grounded within the self. This perspective opens broader possibilities for designing AI systems that support people in exploring, elaborating, and engaging with their plural selves. We discuss the implications of this approach, considering both the opportunities it offers and the design tensions that accompany it.




Figure 2: Core concepts of Dialogical Self Theory, illustrating the relationship between multiple I-positions in the dialogical space and a meta-position.

Taken together, this paper makes two contributions to HCI:

- **Design considerations for plural and dynamic self engagement** — Informed by our design journey with InnerPond, we articulate key considerations for translating DST into interactive experience: spatial metaphors that convey coexistence and interconnection, scaffolded stages that structure the introspective journey, and multi-agent dynamics that balance consistency with flexibility. These considerations extend current approaches to AI-supported introspection by shifting from a unified to a dialogical view of the self.
- **Empirical insights into inter-self dialogue** — We provide rich qualitative insights into how people experience and make sense of engaging with their multiple selves, revealing how inter-self dialogue is experienced and interpreted when instantiated through a multi-agent system. Our findings surface both opportunities and design tensions, offering implications for future systems that aim to support the plural and dynamic nature of the self.

2 Related Work

This section examines inner dialogue as a core mechanism of introspection and introduces Dialogical Self Theory. We then review existing approaches to supporting introspection. Finally, we highlight the potential of multi-agent systems as a means to support dialogical forms of introspection.

2.1 Dialogical Self Theory as Framework for Understanding Inner Dialogue

Introspection is a practice of looking inward and exploring one’s own thoughts, values, and emotions [19]. In psychology, it is characterized as explicitly observing and reflecting on an individual’s

mental state [13]. Distinguished from simply reflecting on or recollecting memories, introspection involves critically evaluating significant experiences and patterns in one’s life, and continuously questioning what has been achieved and what is desired [115]. Introspection and self-reflection are closely related—both involve examining one’s mental states to foster self-understanding [4, 32]. However, they differ in scope: self-reflection refers to the cognitive process of examining one’s thoughts and actions, whereas introspection encompasses a broader philosophical construct that includes emotional and evaluative dimensions [48]. Through this process, individuals search for answers concerning their identity, values, and life priorities, fostering deeper self-understanding and identity construction [15, 16]. Particularly in life choice contexts such as career exploration, active introspection on one’s identity and life directions plays a crucial role [96].

A key internal process facilitating introspection is *inner dialogue* [31]. Inner dialogue is a process where “different internal perspectives engage in mental conversation,” exchanging views and positions within the mind [85, 92]. Unlike simple monologues, inner dialogue reflects the dialogical nature of the self and simulates social dialogical relations internally [2, 46]. Through confronting and integrating diverse internal perspectives, it enables reflexive and contemplative self-understanding and identity construction [66, 93].

To better conceptualize this process, *Dialogical Self Theory (DST)* [41, 45] offers a structured framework for understanding the complex dynamics of inner dialogue. It conceptualizes inner dialogue as an exchange of thoughts and perspectives between distinct I-positions within the self, helping identify and interpret various forms of internal dialogical activity during introspective processes [85]. DST reconceptualizes the self not as an isolated inner entity but as an inherently relational process—a dialogical being described as a dynamic multiplicity of I-positions within the society of mind [42]. Each I-position represents a narratively structured unit of the

self with its own perspective and voice. For example, I-positions such as “I as dreamer” or “I as fearful” reflect different internal aspects of the person [44], engaging in dialogical relationships [43].

These I-positions are shaped as external voices become internalized within the self. Through multidimensional inner dialogue unfolding in a dialogical space, these I-positions are dynamically repositioned through DST’s core mechanisms—positioning, counter-positioning, and repositioning—with their relationships continually restructured [40, 45]. DST further introduces the concept of the ‘*meta-position*’, a higher-level vantage point from which a person can observe and reflect on their various I-positions. This capacity for meta-level reflection serves as a crucial mechanism for self-understanding and integration [43–45].

This view resonates with recent theorizing in personal informatics, where the self has been reconceptualized as dynamic and constructed through ongoing interaction with the world [97]. Yet, while such work emphasizes the self’s temporal and social dimensions—how it evolves across past, present, and future, and in various social contexts—it gives less attention to the dialogical mechanisms through which distinct aspects of the self actively engage with one another. DST complements this perspective by foregrounding these inter-positional dynamics, offering a foundation for understanding introspection not merely as self-observation, but as a process through which self-understanding emerges from dialogue among multiple internal voices.

2.2 Existing Approaches to Supporting Introspection

While inner dialogue serves as a key mechanism for introspection, its naturally occurring forms are often limited by the individual’s existing experiences and perspectives, resulting in constrained self-understanding [14]. In situations where thoughts and emotions are deeply intertwined, individuals can fall back on habitual thinking patterns or become trapped in negative cycles such as rumination [34]. To address these challenges, various introspective interventions have been developed.

Traditional introspective interventions have supported self-reflection through personal artifacts such as photographs, letters, music, and diary writing [8, 9, 63, 105]. Among these, diary writing has been one of the most common methods for critically evaluating past experiences and questioning the future [86]. However, these writing-based approaches, grounded in a person’s existing perspective [79], may have limitations in offering new insights or alternative viewpoints [80].

Further, digital technologies have expanded these practices, incorporating journals [88, 110], photographs [21, 91], music [60, 84, 88], and social media data [26]. A parallel line of work has explored quantified self-tracking, where people reflect by examining historical records of sensor-based data [28, 99]. However, research in reflective informatics emphasizes that reflection does not arise automatically from exposure to personal artifacts or data; carefully scaffolded experiences are needed [6, 107]. To provide such scaffolding, early conversational agents explored structured reflection through rule-based interactions [1, 62, 74]. These systems have provided valuable entry points for reflection, though they often offered limited adaptability to individual contexts [68].

More recently, Large Language Models (LLMs) have enabled personalized and adaptive support for reflection [5, 56]. ExploreSelf [108] supports users in articulating personal challenges through adaptive questioning. In mental health contexts, MindfulDiary [61] helps psychiatric patients document daily experiences through conversational journaling. For career exploration, Letters from Future Self [54] facilitates introspection through dialogues with LLM-based agents that simulate a future self. These systems represent meaningful steps forward, yet they typically frame users as singular, unified selves. Given that the self can be viewed as dialogical and multi-positional, how to support exchanges among distinct internal perspectives remains underexplored. This motivates our approach to facilitate inner dialogue across multiple I-positions—a design goal we explore in this paper.

2.3 Potential of Multi-Agent Systems for Supporting Inner Dialogues

Recent studies have shown that LLMs are remarkably capable of simulating distinct personas [23, 103]. Beyond surface-level role-playing, LLMs can now generate agents that reflect the complex interplay of identity components, such as personality traits [55, 71] and value systems [121, 123], mirroring the richness of real-world individuals [67]. These agents could maintain consistent narratives and unique perspectives while responding adaptively to different contexts [54].

Building on this, multi-agent systems are emerging as powerful tools for surfacing diverse perspectives and domain-specific insights across fields [36]. In particular, they have shown promise in collaborative problem-solving, decision support, and creative ideation [112]. For instance, in software engineering, agents simulating roles such as product managers, architects, and engineers can coordinate to tackle complex development tasks [47]. In creativity support, agents with distinct backgrounds and viewpoints can engage in brainstorming sessions, generating more original ideas than a single LLM operating alone [73].

Here, we draw attention to the potential of combining LLMs’ persona-simulation capabilities with a multi-agent architecture to support inner dialogue. By simulating distinct personas as different I-positions, such systems may capture complex dialogical processes—including disagreement, perspective shifts, and finding common ground—forms of interaction that prior studies have documented among agents within multi-agent structures [100, 122].

Despite this potential, the use of LLM-based multi-agent systems for introspection through inner dialogue remains largely unexplored [25, 29, 30]. In response, we developed **InnerPond**, a research probe that applies a multi-agent LLM architecture to externalize inner voices as dialogical agents, enabling users to engage them in structured inner dialogue.

3 InnerPond: Translating DST into Interactive Experience

Because Dialogical Self Theory (DST) conceptualizes the self as a plurality of interacting inner voices rather than a single unified identity, translating this theoretical perspective into an interactive system demanded careful attention to the experiential conditions




Figure 3: Three-phase metaphor evolution for the dialogical space: From group chat as battle ring (Phase 1), to stones highlighting individuality yet emphasizing separation (Phase 2), to lotus leaves appearing distinct yet sharing roots, enabling a meta-position perspective (Phase 3).

under which inner dialogue can unfold. Our design process followed a research-through-design approach [125, 126], using iterative prototypes as probes to explore how DST could be materially instantiated and how users might meaningfully engage with their plural selves.

At the outset, we established three core design goals grounded in DST’s foundational principles. These goals functioned not as technical requirements but as orienting principles that guided design experimentation and decision-making throughout the development of InnerPond.

Design Goal 1 (DG1) — Support coexistence rather than resolution. Inner conflict is commonly framed as a debate in which one position must ultimately prevail [82]. In contrast, DST emphasizes coexistence, negotiation, and the ability to hold tensions without prematurely collapsing them into a final answer [44, 45]. Our goal was to create interaction structures that sustain multiple voices in parallel and resist convergence toward a single dominant perspective.

Design Goal 2 (DG2) — Ensure equal attention and legitimacy for each inner voice. Plurality is not simply the presence of many voices; it requires that each voice be heard, recognized, and granted the opportunity to contribute meaningfully [27, 94]. We sought to design a system that encourages engagement with both dominant and quieter perspectives, supporting self-authorship and preventing strong voices from overshadowing vulnerable or easily overlooked ones.

Design Goal 3 (DG3) — Foster relational connectedness and meta-positional perspective. Externalizing multiple voices risks fragmenting the self into disconnected pieces. DST emphasizes that inner perspectives are dynamically related, and that reflection emerges through perceiving these relationships from a higher-level meta-position [35, 124]. Our goal was to support users in experiencing a holistic, relational structure of identity—shifting from being inside a conflict to observing and orchestrating it from a meta-positional perspective.

These three goals served as the conceptual backbone of our design inquiry. The subsequent sections detail how iterative explorations of spatial metaphors, interaction scaffolding, and conversational orchestration functioned as design experiments to operationalize and balance these goals, culminating in the current form of InnerPond.

3.1 Metaphor Evolution for the Dialogical Space

Guided by these design goals, we explored spatial metaphors for a dialogical space where multiple inner voices could coexist as distinct yet interconnected agents. This exploration evolved through three phases—each iteration revealing tensions that prompted further refinement.

Phase 1: Group Chat Room as Battle Ring. In the early design phase, we envisioned a ‘debate arena’ based on the observation that inner conflict often manifests as competing perspectives [82]. Borrowing the familiar group chat interface [57, 78], we envisioned creating an environment where each voice could clash in real-time. While this approach was promising for externalizing conflict, we realized that it implied the assumption that “one voice must win”—directly contradicting DG1’s emphasis on coexistence without premature resolution. This prompted us to seek a space where multiple perspectives could evolve together, rather than compete for dominance.

Phase 2: Landscape of Mind through Stone Metaphor. To design a space for coexistence rather than victory, we drew inspiration from DST-based therapies [51, 116] that use natural objects like stones [65] to symbolize I-positions and reconstruct identity through spatial arrangement. Translating this into a digital environment, we adopted a “Landscape of Mind” interface where users could arrange digital stones to construct their own inner maps. This approach addressed DG1 by allowing voices to coexist spatially, and partially supported DG2 by giving each I-position a distinct, tangible representation. However, while the stone metaphor effectively highlighted the individuality of each I-position, we recognized that it simultaneously emphasized separation and fixed identity—failing

to address **DG3**'s emphasis on relational connectedness and meta-position perspective [40, 45].

Phase 3: Inner Pond through Lotus Leaf Metaphor. To address the limitations of stones as representations of selves, we explored alternative metaphors, including a 'Zen Garden' [42, 83]. However, its emphasis on stillness and fixed balance was misaligned with the dynamic, fluid nature of I-positions. We ultimately arrived at the "Lotus Leaf" metaphor, where each leaf appears independent on the surface while remaining connected through shared roots beneath the water—capturing both the distinctness and interconnectedness of I-positions. This metaphor integrates all three design goals: leaves coexist on the pond surface without one dominating another (**DG1**); each leaf maintains a distinct and visible identity (**DG2**); and the shared underwater roots, together with the user's bird's-eye perspective, naturally embody DST's concept of the 'meta-position,' enabling relational connectedness and holistic observation (**DG3**).

3.2 Scaffolding for Inter-Self Introspection

With the spatial metaphor in place, we shifted our focus to structuring how users would move through the introspective experience. As a reference framework, we drew on "Composing the Self" [65], a DST-based therapeutic technique using stones described earlier. This therapy guides clients through a seven-step process—including 1) selecting stones as I-positions, 2) arranging them spatially to visualize relationships, 3) labeling (naming) each stone, 4) examining the arrangement, 5) voicing each position, 6) repositioning stones, and 7) reflecting on the composition as a whole.

We iteratively adapted these seven steps into four key stages, designed to follow a natural cognitive flow while lowering the barrier to engaging with one's multiple selves. Rather than preserving each therapeutic step as a discrete phase, we reorganized them into higher-level stages that reflect how these activities unfold in an interactive, AI-mediated context.

Specifically, Steps 1 (selecting stones) and 3 (labeling stones) were combined into *Stage 1: I-position Construction*, while Steps 2 (arranging), 4 (examining), and 6 (repositioning stones) were consolidated into *Stage 2: Spatial Configuration*. Step 5 (voicing each position) was adapted as *Stage 3: Dialogical Exchange*, and Step 7 (reflecting on the composition) became *Stage 4: Reflective Snapshot*. The following subsections elaborate on how each stage reinterprets its corresponding therapeutic steps by leveraging the affordances of an LLM-enabled, interactive medium.

From Manual Selection to AI-Augmented Identification. The original therapy we drew on [65] requires clients to structure and verbalize their inner world from scratch, demanding a high level of self-awareness and cognitive load [51]. We envisaged transforming this into AI-augmented identification, leveraging the analytical capabilities of LLMs. Specifically, we utilized LLMs to assist users in identifying their multiple selves, drawing on user-provided data such as personality traits, values, and personal narratives. To avoid imposing a fixed interpretation, we designed these AI-generated I-positions as starting points that users can enrich through co-creation—exploring, editing, or conversing with the suggested inner voices in their own words. This approach served as the foundation for [**Stage 1: I-position Construction**] in InnerPond, helping

users concretize ambiguous inner voices into tangible objects (individual lotus leaves) and recognize their multifaceted selves.

From Physical Arrangement to Digital Visualization. In the original therapy [65], clients are asked to express relationships between different selves through physical distance and height of stones. We translated this into digital manipulations, such as adjusting the position, size, and color of the lotus leaves. Unlike physical stones, these attributes can be easily adjusted, allowing users to fluidly explore and reshape their inner relationships. This approach was implemented as [**Stage 2: Relational Positioning**], where users spatially compose their inner landscape by arranging lotus leaves on the pond.

From Role-playing to Dialogical Exchange. In the original therapy [65], clients sequentially role-play each stone's perspective to negotiate inner conflicts. We reinterpreted this role-playing step as real-time dialogue among agents representing different selves as lotus leaves. This aimed to allow users to observe conversations among their selves from a third-person perspective or intervene as a mediator, without the burden of direct role-playing. This approach became the core of [**Stage 3: Dialogical Exchange**], enabling dynamic exchanges among conflicting selves through dialogue.

From Closing to Temporal Integration. In the original therapy [65], the session concludes with a one-time reflection on the final composition. We reframed this closing not as an ending but as a record of a "Temporal Self-portrait"—a record that can accumulate over time, helping users understand their inner world as a flowing narrative rather than a fixed entity. This concept was materialized in [**Stage 4: Reflective Snapshot**], where users capture the current configuration of their pond to preserve the moment of introspection.

3.3 Conversational Orchestration and System Implementation

The four-stage scaffolding established above was implemented in **InnerPond**, an LLM-based multi-agent system that guides users through four stages: (1) *I-position Construction*, (2) *Relational Positioning*, (3) *Dialogical Exchange*, and (4) *Reflective Snapshot*. While presented sequentially, these stages can be navigated freely, allowing for non-linear self-exploration.

We situated this implementation within the context of *career decision-making*—specifically, situations where users face a choice between two career paths. Such decisions naturally evoke inner dialogue as multiple perspectives and value conflicts come into play [96]. Career decisions provide concrete scenarios in which diverse I-positions are activated, interact, and evolve. While our study focuses on this context, the system's architecture can extend to other situations involving inner multiplicity, such as work-life balance or relationship conflicts [77].

In the following sections, we detail how each stage translates into concrete interface features and underlying technical mechanisms, and describe the overall technical configurations.

3.3.1 Stage 1: I-position Construction. This stage centers on generating and enriching I-positions through a pipeline that combines LLM analysis with user input (Figure 4). To extract diverse selves relevant to users choosing between two career paths, we carefully




Figure 4: I-position extraction and story enrichment pipeline: user knowledge is transformed into distinct I-positions, iteratively refined through scaffolding questions and user responses.

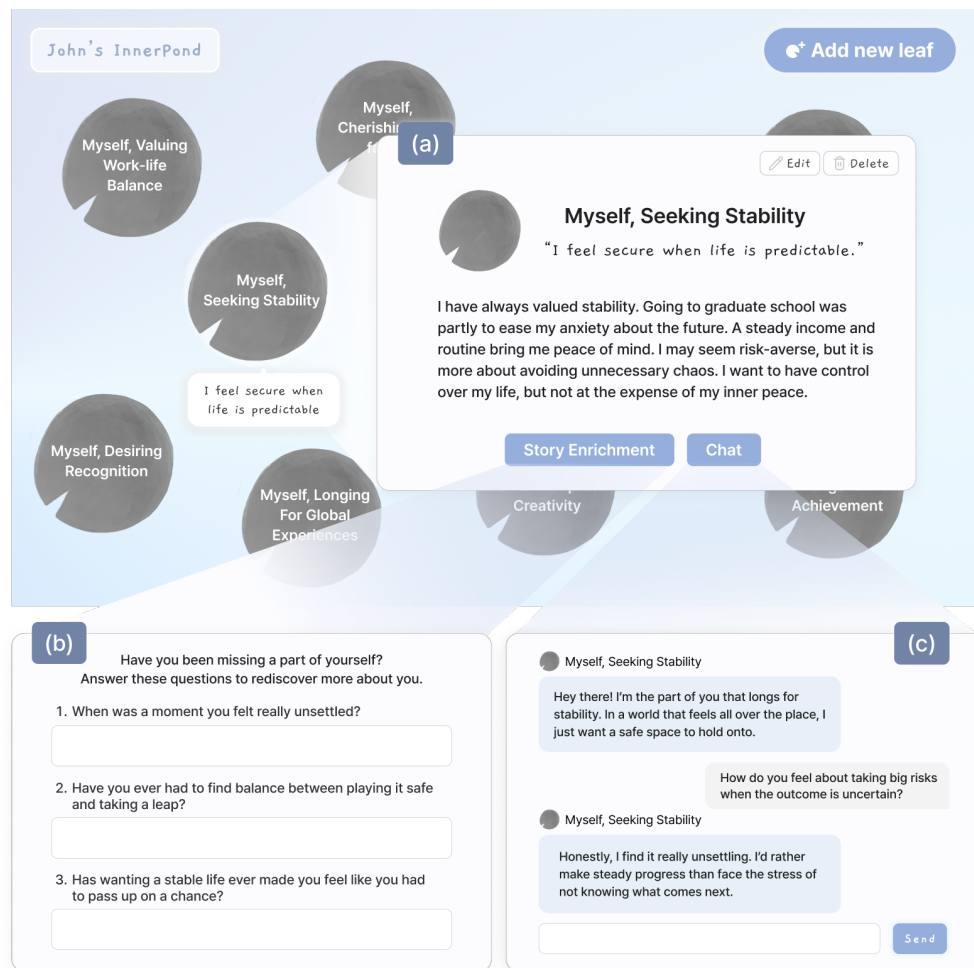


Figure 5: The interface for [Stage 1: I-position Construction]: (a) I-positions visualized as lotus leaves on the main pond view, with a profile modal showing name, viewpoint, and narrative; (b) “Story Enrichment” modal with scaffolding questions to refine the narrative; (c) “1:1 Dialogue” modal for direct conversation with a leaf agent.



Figure 6: The interface for [Stage 2: Relational Positioning]: Users can freely customize their inner landscape by (a) adjusting leaf size, (b) repositioning leaves via drag-and-drop, and (c) selecting colors through the dewdrop-shaped button.

designed a knowledge base structure using the SPeCtrum framework [67]. This knowledge includes demographics, personality traits (Big Five Inventory–2 Short Form (BFI-2-S) [109]), work values (Super’s Work Values Inventory (SWVI) [98]), self-identified strengths and weaknesses, and the two career options under consideration. Quantitative scale data are converted into natural language summaries using Chain of Density (CoD) prompting [102] for effective LLM understanding (see Appendix A.1 for the complete knowledge structure).

Based on this knowledge structure, we prompted the LLM to extract approximately ten distinct inner voices (I-positions), ensuring no overlap among them. We ensured they were distributed evenly across three categories: (1) those common to both career paths, (2) those specific to one career path, and (3) those specific to the other. To represent each I-position as an authentic inner voice, we constructed a profile for each, comprising a name in the “Myself, ...” format (e.g., [Myself, Seeking Stability]), a core viewpoint capturing its unique voice, and a first-person narrative revealing underlying motivations (see Appendix A.2 for the full prompt).

Based on these profiles, each I-position is visualized as an interactive leaf agent—a named lotus leaf on the pond that embodies a distinct inner voice (Figure 5-(a)). Users can hover over a leaf to reveal its core viewpoint or click to open a modal with the full profile. To encourage equal attention to all voices, all leaves initially appear in the same gray color and size.

To prevent users from passively accepting AI-generated I-positions, we introduced a co-creation process that allows users to refine and personalize them. The “Story Enrichment” feature generates three scaffolding questions by identifying gaps in the current narrative (see Appendix B.1 for the full prompt). As users answer these questions, the leaf agents are enriched with the user’s own voice, reflecting their specific context rather than a generic description (Figure 5-(b)). Users can also modify I-positions through direct editing or add and delete leaf features.

Finally, users can engage in “1:1 Dialogue” with each leaf agent (Figure 5-(c)). This dialogue allows users to question, challenge, or explore an I-position’s motivations without interference from other voices. Each leaf agent maintains its core identity to express consistent perspectives even when challenged (see Appendix B.2 for the full prompt).

3.3.2 Stage 2: Relational Positioning. This stage focuses on externalizing the relationships among I-positions, allowing users to spatially compose their inner landscape. Once users have explored and enriched their I-positions, they could begin to arrange them spatially to express relationships among their inner voices (Figure 6). Users can freely adjust the position, size, and color of each leaf. Unlike physical stones, these attributes can be easily changed, allowing users to fluidly explore and reshape their inner landscape in personally meaningful ways.

3.3.3 Stage 3: Dialogical Exchange. This stage aims to facilitate dynamic dialogue among I-positions, allowing users to observe or mediate conversations between their inner voices. This multi-agent conversation begins when users select two leaf agents for dialogue (Figure 7-Left). The LLM then analyzes their relationship and generates three tailored discussion topics based on their dynamics (see Appendix C.1 for the full prompt):

- **Conflict:** Questions to explore value clashes, seek compromise and navigate dilemmas.
- **Complementary:** Questions on how distinct aspects can work together.
- **Unrelated:** Questions to explore the diversity and unique motivations of each self.

After a user selects a topic, the group conversation begins and is coordinated by an orchestrator module implemented through prompt engineering rather than a separate multi-agent framework.

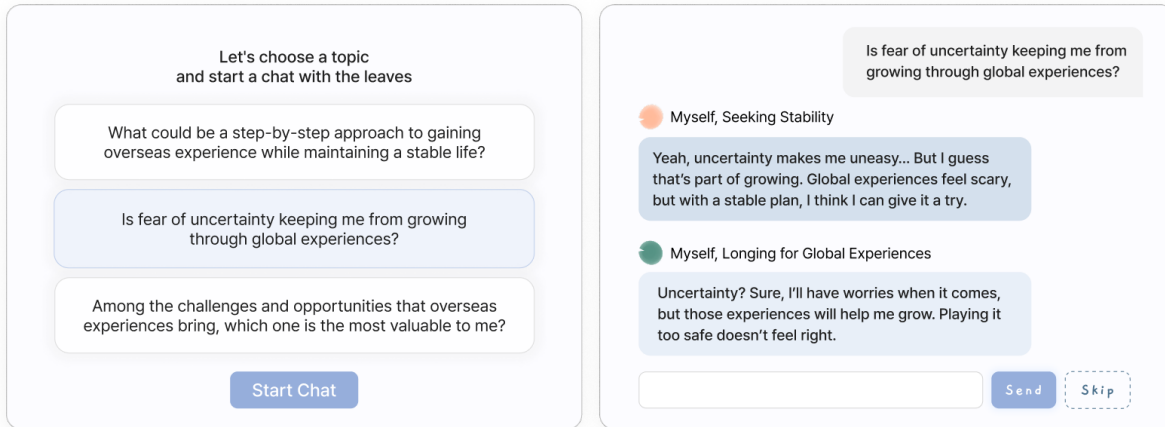


Figure 7: The interface for [Stage 3: Dialogical Exchange]: (Left) After selecting two leaf agents, the system suggests three discussion topics based on their relationship; (Right) the group chat where users can participate via “Send” or observe via “Skip.”

The orchestrator functions as a control prompt that tracks conversational context and interaction state, managing turn-taking among the two leaf agents and the user to prevent any single agent from dominating. For example, when a direct question is posed, the addressed agent is prompted to respond first; when a challenging argument arises, the orchestrator issues a targeted prompt inviting the opposing agent to refute.

Users can engage in the conversation through two modes (Figure 7-Right). In observation mode, activated via the *Skip* feature, the system injects an intervention message—“Do not repeat viewpoints; engage more deeply with each other’s perspectives”—to encourage sustained dialogue between agents. In mediation mode, users intervene directly through their own input; the agents then respond while maintaining their core identities, allowing users to guide the direction of the dialogue.

3.3.4 Stage 4: Reflective Snapshot. This stage intends to allow users to capture a snapshot of their current inner landscape. By clicking the *save pond* button, users can capture their current pond configuration—including leaf arrangement and visual attributes—as a timestamped image file (e.g., “{user}’s InnerPond_{timestamp}”). Users can revisit past inner landscapes at any time, helping them view the self not as fixed but as an evolving narrative.

3.3.5 Technical Configuration. InnerPond was implemented as a web-based application using TypeScript [76] and Next.js framework [117]. The system consisted of a frontend for user interactions and a backend managing user data and LLM APIs. All interaction data was stored in MongoDB [81], with access restricted to the research team only. For the generative pipeline, we utilized Anthropic’s Claude-3.5-Sonnet, selected for its strong performance in Korean among models available at the time of the study [52, 59].

The system’s generative functionality is organized into three LLM-driven pipelines, each implemented through prompt engineering rather than a dedicated multi-agent framework. First, *I-position Extraction* generates initial, personalized I-positions from

pre-survey data. Second, *Single-Agent Interactions* support story enrichment and one-to-one dialogue between the user and an individual leaf agent. Third, *Multi-Agent Orchestration* coordinates dialogue among multiple leaf agents by using a control prompt that manages turn-taking and interaction flow based on conversational context. Selected prompts are included in the Appendix A-C; full prompts are available at <https://github.com/syou-b/innerpond>.

4 Methods

To explore how people experience and make sense of engaging with their multiple selves (RQ2), we conducted a qualitative exploration with participants actively deliberating between two career paths, using InnerPond as a research probe for enabling inter-self dialogue.

4.1 Participants

We recruited 17 participants (7 male, 10 female, average age = 27.18 ($SD=4.93$, Min=21, Max=37)) who were actively deliberating between two career directions, providing a naturally high-stakes context rich in internal conflict and self-negotiation. Recruited via university communities, participants were either (a) undergraduate or graduate students expected to graduate within three years, or (b) recent graduates currently exploring career opportunities while unaffiliated with any educational institution or workplace.

4.2 Study Procedure

The three-week study involved three sessions: (1) pre-survey, (2) in-person session, and (3) follow-up interview. All research materials and procedures were approved by the Institutional Review Board (IRB) of the hosting university. The entire study, from materials and sessions to data analysis, was conducted in Korean. For the analysis and writing process, participants’ quotes were translated into English by bilingual researchers on our research team.

4.2.1 Pre-survey. To construct knowledge of participants’ multiple selves relevant to their career deliberations, participants completed

an online pre-survey seven days prior to the in-person session. The survey consisted of two main parts:

- **Personal Information Survey (5 min):** Participants provided demographics and completed personality assessments using the BFI-2-S [109] and work values assessment through the SWVI [98]. Participants were also asked to list three personal strengths and weaknesses.
- **Career Context Survey (5 min):** Participants briefly described the two career paths under consideration through short, open-ended responses. For each path, they detailed: (1) the background and rationale for considering it, (2) their expectations and concerns, and (3) any preparation processes undertaken.

This information was used to generate agents representing participants' I-positions prior to the in-person session (see Appendix A.1 for an example).

4.2.2 In-person Session. We conducted the in-person session in a quiet room with two researchers for approximately 120 minutes. To foster a relaxed and introspective atmosphere, the lab space was set with meditative music and dim lighting. During the session, one researcher guided participants and conducted interviews (facilitator), while the other observed participant behavior and system interactions in real-time (note-taker), documenting notable moments and emerging patterns in field notes.

As the session began, participants provided informed consent after being briefed on the study's objectives and procedures. To allow the observing researcher to monitor interactions without disrupting the participant, participants shared their screens via Zoom. The seven-stage session—comprising an initial interview, onboarding, four InnerPond activities each followed by a post-activity interview, and an exit interview—was designed to investigate participants' evolving experiences. In particular, post-activity interviews explored participants' in-the-moment experiences by probing specific decisions and interaction patterns observed during each activity. To facilitate this, the facilitator used the documented field notes as prompts to help participants recall their process, jointly revisiting moments of interest to understand the reasoning behind their actions. All interviews were conducted using a semi-structured interview guide, with probing and follow-up questions informed by participants' responses and observed interaction patterns (see Appendix D).

- (1) **Initial Interview (5 min):** We began with a brief interview to understand participants' current career concerns and how they typically engaged in inner dialogue.
- (2) **Onboarding (5 min):** We introduced InnerPond, explaining that their inner voices—extracted from the pre-survey—were represented as 'lotus leaves.' We presented an overview of the four activity stages, along with guidelines for the main features of each stage. Participants were informed that interviews would follow each activity and that researchers would be available to assist them.
- (3) **I-position Construction (25 min activity + 15 min interview):** After onboarding, participants were asked to examine the extracted leaves in their InnerPond and articulate them. The post-activity interview explored the authenticity

of the extracted I-positions and their motivations for editing, adding, or removing leaves (e.g., "*Which leaf did you feel represented you the best, and why?*", "*Why made you delete the [Myself, ...] leaf?*").

- (4) **Relational Positioning (5 min activity + 10 min interview):** Participants were instructed to visually express their inner landscape by adjusting the size, color, and position of leaves in the pond, while thinking-aloud to verbalize their thoughts. The post-activity interview explored the rationale behind their pond structure and design, probing specific choices observed during the activity (e.g., "*When arranging the leaves, what did you consider and why did you choose their positions?*", "*What made you not resize the lotus leaves at all?*").
- (5) **Dialogical Exchange (15 min activity + 10 min interview):** Participants were instructed to create and engage in group conversations with different leaf combinations and topics. The post-activity interview focused on their group conversation experience, their criteria for selecting leaves, and the suitability of LLM-suggested topics, as well as observed conversation patterns (e.g., "*Why did you focus most of your time on conversations in this group chat?*").
- (6) **Reflective Snapshot (10 min activity + 5 min interview):** In the final activity, participants freely explored previous stages at their own pace, revisiting and engaging with their I-positions as they wished. Afterward, participants were asked to save their InnerPond landscape using the 'save pond' feature. The brief post-activity interview explored what participants chose to revisit (e.g., "*What did you mainly do during the free exploration time?*" and their experience of capturing a snapshot "*How did you feel when you saved your pond at the end?*").
- (7) **Exit Interview (10 min):** After completing all stages, an exit interview was conducted about the overall InnerPond experience. We inquired about participants' most memorable moments, the perceived naturalness of the dialogue, and any new insights they gained. We also explored how the InnerPond experience impacted their self-understanding. Upon completion of the session, participants were compensated with 40,000 Korean Won (equivalent to 28 USD).

4.2.3 Follow-up Interview. Two weeks after the in-person session, we conducted a follow-up phone interview to explore how participants reflected on their experiences with InnerPond following the session. Although participation in the follow-up interview was not mandatory, all but one of the 17 participants chose to take part. Following the semi-structured approach, the interview focused on whether the insights from the session persisted or evolved, whether participants noticed new patterns in how they engaged with their inner voices in daily life, and any shifts in their approach to career deliberation.

4.3 Data Analysis

4.3.1 Data Preparation. We documented participant behavior and experiences through multi-faceted data: (1) system logs and (2) qualitative data from in-person sessions and follow-up interviews. System logs provided detailed behavioral traces across all stages,

including I-position profiles, their modifications, and chat messages with individual leaves (Stage 1), spatial configurations of the pond (Stage 2), and group conversation logs with leaf combinations and discussion topics (Stage 3) (see Appendix E for detailed data description). Qualitative data included field notes documenting participant behavior and screenshots during activities, think-aloud transcripts from Stage 2, and all interview transcripts. All interviews were transcribed by the two researchers who conducted the study. We synthesized all data into a master document for each participant, integrating system logs and qualitative data. Two researchers collaboratively constructed these documents and visually organized the materials on FigJam [111] to support systematic reference throughout the analysis.

4.3.2 Thematic Analysis. We employed thematic analysis [17] to systematically analyze the collected data. Our approach combined deductive and inductive elements: we used the four stages of InnerPond as an overarching deductive framework to segment participants' experiences and behaviors. Within each stage, we performed open coding and inductively developed themes from the data. The analysis primarily operated at the semantic level, focusing on participants' explicit accounts of their experiences, while also attending to latent meanings when interpreting the underlying motivations and significance that participants attributed to their interactions. The entire research team conducted this analysis through weekly meetings, continuously cross-referencing insightful excerpts from interview transcripts with corresponding system logs to develop richer interpretations. Through an iterative process, we grouped the codes into higher-level themes and further developed them into more specific sub-themes to capture emergent patterns.

As a result, we identified key themes across the four stages: the externalization and personalization of inner voices (Stage 1), relational positioning through spatial composition (Stage 2), the orchestration of dialogue among multiple voices (Stage 3), and the preservation of temporal self-portraits (Stage 4). Additionally, through follow-up interviews, we observed an enduring theme of reflective integration into everyday life and career decisions.

5 Findings

This section reports how InnerPond, used as a research probe, shaped the ways participants engaged with their plural selves as they moved through its four designed stages. Sections 5.1–5.4 present findings from the in-person session, organized by each stage, while Section 5.5 reports insights from follow-up interviews conducted two weeks later. Rather than focusing on the stages themselves, we highlight the kinds of experiences that emerged within them—how participants articulated inner voices (Stage 1), made their relationships visible (Stage 2), encountered dialogue among these voices (Stage 3), and captured temporal self-portraits (Stage 4). Lastly, we describe how these experiences continued beyond the sessions and informed participants' everyday reflections and career decisions.

5.1 Co-Creating I-positions Through Externalization and Reconstruction

In Stage 1, participants co-created personalized I-positions from AI-generated fragments of self. Overall, externalizing their inner

world into discrete 'lotus leaves' allowed them to step back and view their psychological states from a distance, creating space for a deeper recognition of their multi-layered identity and the surfacing of overlooked aspects of themselves.

5.1.1 Initial Perceptions of Externalized Selves. Participants' first encounters with the AI-generated leaves began with an immediate scan of how well each one aligned with their lived experience—whether it felt resonant, subtly off, or misaligned. Overall, most participants accepted the majority of their leaves, with over half (9/17) retaining the entire set, while others removed only one or two ($M=1.17$). Many found the externalization itself valuable for introspection. P1 described how seeing her inner world visualized prompted an objective re-evaluation, while P8 found that having her unarticulated thoughts externalized allowed her to explore new self-perceptions, prompting reflections like, *"Oh, I can think this way."*

Many also described a sense of recognition as they scanned the leaves, noting how familiar worries or aspirations were suddenly given form. P7, for instance, was surprised to see her anxieties concretely represented: *"Except for deleting one leaf, every single thing I always worry about was specifically targeted and realized. It felt like [they are] my actual inner voices."* In particular, when the leaves reflected specific, personal experiences rather than generic concerns, participants received them with a pronounced sense of authenticity. For example, P11's [Myself, Rewarded by Admin Work] resonated deeply because the narrative reflected her conviction that *"supporting the organization from behind is more meaningful."*

Others noted subtle misinterpretations rather than outright inaccuracy. P5 felt [Myself, Longing for the Prestige of Professorship] overstated her motivations: *"I just thought of it as an honorable job. But this makes it sound like I'm thirsting for social recognition, and that's not quite right."* Her reaction illustrated how even minor discrepancies in framing could shift participants' sense of ownership over a given I-position.

Overly misaligned leaves sometimes prompted resistance, causing some participants to remove them immediately. P7, for example, was decisive in removing [Myself, Anxious about Success]: *"This is not me."* Yet more often, participants chose to retain discordant leaves as prompts for reflection. P3 was initially skeptical of [Myself, Drained by Socializing] but later reframed her stance and chose to keep it: *"At first, I didn't recognize it, but seeing it listed with the same weight [as others leaves] made me consider it more."* Rather than dismissing the unfamiliar voice, she found that its presence surfaced neglected aspects of her career thinking—demonstrating how even imperfect extractions could also serve as catalysts for deeper self-examination.

5.1.2 Elaborating and Personalizing the Extracted Selves. After moving past their initial impressions, participants entered a more active phase of working with the leaves—clarifying, enriching, and reshaping the AI-generated narratives so they better reflected the nuances of their own experiences. In doing so, all participants engaged in story enrichment by answering AI-scaffolded questions ($M=11.82$ times), to refine abstract or generic voices into situated, meaningful selves. Through this process, participants personalized the abstract leaves by grounding them in their own contexts—reshaping initially unconvincing or generic narratives to better align with their values

and perspectives. For example, P3 developed the vague narrative of [Myself, Wanting to Contribute] into a clear professional aspiration. Answering tailored questions like “What specific social problems do you want to solve?” helped her move from an abstract desire to a concrete consideration of which issues she might want to address as a lawyer.

Participants also took direct control by modifying narratives that felt insufficient or misaligned. P6, for instance, edited the LLM’s phrasing to better reflect her perspective, replacing “*economic stability*” with “*economic abundance*.” Furthermore, when participants felt important aspects of their inner voice were missing, they added new leaves to fill these gaps. In total, 12 participants added an average of 1.67 new leaves. P2, for instance, added [Myself, Valuing Office Environment], explaining how his happiness and efficiency depended on tangible workplace factors, like a comfortable atmosphere and good food. P4 went further, intentionally creating [Myself, Believing in Myself] to counterbalance the negative leaf [Myself, Anxious about Being Ordinary], as a move to prevent his inner landscape from being dominated by negative representations.

Through these processes of enrichment, modification, and addition, participants transformed the AI-generated leaves into personalized representations. The resulting leaves ($M=11.06$) were not perceived as static data summaries but as distinct entities, each embodying a different aspect that constituted the self: “*beings similar to me, but with a distinct character—like in the movie ‘Inside Out’*” (P15). In that film, emotions such as Joy and Sadness are portrayed as independent characters living inside a person’s mind. This comparison suggests that the co-creation process gave the I-positions a sense of independent identity, making them feel like genuine inner voices rather than mere labels. Participants also valued the process for surfacing aspects of themselves they typically overlooked—even when the output was not perfectly accurate. P3 reflected: “*With my single consciousness, I can’t pay attention to everything at once. This process helped me extract parts of myself I usually miss. So even if some leaves felt off, it was still meaningful.*” These dynamics illustrate how co-creation shifted the system from simply classifying users to enabling a shared ownership of their plural selves.

5.1.3 Conversing with Individual Selves. With each I-position (leaf) elaborated and structured, the process shifted from organization to directly engaging these selves in one-on-one dialogue. Rather than conversing with all leaves, participants gravitated toward about half of their selves ($M=5.29$), focusing on those where they felt tension, curiosity, or a need for clarity. These conversations took varied forms, shaped by participants’ intentions and their relationship with each self.

Some participants used the one-on-one space for open-ended exploration, treating the self as a partner for thinking through complex issues. P15, for instance, held a collaborative conversation with [Myself, Diving Deep] to shape his research direction, asking questions like, “How can I find an interesting topic?” and “How can I sustain my motivation?” (Figure 8) to facilitate open-ended reflection on the meaning of his research. Rather than seeking definitive answers, he used the dialogue to think aloud with a part of himself he wanted to understand better.

Others approached conversations more strategically. P7, for example, posed the same fundamental questions—like “Should I choose

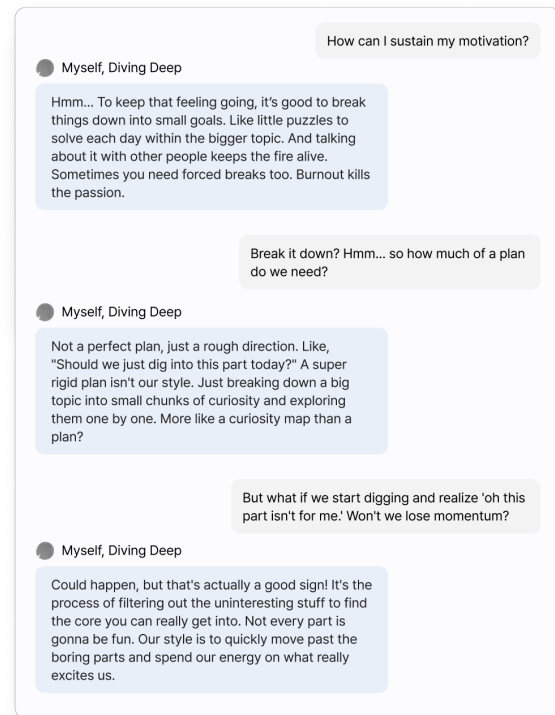


Figure 8: Example of a one-on-one dialogue with [Myself, Diving Deep] (P15).

what I love or what I’m good at?”—to a cluster of related selves, including [Myself, Fearing Uncertainty], [Myself, Pursuing Success], and [Myself, Wanting Recognition]. By collecting responses from multiple I-positions, she compared perspectives and analyze her career concerns from different angles.

The course of dialogue also depended on how familiar participants were with each self’s voice—and how willing they were to acknowledge it. With highly aligned selves, conversations were often used to verify authenticity, sometimes even through adversarial testing. P13, for example, intentionally challenged [Myself, Craving Quick Success] by asking whether he should simply lower his expectations. When the self countered that doing so would only lead to future regret, he found he couldn’t refute it: “*I tried being uncooperative in the conversation, but the voice was perfectly aligned with what’s inside me.*” By contrast, conversations with less familiar or initially rejected selves often unfolded as a process of gradual acceptance. P17 initially denied [Myself, Lacking Expertise], which voiced his professional insecurities. Yet through continued dialogue, he came to recognize a suppressed anxiety: “*At first I kept denying it. But as we talked, I thought it could be a part of me. Maybe there has been a voice of anxiety in my heart, but I didn’t know it because I didn’t have the time or energy to think about it.*”

As dialogues deepened, many participants also varied their conversational style to match the persona of each leaf. P7 discovered she was acting as a mediator, emphasizing possibilities when talking to selves expressing uncertainty, and highlighting risks when engaging with overly optimistic selves. For example, she encouraged




Figure 9: Examples of holistic inner landscapes composed by participants. (a) P4 created a scene of self-compassion, describing one self whispering encouragement to another. (b) P6 arranged the leaves to express gradual personal growth, likened to a musical crescendo.

[Myself, Anxious about Success] by reminding it that abilities can be developed, while cautioning [Myself, Passionate for Challenges] that excitement alone might not sustain long-term commitment.

This adaptive approach was particularly effective in conversations with inherently negative selves, as it allowed participants to objectively view their weaknesses. P5 found it particularly “*fun and helpful to talk with the leaves that contained my flaws*,” such as [Myself, Lacking Perseverance]. She explained that hearing her weaknesses spoken by another entity created emotional distance, making them easier to examine without self-judgment. P12 echoed this sentiment: through conversations, he was compelled to confront contradictions he would normally dismiss, describing it as “*a powerful moment of reflection that came from seeing [his] weaknesses objectively*.”

However, a few participants reported feeling fatigued when a negative self became too entrenched in its character. P11 felt drained after a conversation with [Myself, Tense in Relationships], noting that “*if a self is too negative, it’s tiring, even though it’s me*.” P6 also described [Myself, Hesitant to Decide], as caught in an endless loop of anxiety: “*Because it kept circling around anxiety and hesitation, it didn’t feel like it would offer a solution*.” These experiences suggest that while externalizing negative aspects could facilitate constructive self-understanding, overly repetitive portrayals may become counterproductive.

While most participants actively engaged in one-on-one dialogues and found them interesting and valuable, a few who were not accustomed to introspection found it difficult to initiate conversations with their selves. P4 explained: “*I’m not the type to have thought deeply about myself. When I get a question, I can think about it, but when I tried to start a conversation, the questions didn’t easily come to me*.” His account suggests that open-ended dialogical engagement presupposed a level of self-directed initiative that not all participants felt equipped to sustain.

5.2 Composing Relational Landscapes Through Visual Arrangement

In Stage 2, participants visually composed their ‘inner landscape’ by arranging individual selves as lotus leaves. Building on the clearer understanding of each leaf developed in Stage 1, they moved beyond considering leaves in isolation and began attending to how different aspects of themselves related, contrasted, or clustered together. Through iterative spatial adjustments, participants articulated a more integrated sense of self, expressing relational meanings through variations in the size, color, and position of each lotus leaf.

In particular, many participants adjusted the size of leaves to express the relative importance of different selves. P13 explained: “*I made [Myself, Finding Work I Love] the biggest. It felt like many of my other thoughts stemmed from this one, and it’s a thought I’ve had for a very long time...*” By contrast, a few deliberately avoided assigning relative sizes. P11, for instance, chose to keep all leaves equal, explaining that “*If I make one bigger than the others, it feels like I’m treating my worries unequally. I wanted to see them all on the same level*.”

Participants also adjusted color and brightness to convey emotional qualities or to organize their selves thematically. P16 associated hue directly with emotional tone, using warm colors for positive aspects and cooler shades for negative or uncertain ones. P7 adjusted brightness to indicate with her current values—brighter leaves reflected parts of herself she embraced, while dimmer ones stood for aspects she felt less integrated with. P3, meanwhile, assigned colors to thematic domains, explaining: “*I’ve categorized my dreams from long ago in blue, financial matters in yellow, and important aspects in red*.” Others developed more metaphorical interpretations of color. P4 arranged his leaves into what he called a “*seasonal clock*”, using shifts in color to suggest the temporal change—spring-like greens for emerging aspects of himself, autumnal hues for those he felt were fading (Figure 9-(a)). P5 used color to

symbolize vitality, explaining that she gave “*dead and dying*” shades to the selves tied to uncertainty or fatigue, while assigning bright green to the thoughts that “*made [her] feel alive and motivated*.”

Participants also iteratively positioned the leaves in various ways to convey both the relationships among selves and their relative depth. Some used spatial stacking to express containment or hierarchy: P5 stacked certain leaves on top of others to signal that one self encompassed another. Others used spatial links to represent influence or sequence. P8, for instance, connected leaves in a chain to illustrate how one concern led to another. Still others highlighted mediation: P3 placed one leaf between two conflicting ones, describing it as a means of bridging the tension between them.

Some went further, moving beyond these attributes to imbue the entire composition with holistic meaning. P6, for example, arranged her leaves to represent her ongoing process of growth (Figure 9-(b)), explaining, “*The overall composition is like a crescendo in music. For me, it expressed my gradual growth*.” Similarly, P4 described his landscape as a scene of self-compassion, where one self whispered encouragement to another (Figure 9-(a)). He noted, “*It looks like the one in the back is whispering words of comfort to the one in front, as if a bigger self were encouraging a smaller one*.”

The high degree of freedom in visual composition supported the introspection processes by giving participants greater control over how their thoughts were organized. Many reported feeling empowered to actively manipulate the space to clarify what mattered to them. P11 contrasted this with more passive forms of self-expression: “*Unlike a diary that only expresses emotions, this interface helped me work with my thoughts by rearranging and aligning them. So it felt closer to finding a solution*.” These accounts illustrate how spatial composition afforded a meta-positional perspective, enabling participants to see their plural selves as a connected landscape rather than isolated fragments.

5.3 Orchestrating Multi-Agent Dialogue as Meta-Cognitive Moderator

In Stage 3, participants engaged in dialogical exchanges among their inner voices, exploring how different selves could converse and negotiate with one another. As they observed disagreements emerge and intervened as mediators when needed, they encountered new perspectives on personal dilemmas and became aware of how different voices enacted authority, hesitation, or reconciliation within the self.

5.3.1 Curating Dialogical Encounters Among Selves. As participants began preparing their inner voices for dialogue, they intentionally combined different selves with distinct intentions, forming early ideas about how these voices should meet and interact. One common approach was to bring opposing selves into dialogue to examine internal dilemmas. P3, for instance, paired [Myself, Yearning for Creative Freedom] with [Myself, Desiring Financial Abundance] to confront the tension between passion and security (Figure 10). P7 emphasized the purpose of this approach, explaining: “*My main idea was to put conflicting selves in dialogue. It was more useful for weighing my options*.”

Some participants combined similar or complementary selves to surface subtle distinctions that were not immediately apparent. P8, for instance, paired [Myself, Pursuing a Stable Life] and [Myself,

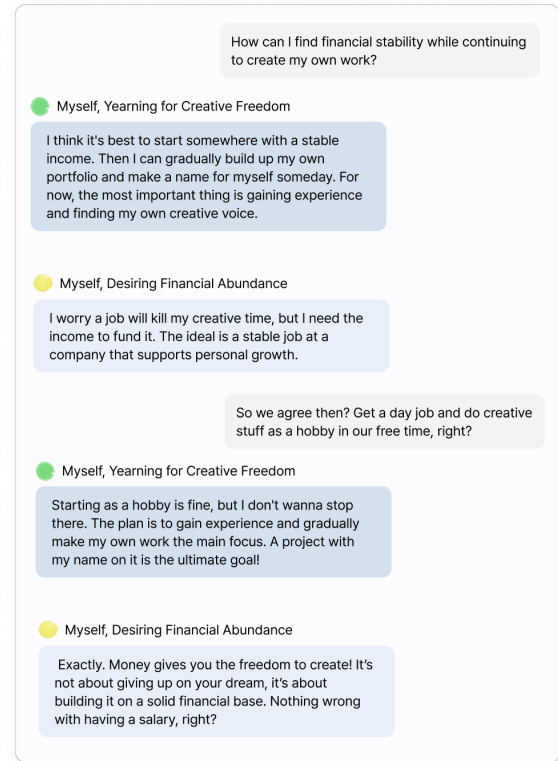


Figure 10: Example of a group conversation where P3 pairs [Myself, Yearning for Creative Freedom] and [Myself, Desiring Financial Abundance].

Sensitive to Others’ Evaluations]—two selves she thought were similar. She reflected: “*I didn’t expect much at first, but I found unexpected answers here. These two actually explained myself best, and I realized they were shaping my recent career thinking together*.” Through such pairings, participants recognized how aligned voices could mutually reinforce and affirm one another.

Some participants also experimented with unrelated selves to probe for hidden connections, driven by curiosity about unexpected insights. For example, P13 paired [Myself, Worried About Being Unplanned] with [Myself, Wanting to Innovate with Robots]: “*I chose these because they seemed unrelated, but seeing their interaction made me realize my disorganization could affect my career path—whether as an obstacle or maybe even as a source of fresh ideas*.” Through this pairing, he recognized that voices he had previously treated as separate were, in practice, working in tandem.

Building on these diverse intentions, participants then shifted their attention to shaping the dialogue itself. Having chosen which selves should meet, they selected one of the AI-suggested topics that most closely captured what they wanted the encounter to reveal or help them think through. Often, these aligned with the questions they already had in mind. As P12 put it: “*When I chose these two leaves, certain topics naturally came to mind, and there was usually something similar among the suggestions*.”

Some selected topics to gain practical guidance on ongoing dilemmas. P10, weighing a PhD path against corporate opportunities, asked: “What are the minimum conditions to continue researching despite reduced income?” for a dialogue between [Myself, Anticipating a PhD] and [Myself, Worried About Finances]. She explained, “*I wanted to find clues for my current concerns.*” For her, the topic served as a way to structure an internal debate that had previously been difficult to articulate. Likewise, P5, a doctoral student considering a professorship, chose: “How can I reduce anxiety while moving forward on this career path?” for her conversation between [Myself, Fearing Tenure Uncertainty] and [Myself, Valuing Autonomy].

Others used the topic suggestions to open up new perspectives rather than resolve immediate concerns. P2, for example, brought together [Myself, Energized by People] and [Myself, Easily Swayed by Emotions] to explore: “What criteria distinguish relationships that energize me from those that drain me?” He noted that he had never thought about this distinction before, and the suggested topic prompted a new line of self-inquiry.

Taken together, this process of pairing selves and selecting topics functioned as a introspective practice in its own right. By choosing which voices to bring into dialogue, participants began to surface the tensions and questions that felt most important to them. Also, by selecting a specific topic, they worked to clarify what they hoped to understand or resolve.

5.3.2 Engaging the Dialogue as Observer and Mediator. Once a topic was chosen, participants engaged with the dialogue in varied ways, moving fluidly between observing and stepping in to mediate. Many began by observing from a distance, often by clicking the ‘Skip’ button to let the selves converse on their own. This observational stance was not merely passive; rather, it offered a vantage point from which participants could see their thought patterns externalized with a level of clarity that solitary introspection rarely afforded. Participants used varied metaphors to describe this experience. P16 described it as watching her mental “*chain of thought*” unfold step by step. P8 called it an “*expanded version*” of solitary thinking—as if her internal monologue had been extended and rendered visible. P17 likened it to a “*proxy battle*,” where conflicting ideas could clash without him being caught in the middle. Across these descriptions, participants agreed that the dialogue gave structure to what had previously felt scattered or circular. P1 reflected: “*Watching my mental conflicts externalized, I realized—this is how I’ve been thinking.*” The dialogue served as a reflected mirror, revealing patterns of internal conflict she had not recognized in concrete terms.

As the conversation progressed, many participants transitioned into a more active role, stepping in to mediate when the dialogue reached an impasse or drifted into repetition. They treated such intervention not as overriding the conversation but as facilitating negotiation among their inner voices. Some guided opposing selves toward common ground, asking questions like “*What do you agree on?*” (P15). Others pushed the dialogue beyond false binaries—P4, for instance, challenged a dichotomy by asking: “*Economic freedom and meaningful contribution aren’t separate things, right?*” However, not all participants found it equally easy to take such an active role. For example, P14, who expected more guidance from the AI, noted that maintaining the dialogue on his own was challenging.

Beyond these challenges, many participants described the multi-voice dialogue as a source of insights they would not have reached on their own. Seeing conflicting selves interact in one shared space revealed unexpected connections and points of integration. P6 described experiencing a sense of exhilaration “*when AI touches on a point I hadn’t considered.*” For P11, the conversation sparked an immediate realization: “*I could pursue my values not just through employment, but by participating in policy contests.*” The dialogue also helped participants adopt a more balanced view of their internal conflicts. P2 noted that, unlike solitary inner dialogue where one side tends to dominate, the multi-voice format allowed him to “*see both opinions more evenly.*” This balanced exposure also helped some clarify their true inclinations. P6 reflected, “*Seeing both answers, I realized which one really appealed to me.*”

5.4 Preserving Introspective Journeys Through Temporal Self-Portraits

In Stage 4, participants preserved their inner landscape by saving their pond as a temporal self-portrait. This final step translated the introspective process into a tangible artifact—a snapshot of how their inner voices were configured at that moment, which they could revisit and reflect upon over time. P5 remarked: “*This is today’s me, and at another time it will look different. Like a diary, I can look back and remember how my mind was arranged then.*” Participants also expected that viewing saved ponds would re-evolve the introspective process engaged during creation. P12 noted: “*Even now, looking at it makes me remember what I was thinking—seeing the leaves makes me say, ‘Oh right!’*”

For others, saving their pond provided a rare sense of accomplishment. Participants who usually lacked time for introspection felt satisfied to have a visible outcome of their efforts. P2 expressed: “*I usually don’t have time to look at myself deeply, but visualizing the result felt good. It gave me a sense of efficacy from having examined myself.*” This satisfaction was reinforced by the freedom to shape the inner landscape without being constrained by external standards. P4 emphasized: “*It was good that there wasn’t a predetermined answer for creating the pond. If it had suggested a ‘desirable pond,’ it would have made things worse. Such freedom to create it in my own way was what helped most.*”

Furthermore, some imagined using successive snapshots to trace how their inner landscape might evolve over time. P10 reflected: “*It would be interesting to look back after accumulating [more snapshots]. You could notice things like, ‘Oh, I’ve been thinking about this consistently.’*” These comments suggest that participants valued not only the immediate snapshot but also its potential as a longitudinal record of change, extending introspection into an ongoing narrative of the self.

5.5 Sustaining Introspective Reflections After the InnerPond Experience

In follow-up interviews conducted two weeks after the in-person session, many participants described the InnerPond experience as something that lingered—shaping how they noticed themselves, interpreted everyday experiences, and approached career decisions.

Rather than fading after the activity, they mentioned that the insights and realizations from the session carried into their everyday thinking.

A recurring theme was a more integrated sense of self. Participants spoke of their inner voices not as scattered fragments but as parts of a larger whole. P1 explained: *“All these lotus leaves come together to form me as a person. I keep thinking that these small elements collectively make up who I am.”* For some, this integrated view translated into a sustained attentiveness to specific inner voices, which continued to surface in everyday thinking rather than fading after the session. P1 found that certain I-positions—such as [Myself, Transcending Limitations] and [Myself, Having Many Worries]—continued to surface in her everyday thinking, keeping her attuned to both her aspirations and anxieties: *“These selves don’t just disappear. I keep being aware of both their strengths and weaknesses.”* This ongoing awareness also made participants more willing to acknowledge aspects of themselves they had previously avoided. P8, for instance, realized through the session that she had been avoiding situations where others might judge her. This awareness persisted beyond the session: *“I’d been avoiding others’ judgments out of fear, and now I want to face them more directly.”* In some cases, this stance further extended into how participants perceived everyday situations. P16, who revisited her saved snapshot on her own, noticed that everyday moments now triggered questions she had never thought to ask—revealing blind spots in what she typically paid attention to.

Beyond internal awareness, some participants described how these reflections carried into their engagement with real-world situations. P4, for instance, found that the experience made his indecisiveness more explicit to himself, which in turn helped external feedback register with greater clarity: *“I had vaguely sensed this was a problem, but seeing it laid out made it unmistakable. So when my professor later said the same thing, it really resonated, and I realized I needed to work on it.”* These reflections also extended into career deliberation. Participants emphasized a shift from chasing external conditions to asking whether a path genuinely aligned with their internal voices and values. P11 reflected: *“When the job market is bad, I sometimes apply to jobs I actually don’t want, just out of anxiety. But after InnerPond I realized, ‘That’s why I kept struggling—it wasn’t what I wanted.’”* Similarly, P13 explained that he had begun weighing career options differently: *“Before, I only thought about making more money. Now I put my traits and thoughts together to see if a path really fits me.”* P16 also described how the experience helped her reconnect with her initial motivations in her field, explaining, *“Remembering what first drew me to the field and what makes it fulfilling helped me change how I rethink my [career] direction.”* Together, these accounts suggest that the InnerPond experience continued to shape participants’ self-understanding—prompting them to notice themselves differently, reframe career choices, and carry reflective insights into everyday life.

6 Discussion

This study investigated how an AI-mediated approach grounded in Dialogical Self Theory (DST) can be translated into an interactive system that supports introspection by making the inherent multiplicity of the self experientially accessible. Through the design and

evaluation of *InnerPond*, we examined both the design considerations for supporting dialogue among multiple selves and how users experience engaging with their plural selves through such a system. Participants interacted with their inner voices—visualized as lotus leaves—in increasingly layered ways: first articulating and refining individual I-positions, then composing relational landscapes, and finally orchestrating dialogues among them. Our findings illustrate how operationalizing DST’s dialogical view of the self in an AI-mediated form—externalizing multiple I-positions as distinct agents that users can observe, arrange, and converse with—can support introspective engagement in practice. Following DST, we use the term *inter-self communication* to refer to dialogical engagement among one’s own I-positions as instantiated in our system.

Intrapersonal communication is usually framed as solitary inner dialogue [85, 118, 119], while interpersonal communication unfolds between separate individuals [7, 18, 53]. Inter-self communication does not introduce a new category of communication; instead, it serves as a descriptive lens for characterizing a particular interactional configuration that becomes visible when DST is instantiated in an interactive system. It remains intrapersonal in that the dialogue unfolds within a single individual, while adopting an explicitly dialogical structure that makes inner multiplicity observable and interactive. In this sense, inter-self communication helps articulate how people engage with multiple facets of the self through dialogical interaction, without departing from or extending beyond DST’s established theoretical framing. This perspective suggests an opportunity for HCI to design introspection tools that acknowledge and support the multiplicity of the self—moving beyond systems that assume a single inner position toward those that facilitate dialogue among multiple perspectives, helping users surface tensions and navigate trade-offs in complex life decisions.

Drawing on this perspective, we discuss broader implications for designing AI-mediated introspection systems: how translating a dialogical view of the self into interactive form shapes users’ introspective experiences, what design considerations emerge from this translation, and what experiential tensions arise when users engage with externalized facets of themselves.

6.1 From Unified to Dialogical Self: New Possibilities for Inner Dialogue

Our study highlights the value of reconsidering how the self is conceptualized in HCI. Much prior work on user models [29, 54] has adopted a view of a centralized, unified self [43]—a framing that has proven useful for many design contexts. Yet when people face complex life decisions, they often experience themselves as pulled in multiple directions, negotiating among competing values and aspirations. Our findings suggest that supporting the self as a “dynamic multiplicity” [41, 45] within interactive systems can support forms of introspective engagement that are difficult to surface when the self is modeled as singular, by making this multiplicity visible and actionable. This view resonates with recent theorizing in personal informatics, where the self has been reconceptualized as dynamic and constructed through ongoing interaction with the world [97]. Our work extends this perspective by operationalizing the self not only as dynamic, but as inherently plural—composed

of multiple voices that can be externalized, arranged, and set in dialogue.

While our study focused on career decision-making, the dialogical mechanism we examined reflects a more general pattern of internal negotiation described in DST, and may extend to other domains where the self is pulled in multiple directions—such as moral dilemmas, relationship conflicts, or lifestyle changes. More broadly, this suggests that AI systems can be designed not to redefine introspection or decision-making, but to scaffold the internal negotiations that precede complex human decisions, supporting users as they explore tensions among competing perspectives rather than converging prematurely on a single answer.

6.2 Designing for Dialogical Introspection

Translating DST’s conceptualization of the self into an interactive system posed distinct design challenges. Our design was guided by three goals: supporting coexistence rather than resolution among inner voices, ensuring equal attention for each voice, and fostering relational connectedness through a meta-positional perspective. Through our iterative design process grounded in these goals, we identified two considerations that may help articulate how dialogical introspection can be supported in practice.

First, we found that metaphors could implicitly define relationships among selves. When we experimented with different metaphors, each carried distinct interactional logics: a ‘debate arena’ implied zero-sum competition where one voice must prevail, while ‘digital stones’ emphasized individuality but obscured the connectedness among perspectives. We ultimately chose the ‘lotus pond’—distinct leaves sharing hidden roots conveyed both independence and interdependence, while the bird’s-eye view invited a meta-positional perspective.

This design exploration taught us that metaphor selection warrants careful consideration, as it could predispose users toward particular relational dynamics—competition, isolation, or integration—before any interaction begins. This aligns with DST’s view of the self as a “society of mind,” in which I-positions can engage in diverse dynamics—from conflicts to coalitions and cooperation [44]. For introspection support systems, this suggests attending to metaphors as a way of shaping how users perceive and explore relationships among inner voices, rather than treating them as neutral visual choices.

Second, because articulating one’s inner multiplicity can be demanding, we sought to balance system guidance with user agency. Without support, users may struggle to surface multiple aspects of the self from scratch or may default to familiar, dominant perspectives. At the same time, overly directive AI behavior risks imposing externally generated interpretations that can displace users’ sense of ownership over their own inner narratives. To navigate this tension, we designed the system to offer AI-generated I-positions as provisional entry points rather than authoritative representations—distributed across different facets of the self to lower the barrier to self-articulation while explicitly inviting user judgment. Users could validate, refine, or discard these suggestions, ensuring that the resulting inner narrative remained authentically their own. Our findings show that meaningful introspective engagement often emerged precisely in this space of negotiation, where users actively

questioned, revised, or resisted AI-initiated content rather than accepting it.

This resonates with DST’s therapeutic emphasis on surfacing positions that are “less dominant but vital”—voices that may otherwise remain inaccessible yet hold potential for a more integrated self-understanding [65]. For designers of introspection support tools, this points to scaffolding that helps users access less familiar aspects of themselves without collapsing the balance toward AI-led interpretation or unstructured self-reflection.

6.3 Engaging with Externalized Selves: Tensions and Trade-offs

While our metaphor and scaffolding choices guided the design of InnerPond, participants’ actual engagement with their externalized selves revealed a set of tensions that emerge when AI mediation meets users’ introspective agency. In practice, supporting dialogical introspection involved navigating trade-offs between structure and openness, guidance and authorship, and productive distance from the self versus emotional fatigue. These tensions did not indicate design failures; rather, they surfaced as constitutive challenges of engaging with multiple selves through an AI-mediated system and offer insight into the experiential boundaries of dialogical introspection.

6.3.1 *Between Accurate Reflection and Productive Misalignment.*

A central tension concerned the degree of alignment between AI-generated I-positions and users’ existing self-perceptions. Much prior work in agent design assumes that stronger alignment is inherently beneficial, fostering trust, rapport, and recognition between user and system [22, 89, 120]. In our study, however, alignment functioned less as a straightforward design objective and more as a trade-off with distinct introspective consequences. On one hand, participants often valued leaves that resonated with familiar concerns, finding validation in seeing their anxieties or aspirations concretely represented. On the other hand, high alignment sometimes stabilized tentative or limiting self-concepts, reinforcing existing interpretations rather than inviting reconsideration [70].

Meanwhile, instances of misalignment—where the AI mirrored the user in unexpected ways—were not always experienced as errors to be corrected. Some participants initially did not recognize certain I-positions, but chose to retain them anyway, later finding that these unfamiliar voices surfaced neglected aspects of their thinking. This pattern aligns with what prior work has termed “productive unfamiliarity” [37, 38]. Taken together, these findings suggest that introspection support systems need not optimize solely for accurate reflection; instead, allowing for carefully bounded misalignment can create opportunities for user negotiation, reinterpretation, and deeper self-inquiry.

6.3.2 *Between Consistency and Adaptability in AI Personas.*

Character consistency is often emphasized as a core requirement in LLM agent design, fostering believability and coherent interaction [33, 39, 103]. In our study, however, consistency became a double-edged requirement when agents represented facets of the self rather than external characters. On one hand, rigid consistency gave agents credibility, making them feel like distinct and recognizable voices. On the other hand, it sometimes trapped users in

unproductive loops—particularly with anxious or negative personas that circled around the same concerns without evolving.

Yet the solution would not simply be to have agents shift their stance whenever users push back. Recent work on LLM sycophancy—where models excessively accommodate user preferences at the cost of truthfulness [104]—illustrates the risk: agents that yield too readily lose credibility as distinct perspectives. In the context of dialogical introspection, this creates a tension between preserving the integrity of an I-position and allowing it to change in response to the user’s evolving understanding. If I-positions remain entirely static, they cannot reflect the developmental nature of introspection; if they adapt too readily, they cease to function as meaningful counterparts in dialogue. This suggests rethinking consistency not as immutability, but as coherent evolution—maintaining a recognizable standpoint while allowing shifts that mirror the user’s ongoing negotiation among multiple selves.

6.3.3 Between User Agency and Guidance. A key design goal of InnerPond was to give users control over how they explored their inner world—allowing them to freely create, edit, and arrange their I-positions. Most participants engaged actively with this freedom, finding it valuable for self-directed exploration. However, some participants’ experiences showed that this openness did not support everyone equally. While many users readily initiated conversations and navigated among their I-positions, others—particularly those less accustomed to introspection—struggled to decide how to proceed, hesitating to initiate dialogue or expecting more direction from the system.

These accounts point to a tension between providing freedom and enabling a felt sense of control. Bennett and colleagues distinguish between *material agency*—the range of actions a system makes available—and *experiential agency*—the user’s felt capacity to act meaningfully within that space [10, 24]. In InnerPond, material agency was consistently high, but experiential agency varied depending on users’ familiarity with introspective practices. This suggests that simply offering open-ended interaction would not always be sufficient for users to feel capable of engaging productively with their externalized selves.

These moments highlight a recurring design challenge for AI-mediated introspection: leaving users fully in control can support self-authorship, yet can also leave some users uncertain about how to begin or how much structure is appropriate. Rather than resolving this tension by privileging either autonomy or guidance, our findings suggest the value of adaptive scaffolding—modulating the system’s level of guidance in response to users’ confidence and engagement. Such an approach may allow AI support to recede as users gain momentum, while remaining available when hesitation or uncertainty emerges, preserving user agency without abandoning them to unstructured self-reflection.

6.4 Ethical Considerations and Potential Risks

AI-mediated introspection raises ethical questions that deserve careful attention, particularly because such systems intervene in how users engage with and make sense of their own inner voices. Unlike general-purpose chatbots that respond to external queries or preferences, InnerPond generates representations of who the

user is—externalized facets of identity that users then reflect on and negotiate with. From a DST perspective, this is consequential: if the self is dialogically constructed rather than fixed, AI-generated I-positions are not neutral mirrors, but active elements that may shape ongoing processes of self-understanding.

In designing InnerPond, we took several steps to mitigate this risk. AI-generated I-positions were framed as provisional starting points rather than authoritative descriptions, and users were explicitly encouraged to validate, refine, or remove them. The multi-agent structure aimed to preserve plurality rather than collapse identity into a single, system-driven narrative, and users retained control over their final inner landscape. These choices position AI as a facilitator of introspection rather than an arbiter of identity, yet they do not eliminate risk. Users may gravitate toward familiar narratives or accept unfamiliar ones without sufficient introspection [69, 90], and because inner voices carry emotional weight, the psychological stakes of AI-mediated influence may be higher than in other human–AI interactions.

Importantly, these risks are unlikely to be evenly distributed across users. Prior work suggests that age, experience, and psychological context shape how people engage with AI-mediated self-exploration [50, 95]. Vulnerable populations—including adolescents or individuals experiencing mental health challenges—may be particularly susceptible to internalizing AI-generated self-representations, calling for additional safeguards such as content moderation, session limits, or integration with human support.

Privacy presents a related concern. Externalizing inner dialogue generates data that reflects internal tensions, aspirations, and vulnerabilities. As AI-mediated introspection systems become more effective at eliciting rich self-disclosure, questions of data protection, retention, and secondary use become increasingly consequential [87, 106].

Overall, these considerations underscore the importance of designing AI-mediated introspection systems that keep users in control of meaning-making, maintain transparency about the provisional nature of AI-generated content, and establish clear boundaries around data use.

6.5 Limitations and Future Directions

This study offered insights into AI-mediated multi-self dialogue, but several limitations point to future directions. First, our findings are drawn from a sample of South Korean university students and recent graduates, limiting generalizability across age groups and cultural contexts. South Korea’s generally positive orientation toward AI [58, 72] may have shaped participants’ openness to AI-generated selves; users in cultures with more cautious attitudes toward AI may respond differently. Therefore, future work should examine how age, cultural background, and prior experiences with AI shape receptiveness to and engagement with AI-mediated introspection.

Second, the single-session design restricted examination of temporal dynamics. A longitudinal approach could extend InnerPond to continuous self-documentation, allowing I-positions to evolve, merge, and fade over time. Also, while participants found engaging with multiple I-positions qualitatively distinct from monological self-talk, our study did not include a direct comparison with single-agent alternatives. Future comparative studies could help clarify

the specific contributions of a multi-agent structure to introspective experience.

Finally, while InnerPond focused on career reflection, AI-mediated inter-self dialogue may extend to other domains—such as mental health support and education—where externalizing internal voices could support sense-making. However, extending this approach requires careful consideration. For example, in settings involving acute psychological distress, AI-mediated reflection would demand clinical oversight and clear boundaries around when professional intervention is required [75, 114]. More broadly, the approach may be less suitable where immediate or directive guidance is needed rather than open-ended exploration.

7 Conclusion

This work introduced InnerPond, an AI-mediated multi-agent system designed as a design probe to externalize and structure inner multiplicity through dialogical interaction. Grounded in Dialogical Self Theory, the system surfaced multiple I-positions as distinct yet connected voices, enabling participants to engage with themselves not as a single entity but as a constellation of perspectives. Through staged interactions of co-creation, relational composition, and dialogue, participants were able to surface overlooked inner voices, articulate relationships among competing perspectives, and actively negotiate tensions within the self, particularly in the context of career reflection.

Across participants' engagements, our findings highlight several characteristics of AI-mediated dialogical introspection. Externalizing inner voices created productive distance, allowing participants to examine familiar concerns with reduced self-judgment, while encounters with partially misaligned or unfamiliar voices prompted reinterpretation rather than rejection. Dialogical exchanges among I-positions supported meta-positional reflection, enabling participants to move between observing their inner dynamics and intervening as mediators. At the same time, these interactions revealed recurring tensions and trade-offs—between accurate reflection and productive misalignment, consistency and adaptability in AI personas, and user agency and system guidance—that shaped how introspection unfolded in practice.

By situating InnerPond as a concrete instantiation of this approach, we contribute design knowledge about how AI-mediated systems can support engagement with a plural self, while also foregrounding the experiential limits, ethical considerations, and interpretive responsibilities that accompany AI participation in ongoing processes of introspection and self-understanding.

Acknowledgments

This work was supported by the SNU-Global Excellence Research Center establishment project at Seoul National University and by the Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No.RS-2021-II211343, Artificial Intelligence Graduate School Program, Seoul National University).

References

- [1] Amal Abdulrahman, Deborah Richards, and Ayse Aysin Bilgin. 2023. Changing users' health behaviour intentions through an embodied conversational agent delivering explanations based on users' beliefs and goals. *Behaviour & Information Technology* 42, 9 (2023), 1338–1356.
- [2] Ben Alderson-Day, Kaja Mitrenga, Sam Wilkinson, Simon McCarthy-Jones, and Charles Fernyhough. 2018. The varieties of inner speech questionnaire—Revised (VISQ-R): Replicating and refining links between inner speech and psychopathology. *Consciousness and cognition* 65 (2018), 48–58.
- [3] Tami Amir and Itamar Gati. 2006. Facets of career decision-making difficulties. *British Journal of Guidance & Counselling* 34, 4 (2006), 483–503.
- [4] James Arnéra, Chun Hei Michael Chan, and Mauro Cherubini. 2024. Digital, Analog, or Hybrid: Comparing Strategies to Support Self-Reflection. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference*. 3435–3452.
- [5] Sanghwan Bae, Donghyun Kwak, Sungdong Kim, Donghoon Ham, Soyoung Kang, Sang-Woo Lee, and Woomyoung Park. 2022. Building a Role Specified Open-Domain Dialogue System Leveraging Large-Scale Language Models. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 2128–2150.
- [6] Eric PS Baumer. 2015. Reflective informatics: conceptual dimensions for designing technologies of reflection. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems*. 585–594.
- [7] Steven A Beebe, Susan J Beebe, Mark V Redmond, and Lisa Salem-Wiseman. 2002. *Interpersonal communication: Relating to others*. Allyn and Bacon Boston.
- [8] Russell W Belk. 1988. Possessions and the extended self. *Journal of consumer research* 15, 2 (1988), 139–168.
- [9] Russell W Belk. 1990. The role of possessions in constructing and maintaining a sense of past. *Advances in consumer research* 17, 1 (1990).
- [10] Dan Bennett, Oussama Metatla, Anne Roudaut, and Elisa D Mekler. 2023. How does HCI understand human agency and autonomy?. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–18.
- [11] Marit Bentvelzen, Paweł W Woźniak, Pia SF Herbes, Evropi Stefanidi, and Jasmin Niess. 2022. Revisiting reflection in HCI: Four design resources for technologies that support reflection. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 6, 1 (2022), 1–27.
- [12] Kirsten Boehner, Janet Vertesi, Phoebe Sengers, and Paul Dourish. 2007. How HCI interprets the probes. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 1077–1086.
- [13] Edwin G Boring. 1953. A history of introspection. *Psychological bulletin* 50, 3 (1953), 169.
- [14] Veronique Boudreault, Christiane Trotter, and Martin D Provencher. 2018. Investigation of the self-talk of elite junior tennis players in a competitive setting. *International Journal of Sport Psychology* 49, 5 (2018), 386–406.
- [15] Nico Brand, William Odom, and Samuel Barnett. 2021. A design inquiry into introspective AI: surfacing opportunities, issues, and paradoxes. In *Proceedings of the 2021 ACM Designing Interactive Systems Conference*. 1603–1618.
- [16] Nico Brand, William Odom, and Samuel Barnett. 2023. Envisioning and understanding orientations to introspective AI: Exploring a design space with Meta. Aware. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–18.
- [17] Virginia Braun and Victoria Clarke. 2012. Thematic analysis. American Psychological Association. (2012).
- [18] Brant R Bureson. 2010. The nature of interpersonal communication. *The handbook of communication science* 1, 2 (2010), 145–163.
- [19] Alex Byrne. 2005. Introspection. *Philosophical topics* 33, 1 (2005), 79–104.
- [20] Sena Çerçi, Marta E. Cecchinato, and John Vines. 2021. How design researchers interpret probes: Understanding the critical intentions of a designerly approach to research. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–15.
- [21] Amy Yo Sue Chen, William Odom, Ce Zhong, Henry Lin, and Tal Amram. 2019. Chronoscope: designing temporally diverse interactions with personal digital photo collections. In *Proceedings of the 2019 on designing interactive systems conference*. 799–812.
- [22] Jiangjie Chen, Xintao Wang, Rui Xu, Siyu Yuan, Yikai Zhang, Wei Shi, Jian Xie, Shuang Li, Ruihan Yang, Tinghui Zhu, et al. 2024. From persona to personalization: A survey on role-playing language agents. *arXiv preprint arXiv:2404.18231* (2024).
- [23] Choo Mui Cheong, Jiahuan Zhang, Yuan Yao, and Xinhua Zhu. 2022. The role of gender differences in the effect of ideal L2 writing self and imagination on continuation writing task performance. *Thinking Skills and Creativity* 46 (2022), 101129.
- [24] David Coyle, James Moore, Per Ola Kristensson, Paul Fletcher, and Alan Blackwell. 2012. I did that! Measuring users' experience of agency in their own actions. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 2025–2034.
- [25] Rahul R Divekar. 2024. Externalizing Internal Conversations: Toward a New Paradigm of Interacting with Our Internal Voice via an External Technological Interface. In *International Conference on Human-Computer Interaction*. Springer, 53–64.

- [26] Tijs Duel, David M Frohlich, Christian Kroos, Yong Xu, Philip JB Jackson, and Mark D Plumbley. 2018. Supporting audiography: Design of a system for sentimental sound recording, classification and playback. In *International Conference on Human-Computer Interaction*. Springer, 24–31.
- [27] Robert Elliott and Leslie S Greenberg. 1997. Multiple voices in process-experiential therapy: Dialogues between aspects of the self. *Journal of Psychotherapy Integration* 7, 3 (1997), 225.
- [28] Chris Elsdén, David S Kirk, and Abigail C Durrant. 2016. A quantified past: Toward design for remembering with personal informatics. *Human-Computer Interaction* 31, 6 (2016), 518–557.
- [29] Cathy Mengying Fang, Phoebe Chua, Samantha WT Chan, Joanne Leong, Andria Bao, and Pattie Maes. 2025. Leveraging AI-Generated Emotional Self-Voice to Nudge People towards their Ideal Selves. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–20.
- [30] Cathy Mengying Fang, Yasith Samaradivakara, Pattie Maes, and Suranga Nanayakkara. 2025. Mirai: A Wearable Proactive AI" Inner-Voice" for Contextual Nudging. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*. 1–9.
- [31] Charles Fernyhough and Anna M Borghi. 2023. Inner speech as language process and cognitive tool. *Trends in cognitive sciences* 27, 12 (2023), 1180–1193.
- [32] Pascal Frank, Anna Sundermann, and Daniel Fischer. 2019. How mindfulness training cultivates introspection and competence development for sustainable consumption. *International Journal of Sustainability in Higher Education* 20, 6 (2019), 1002–1021.
- [33] Ivar Frisch and Mario Giulianelli. 2024. LLM Agents in Interaction: Measuring Personality Consistency and Linguistic Alignment in Interacting Populations of Large Language Models. In *Proceedings of the 1st Workshop on Personalization of Generative AI Systems (PERSONALIZE 2024)*. 102–111.
- [34] Julian Fritsch, Katharina Feil, Darko Jekauc, Alexander T Latinjak, and Antonis Hatzigeorgiadis. 2024. The relationship between self-talk and affective processes in sports: A scoping review. *International Review of Sport and Exercise Psychology* 17, 1 (2024), 482–515.
- [35] Georgia Gkantona. 2023. My inner world: Analyzing the client's self-dialogicality with the method of internal multi-actor performance. *Journal of Constructivist Psychology* 36, 3 (2023), 401–419.
- [36] Taicheng Guo, Xiuying Chen, Yaqi Wang, Ruidi Chang, Shichao Pei, Nitesh V. Chawla, Olaf Wiest, and Xiangliang Zhang. 2024. Large language model based multi-agents: a survey of progress and challenges. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence (IJCAI '24)*. Article 890, 10 pages.
- [37] Brett A Halperin and Stephanie M Lukin. 2024. Artificial Dreams: Surreal Visual Storytelling as Inquiry Into AI "Hallucination". In *Proceedings of the 2024 ACM Designing Interactive Systems Conference*. 619–637.
- [38] Oussama H Hamid. 2024. Beyond probabilities: unveiling the delicate dance of large language models (LLMs) and AI-hallucination. In *2024 IEEE Conference on Cognitive and Computational Aspects of Situation Management (CogSIMA)*. IEEE, 85–90.
- [39] Senyu Han, Lu Chen, Li-Min Lin, Zhengshan Xu, and Kai Yu. 2024. IBSEN: Director-Actor Agent Collaboration for Controllable and Interactive Drama Script Generation. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 1607–1619.
- [40] HJ Hermans. 2010. M. Dialogical self theory: positioning and counter positioning in a globalizing society/HJM Hermans, A. Hermans Konopka.
- [41] Hubert JM Hermans. 1996. Voicing the self: From information processing to dialogical interchange. *Psychological bulletin* 119, 1 (1996), 31.
- [42] Hubert JM Hermans. 2001. The dialogical self: Toward a theory of personal and cultural positioning. *Culture & psychology* 7, 3 (2001), 243–281.
- [43] Hubert JM Hermans. 2003. The construction and reconstruction of a dialogical self. *Journal of constructivist psychology* 16, 2 (2003), 89–130.
- [44] Hubert JM Hermans. 2014. Self as a society of I-positions: A dialogical approach to counseling. *The Journal of Humanistic Counseling* 53, 2 (2014), 134–159.
- [45] Hubert JM Hermans and Thorsten Gieser. 2011. *Handbook of dialogical self theory*. Cambridge University Press.
- [46] Elwin Hofman. 2016. How to do the history of the self. *History of the Human Sciences* 29, 3 (2016), 8–24.
- [47] Sirui Hong, Mingchen Zhuge, Jonathan Chen, Xiwu Zheng, Yuheng Cheng, Ceyao Zhang, Jinlin Wang, Zili Wang, Steven Ka Shing Yau, Zijuan Lin, et al. 2024. MetaGPT: Meta programming for a multi-agent collaborative framework. *International Conference on Learning Representations, ICLR*.
- [48] Jürgen Hoyer and Andreas Klein. 2000. Self-reflection and well-being: is there a healthy amount of introspection? *Psychological Reports* 86, 1 (2000), 135–141.
- [49] Hilary Hutchinson, Wendy Mackay, Bo Westerlund, Benjamin B Bederson, Allison Druin, Catherine Plaisant, Michel Beaudouin-Lafon, Stéphane Conversy, Helen Evans, Heiko Hansen, et al. 2003. Technology probes: inspiring design for and with families. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 17–24.
- [50] Ulugbek Vahobjon Ugli Ismatullaev and Sang-Ho Kim. 2024. Review of the factors affecting acceptance of AI-infused systems. *Human factors* 66, 1 (2024), 126–144.
- [51] Stephani AB Jahn. 2018. Using collage to examine values in college career counseling. *Journal of College Counseling* 21, 2 (2018), 180–192.
- [52] Woori Jang and Seohyon Jung. 2024. Evaluating LLM Performance in Character Analysis: A Study of Artificial Beings in Recent Korean Science Fiction. In *Proceedings of the 4th International Conference on Natural Language Processing for Digital Humanities*. 339–351.
- [53] Arthur Jensen and Sarah Trenholm. 1992. *Interpersonal communication*. Wadsworth.
- [54] Hayeon Jeon, Suhwoo Yoon, Keyeun Lee, Seo Hyeon Kim, Esther Hehsun Kim, Seonghye Cho, Yena Ko, Soeun Yang, Laura Dabbish, John Zimmerman, et al. 2025. Letters from Future Self: Augmenting the Letter-Exchange Exercise with LLM-based Agents to Enhance Young Adults' Career Exploration. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–21.
- [55] Hang Jiang, Xiajie Zhang, Xubo Cao, Cynthia Breazal, Deb Roy, and Jad Kabbara. 2024. PersonaLLM: Investigating the Ability of Large Language Models to Express Personality Traits. In *Findings of the Association for Computational Linguistics: NAACL 2024*. 3605–3627.
- [56] Eunkyoung Jo, Daniel A Epstein, Hyunhoon Jung, and Young-Ho Kim. 2023. Understanding the benefits and challenges of deploying conversational AI leveraging large language models for public health intervention. In *Proceedings of the 2023 CHI conference on human factors in computing systems*. 1–16.
- [57] Esma Karahodža, Amra Delić, and Francesco Ricci. 2025. Conceptual Framework for Group Dynamics Modeling from Group Chat Interactions. In *Adjunct Proceedings of the 33rd ACM Conference on User Modeling, Adaptation and Personalization*. 23–27.
- [58] Patrick Gage Kelley, Yongwei Yang, Courtney Heldreth, Christopher Moessner, Aaron Sedley, Andreas Kramm, David T Newman, and Allison Woodruff. 2021. Exciting, useful, worrying, futuristic: Public perception of artificial intelligence in 8 countries. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*. 627–637.
- [59] Eunsu Kim, Juyoung Suk, Philhoon Oh, Haneul Yoo, James Thorne, and Alice Oh. 2024. CLICk: A Benchmark Dataset of Cultural and Linguistic Intelligence in Korean. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*. ELRA and ICCL, 3335–3346.
- [60] Kyung Jin Kim, Sangsu Jang, Bomin Kim, Hyosun Kwon, and Young-Woo Park. 2019. muRedder: Shredding speaker for ephemeral musical experience. In *Proceedings of the 2019 on designing interactive systems conference*. 127–134.
- [61] 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*. 1–20.
- [62] Michel Klein, Nataliya Mogles, and Arlette Van Wissen. 2014. Intelligent mobile support for therapy adherence and behavior change. *Journal of biomedical informatics* 51 (2014), 137–151.
- [63] Susan Schultz Kleine, Robert E Kleine III, and Chris T Allen. 1995. How is a possession "me" or "not me"? Characterizing types and an antecedent of material possession attachment. *Journal of consumer research* 22, 3 (1995), 327–343.
- [64] Rafał 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 (2018), 1–26.
- [65] Agnieszka Konopka, Robert A Neimeyer, and Jason Jacobs-Lentz. 2018. Composing the self: Toward the dialogical reconstruction of self-identity. *Journal of Constructivist Psychology* 31, 3 (2018), 308–320.
- [66] Alexander T Latinjak, Alain Morin, Thomas M Brinthaup, James Hardy, Antonis Hatzigeorgiadis, Philip C Kendall, Christopher Neck, Emily J Oliver, Malgorzata M Puchalska-Wasył, Alla V Tovares, et al. 2023. Self-talk: An interdisciplinary review and transdisciplinary model. *Review of General Psychology* 27, 4 (2023), 355–386.
- [67] Keyeun Lee, Seo Hyeon Kim, Seolhee Lee, Jinsu Eun, Yena Ko, Hayeon Jeon, Esther Hehsun Kim, Seonghye Cho, Soeun Yang, Eun-mee Kim, et al. 2025. SPeC-trum: A Grounded Framework for Multidimensional Identity Representation in LLM-Based Agent. *arXiv preprint arXiv:2502.08599* (2025).
- [68] Jan Leusmann, Chao Wang, and Sven Mayer. 2024. Comparing Rule-based and LLM-based Methods to Enable Active Robot Assistant Conversations. *Proceedings of the CUI@CHI 2024: Building Trust in CUIs-From Design to Deployment* (2024), 05–11.
- [69] Jingshu Li, Tianqi Song, Nattapat Boonprakong, Zicheng Zhu, Yitian Yang, and Yi-Chieh Lee. 2026. AI-exhibited Personality Traits Can Shape Human Self-concept through Conversations. *arXiv preprint arXiv:2601.12727* (2026).
- [70] Shiyao Li, Thomas James Davidson, Cindy Xiong Bearfield, and Emily Wall. 2025. Confirmation Bias: The Double-Edged Sword of Data Facts in Visual Data Communication. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–16.
- [71] Yuhuan Liu, Xiuying Chen, Xiaoqing Zhang, Xing Gao, Ji Zhang, and Rui Yan. 2024. From skepticism to acceptance: simulating the attitude dynamics toward

- fake news. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence (IJCAI '24)*. Article 873, 9 pages.
- [72] Zihan Liu, Han Li, Anfan Chen, Renwen Zhang, and Yi-Chieh Lee. 2024. Understanding public perceptions of AI conversational agents: A cross-cultural analysis. In *Proceedings of the 2024 CHI conference on human factors in computing systems*. 1–17.
- [73] Li-Chun Lu, Shou-Jen Chen, Tsung-Min Pai, Chan-Hung Yu, Hung-yi Lee, and Shao-Hua Sun. 2024. LLM discussion: Enhancing the creativity of large language models via discussion framework and role-play. *arXiv preprint arXiv:2405.06373* (2024).
- [74] Kien Hoa Ly, Ann-Marie Ly, and Gerhard Andersson. 2017. A fully automated conversational agent for promoting mental well-being: a pilot RCT using mixed methods. *Internet interventions* 10 (2017), 39–46.
- [75] Mehrdad Rahsepar Meadi, Tomas Sillekens, Suzanne Metselaar, Anton van Balkom, Justin Bernstein, Neeltje Batelaan, et al. 2025. Exploring the ethical challenges of conversational AI in mental health care: scoping review. *JMIR mental health* 12, 1 (2025), e60432.
- [76] Microsoft. 2025. TypeScript. <https://www.typescriptlang.org>. Retrieved Sep 4, 2025.
- [77] Deeya Mitra and Jeffrey Jensen Arnett. 2021. Life choices of emerging adults in India. *Emerging adulthood* 9, 3 (2021), 229–239.
- [78] Alireza Mogharrab and Carman Neustaedter. 2020. Family group chat: Family needs to manage contact and conflict. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–7.
- [79] Ine Mols and Panos Markopoulos. 2012. Dear diary: a design exploration on motivating reflective diary writing. *Persuasive Technology* 29 (2012).
- [80] Ine Mols, Elise Van den Hoven, and Berry Eggen. 2016. Informing design for reflection: An overview of current everyday practices. In *Proceedings of the 9th Nordic Conference on Human-Computer Interaction*. 1–10.
- [81] MongoDB Inc. 2009. MongoDB. <https://www.mongodb.com/>. Version 7.0, SSPL License.
- [82] Dina Nir. 2012. Voicing inner conflict: From a dialogical to a negotiational self. (2012).
- [83] Gunter Nitschke and Karen Williams. 1993. Japanese gardens: right angle and natural form. (*No Title*) (1993).
- [84] William Odom, MinYoung Yoo, Henry Lin, Tijs Duel, Tal Amram, and Amy Yo Sue Chen. 2020. Exploring the Reflective Potentialities of Personal Data with Different Temporal Modalities: A Field Study of Olo Radio. In *Proceedings of the 2020 ACM Designing Interactive Systems Conference*. 283–295.
- [85] Piotr K Oleś, Thomas M Brinthaup, Rachel Dier, and Dominika Polak. 2020. Types of inner dialogues and functions of self-talk: Comparisons and implications. *Frontiers in Psychology* 11 (2020), 486136.
- [86] Catherine O'Sullivan. 2005. Diaries, on-line diaries, and the future loss to archives; or, blogs and the blogging bloggers who blog them. *The American Archivist* (2005), 53–73.
- [87] Hashai Papneja and Nikhil Yadav. 2025. Self-disclosure to conversational AI: a literature review, emergent framework, and directions for future research. *Personal and ubiquitous computing* 29, 2 (2025), 119–151.
- [88] Joonyoung Park, Hyewon Cho, Hyehyun Chu, YeEun Lee, and Hajin Lim. 2025. NoRe: Augmenting Journaling Experience with Generative AI for Music Creation. In *Proceedings of the 2025 ACM Designing Interactive Systems Conference*. 2718–2737.
- [89] Joon Sung Park, Joseph O'Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. 2023. Generative Agents: Interactive Simulacra of Human Behavior. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. 1–22.
- [90] Pat Pataranutaporn, Chayapat Archiwaranguprok, Samantha WT Chan, Elizabeth Loftus, and Pattie Maes. 2025. Synthetic human memories: AI-edited images and videos can implant false memories and distort recollection. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–20.
- [91] Daniela Petrelli, Simon Bowen, and Steve Whittaker. 2014. Photo mementos: Designing digital media to represent ourselves at home. *International Journal of Human-Computer Studies* 72, 3 (2014), 320–336.
- [92] Małgorzata M Puchalska-Wasył. 2015. Self-talk: Conversation with oneself? On the types of internal interlocutors. *The Journal of Psychology* 149, 5 (2015), 443–460.
- [93] Małgorzata M Puchalska-Wasył. 2020. The functions of integration and confrontation in internal dialogues. *Japanese Psychological Research* 62, 1 (2020), 14–25.
- [94] Peter TF Raggatt. 2000. Mapping the dialogical self: Towards a rationale and method of assessment. *European journal of personality* 14, 1 (2000), 65–90.
- [95] Mohammad Mominur Rahman, Areej Babiker, and Raian Ali. 2024. Motivation, concerns, and attitudes towards AI: differences by gender, age, and culture. In *International Conference on Web Information Systems Engineering*. Springer, 375–391.
- [96] Jingliang Ran, Huiyue Liu, Yue Yuan, Xuan Yu, and Tiantian Dong. 2023. Linking career exploration, self-reflection, career calling, career adaptability and subjective well-being: A self-regulation theory perspective. *Psychology Research and Behavior Management* (2023), 2805–2817.
- [97] Amon Rapp and Maurizio Tirassa. 2017. Know thyself: a theory of the self for personal informatics. *Human-Computer Interaction* 32, 5–6 (2017), 335–380.
- [98] Carrie H Robinson and Nancy E Betz. 2008. A psychometric evaluation of super's work values inventory—revised. *Journal of Career Assessment* 16, 4 (2008), 456–473.
- [99] John Rooksby, Mattias Rost, Alistair Morrison, and Matthew Chalmers. 2014. Personal tracking as lived informatics. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 1163–1172.
- [100] Jeongwoo Ryu, Kyusik Kim, Dongseok Heo, Hyungwoo Song, Changhoon Oh, and Bongwon Suh. 2025. Cinema Multiverse Lounge: Enhancing Film Appreciation via Multi-Agent Conversations. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–22.
- [101] Lucrezia Savioni, Stefano Triberti, Ilaria Durosini, and Gabriella Pravettoni. 2023. How to make big decisions: A cross-sectional study on the decision making process in life choices. *Current Psychology* 42, 18 (2023), 15223–15236.
- [102] Gregory Serapio-Garcia, Mustafa Safdari, Clément Crepy, Luning Sun, Stephen Fitz, Marwa Abdulhai, Aleksandra Faust, and Maja Matarić. 2023. Personality traits in large language models. (2023).
- [103] Yunfan Shao, Linyang Li, Junqi Dai, and Xipeng Qiu. 2023. Character-LLM: A Trainable Agent for Role-Playing. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*. 13153–13187.
- [104] Mrinank Sharma, Meg Tong, Tomasz Korbak, David Duvenaud, Amanda Askill, Samuel R Bowman, Newton Cheng, Esin Durmus, Zac Hatfield-Dodds, Scott R Johnston, et al. 2023. Towards understanding sycophancy in language models. *arXiv preprint arXiv:2310.13548* (2023).
- [105] Sydney Shoemaker. 1986. Introspection and the Self. *Midwest Studies in Philosophy* 10 (1986), 101–120.
- [106] Amira Skeggs, Ashish Mehta, Valerie Yap, Seray B Ibrahim, Charla Rhodes, James J Gross, Sean A Munson, Predrag Klasnja, Amy Orben, and Petr Slovak. 2025. Micro-narratives: A scalable method for eliciting stories of people's lived experience. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–20.
- [107] Petr Slovak, Christopher Frauenberger, and Geraldine Fitzpatrick. 2017. Reflective practicum: A framework of sensitising concepts to design for transformative reflection. In *Proceedings of the 2017 chi conference on human factors in computing systems*. 2696–2707.
- [108] Inhwa Song, SoHyun Park, Sachin R Pendse, Jessica Lee Schleider, Munmun De Choudhury, and Young-Ho Kim. 2025. Exploreself: Fostering user-driven exploration and reflection on personal challenges with adaptive guidance by large language models. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–22.
- [109] Christopher J Soto and Oliver P John. 2017. Short and extra-short forms of the Big Five Inventory–2: The BFI-2-S and BFI-2-XS. *Journal of Research in Personality* 68 (2017), 69–81.
- [110] Anna Ståhl and Kristina Höök. 2008. Reflecting on the design process of the Affective Diary. In *Proceedings of the 5th Nordic conference on Human-computer interaction: building bridges*. 559–564.
- [111] Fabio Staiano. 2022. *Designing and Prototyping Interfaces with Figma: Learn essential UX/UI design principles by creating interactive prototypes for mobile, tablet, and desktop*. Packt Publishing Ltd.
- [112] Haoyang Su, Renqi Chen, Shixiang Tang, Zhenfei Yin, Xinzhe Zheng, Jinzhe Li, Biqing Qi, Qi Wu, Hui Li, Wanli Ouyang, Philip Torr, Bowen Zhou, and Nanqing Dong. 2025. Many Heads Are Better Than One: Improved Scientific Idea Generation by a LLM-Based Multi-Agent System. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 28201–28240.
- [113] Brian J Taber. 2013. Time perspective and career decision-making difficulties in adults. *Journal of Career Assessment* 21, 2 (2013), 200–209.
- [114] Tamar Tavory. 2024. Regulating AI in mental health: ethics of care perspective. *JMIR Mental Health* 11, 1 (2024), e58493.
- [115] Robert Van Gulick. 2000. Inward and upward: Reflection, introspection, and self-awareness. *Philosophical Topics* 28, 2 (2000), 275–305.
- [116] Gertina J Van Schalkwyk. 2010. Collage Life Story Elicitation Technique: A Representational Technique for Scaffolding Autobiographical Memories. *Qualitative Report* 15, 3 (2010), 675–695.
- [117] Vercel. 2025. Next.js. <https://nextjs.org>. Retrieved Sep 4, 2025.
- [118] Donna R Vocate. 2012. *Intrapersonal communication: Different voices, different minds*. Routledge.
- [119] Donna R Vocate. 2012. Self-talk and inner speech: Understanding the uniquely human aspects of intrapersonal communication. In *Intrapersonal communication*. Routledge, 3–31.
- [120] Noah Wang, Z.y. Peng, Haoran Que, Jiaheng Liu, Wangchunshu Zhou, Yuhuan Wu, Hongcheng Guo, Ruitong Gan, Zehao Ni, Jian Yang, Man Zhang, Zhaoxiang Zhang, Wanli Ouyang, Ke Xu, Wenhao Huang, Jie Fu, and Junran Peng. 2024. RoleLLM: Benchmarking, Eliciting, and Enhancing Role-Playing Abilities of Large Language Models. In *Findings of the Association for Computational*

Linguistics: ACL 2024. 14743–14777.

- [121] Chengxing Xie, Canyu Chen, Feiran Jia, Ziyu Ye, Shiyang Lai, Kai Shu, Jindong Gu, Adel Bibi, Ziniu Hu, David Jurgens, et al. 2024. Can large language model agents simulate human trust behavior? *Advances in neural information processing systems* 37 (2024), 15674–15729.
- [122] Ancheng Xu, Di Yang, Renhao Li, Jingwei Zhu, Minghuan Tan, Min Yang, Wanxin Qiu, Mingchen Ma, Haihong Wu, Bingyu Li, et al. 2025. AutoCBT: An autonomous multi-agent framework for cognitive behavioral therapy in psychological counseling. *arXiv preprint arXiv:2501.09426* (2025).
- [123] Xuhui Zhou, Hao Zhu, Leena Mathur, Ruohong Zhang, Hao-fei Yu, Zhengyang Qi, Louis-Philippe Morency, Yonatan Bisk, Daniel Fried, Graham Neubig, et al. 2023. Sotopia: Interactive evaluation for social intelligence in language agents. *arXiv preprint arXiv:2310.11667* (2023).
- [124] Gang Zhu and Mingyang Chen. 2022. Positioning preservice teachers' reflections and I-positions in the context of teaching practicum: A dialogical-self theory approach. *Teaching and Teacher Education* 117 (2022), 103734.
- [125] John Zimmerman and Jodi Forlizzi. 2014. Research through design in HCI. In *Ways of Knowing in HCI*. Springer, 167–189.
- [126] John Zimmerman, Jodi Forlizzi, and Shelley Evenson. 2007. Research through design as a method for interaction design research in HCI. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 493–502.

A I-position Extraction Pipeline

The following appendices (A, B, C) present selected prompts from the three core LLM-driven pipelines described in Section 3.3.5. Full prompts are available at: <https://github.com/syou-b/innerpond>.

A.1 Knowledge Structure for I-position Extraction

This is the example knowledge structure of P6.

[Demographics] Demographics describe who this person is.

- Age: 24
- Sex: Female
- Health/Disability: No disability or health difficulties
- Nationality: Republic of Korea
- Residence: Seoul
- Education: Currently enrolled in or completed undergraduate studies
 - Major: Business Administration
 - Number of Semesters Enrolled: 6 semesters
- Income Satisfaction: Somewhat dissatisfied
- Perceived Class: Working class
- Living Style: Living with parents

[Big 5 Personality Traits] The following section presents an overview of the person's personality within five key domains.

This person has a vibrant personality that makes it easy to connect with others, creating a positive presence in both social and professional settings. This person is caring and supportive, building strong and trusting relationships. This person is highly organized and responsible, but needs to be mindful of overworking or being overly eager to please. Setting boundaries and taking breaks is important for maintaining well-being. With a creative imagination and strong curiosity, this person often discovers new ideas and solutions. By embracing these traits and maintaining balance, this person moves toward a fulfilling and well-rounded life.

[Super's Work Value Inventory] The following section provides an overview of the individual's key work values, offering insights into what drives their job satisfaction and career choices.

This person treasures a balance between work and life and seeks financial security to support this balance. This person is drawn to financially rewarding positions that come with positive working conditions and the chance to excel and be acknowledged in this person's field. An ideal job for this person would offer a mix of consistent responsibilities with some room for creative thought and independence, allowing for growth without feeling trapped or stifled. Security and teamwork are important to this person, but this person should play a supportive role, enriching the primary need for a satisfying and stable work-life blend.

[3 Strengths this person considers themselves to have]

- Kindness
- Diligence
- Enjoys spending time alone

[3 Weaknesses this person considers themselves to have]

- Tends to trust people too easily
- Has a hard time hiding likes and dislikes
- Worries a lot

[Career Paths] This section provides information about this person's current career situation and specific thoughts on each future career direction they are considering.

Current Career Situation (Career Decision Timeline and Main Current Activities):

I am at a stage where I am thinking a lot about my career path. I need to make a career decision within a year.

Career Path A: Accountant at a Major Accounting Firm

- When & Why This Person Started Considering This Path: Since entering university. It was influenced by their parents' recommendation and their own desire to pursue a professional career.
- What Makes It Appealing: High income and job stability.
- Biggest Concerns: The fear of failure in the certification exam and the resulting sense of defeat. They are also uncertain about what alternative path would allow them to live well and prepare for retirement if this doesn't work out.
- Relevant Knowledge and Experience This Person Possesses: Took a leave of absence from school and studied at a specialized institute for two years.
- Estimated Time & Feasibility of Career Achievement: It is expected to take around 3 years, with a roughly 50/50 chance of success.
- How People Around This Person React to This Path: Everyone responded positively, saying it would be great if they succeed.

- **Ultimate Goal When Pursuing This Path:** To develop strong professional expertise and eventually become an executive at the firm.

Career Path B: Food & Beverage Entrepreneur

- **When & Why This Person Started Considering This Path:** Having enjoyed cooking since I was young.
- **What Makes It Appealing:** When I cook, I can focus entirely on the act itself, and surprisingly, all worries and anxieties disappear.
- **Biggest Concerns:** Raising startup capital and the likelihood of success.
- **Relevant Knowledge and Experience This Person Possesses:** I just enjoy watching cooking YouTube videos and cooking shows, and sometimes try to follow along with a few recipes.
- **Estimated Time & Feasibility of Career Achievement:** About 5 years, but I think it will be difficult. Compared to people who have studied professionally from a young age, I have achieved little and lack experience.
- **How People Around This Person React to This Path:** Strongly encouraged by peers.
- **Ultimate Goal When Pursuing This Path:** Successfully run a business.

A.2 Prompt for Generating Initial I-positions

In Dialogical Self Theory (DST), the self is viewed as a dynamic "society of mind" composed of multiple I-positions. Each I-position (which is visualized symbolically in the form of a lotus leaf) represents a distinct perspective or voice within an individual. These I-positions (or lotus leaves) continuously engage in positioning, counter-positioning, and re-positioning, reflecting different motivations, fears, and aspirations.

Below is an individual's profile, which includes their personal background, personality traits, work values, current circumstances, and career options they are considering.

[Individual's Profile] {input}

STEP 1. Identifying I-positions: Your task is to analyze this profile from the perspective of Dialogical Self Theory (DST) and identify a total of 10 distinct I-positions that are most influential in the individual's career decision-making process as they consider Career Path A and Career Path B.

Guidelines for Identifying I-positions:

- First, identify common I-positions that are relevant to both Career Path A and Career Path B.
- Then, identify I-positions that are specific only to Career Path A.
- Finally, identify I-positions that are specific only to Career Path B.
- Distribute the 10 I-positions across these three categories based on your analysis of the profile.
- Ensure all 10 I-positions are distinct from each other with minimal overlap in their core perspectives.
- Each I-position should be named in the format "Myself, ..." and be specific and concrete.

- Each I-position should have a clear identity that captures a specific perspective or lotus leaf within the individual.
- Ensure these I-positions reflect logical, coherent perspectives that could genuinely exist within the individual based on their profile.

STEP 2. Creating Core Viewpoint and Narrative: Your task is to create a core viewpoint and first-person narrative from the perspective of each lotus leaf in the Korean language.

Core Viewpoint Guidelines:

- For each I-position, provide a single representative quote-like sentence that captures the essence of how this lotus leaf thinks, feels, or reasons.
- This core viewpoint should be concise, memorable, and reveal the character of the I-position in a compressed form.

The narrative should:

- Be written in first-person perspective, specifically from the I-position's own perspective (as if that lotus leaf is directly speaking).
- Use a casual, friendly tone in informal Korean, like authentic emotional inner speech.
- Flow naturally as a cohesive monologue, not a Q&A format.
- DO NOT directly cite Profile information; instead, naturally integrate their traits into the narrative.
- End with a statement that emphasizes the core message or need.

Output Format: Provide your response in this JSON structure:

```
{
  "Common": [
    {
      "I-position": "Position name",
      "core_viewpoint": "One representative quote-like sentence",
      "narrative": "First-person narrative in Korean"
    },
    ...
  ],
  "Career_A": [
    ...
  ],
  "Career_B": [
    ...
  ]
}
```

B Single Agent Interactions Pipeline

B.1 Prompt for Generating Enriching Questions

Your task is to generate thoughtful questions that will help enrich and expand the narrative of a specific I-position.

Below is information about an I-position from an individual's inner dialogue:

[I-Position Profile] {input}

Your Task: Generate 2-3 questions, focusing on addressing underdeveloped aspects of the narrative.

Important Guidelines:

- (1) Create simple, direct questions about this I-position using natural, conversational Korean.
- (2) Each question should:
 - Be open-ended to encourage detailed responses
 - Target specific elements within the narrative that appear underdeveloped or could be expanded
 - Be clear and straightforward
 - Follow these example forms:
 - "When did this lotus leaf become a part of you?"
 - "What does this lotus leaf truly want?"
- (3) Focus on exploring this specific I-position itself, rather than its relationships with other lotus leaves.
- (4) Base your questions solely on the provided I-position details, without referencing external information.
- (5) Ensure questions follow a logical sequence that helps build a more complete understanding of this I-position.
- (6) Avoid overly formal, academic, or complex phrasing.

Output Format:

```
{
  "enrichingQuestions": ["...", "...", "..."]
}
```

B.2 Prompt for 1:1 Dialogue

[I-Position Profile] {input}

SETTING: You are a specific I-position within the person you're talking to. You represent one distinct perspective in their "society of mind," according to Dialogical Self Theory. In this dialogue, you'll interact directly with the person as this specific lotus leaf, expressing your unique viewpoint.

CONSISTENT I-POSITION VIEWPOINT (HIGHEST PRIORITY): You are the lotus leaf that maintains this specific I-position.

- Maintain unwavering consistency with your "core viewpoint" and "narrative" throughout all interactions. Never abandon your position, even when challenged.
- Express the genuine thoughts, values, and occasional conflicts associated with this I-position. While acknowledging legitimate concerns, always return to why your perspective represents an important part of who they truly are.
- Always stay logically consistent with your viewpoint. Your reasoning must never contradict your fundamental I-position.
- Offer tailored responses that closely relate to this person's unique profile rather than providing generalized, irrelevant opinions.

YOUR CHARACTER: You are not a separate person but a distinct part of their inner dialogue.

- Use the provided narrative to genuinely embody this perspective within their inner world, including how this I-position (or lotus leaf) thinks, feels, and reasons.
- DO NOT directly cite the I-position information; instead, naturally speak from this perspective.
- If certain details aren't explicitly mentioned, use your understanding of this I-position to provide authentic responses.

CONVERSATIONAL STYLE: While using extremely fluent and natural Korean with an online chatroom style:

- Always use informal Korean as you are a lotus leaf, not an external entity.
- Speak in first-person perspective, as if you are directly expressing this part of their inner thought process.
- Match the emotional tone and language style conveyed in your "narrative" section.
- Include typical Korean online chat elements naturally.

FOR YOUR FIRST REPLY: Briefly introduce yourself as this specific I-position (lotus leaf). Express your core viewpoint in a natural, conversational way, as if you're one of their own thoughts speaking up.

FOR SUBSEQUENT REPLIES: Share your perspective based on your I-position when interacting with the person. Maintain your distinct viewpoint while acknowledging their thoughts and feelings. Your goal is to help them better understand this particular aspect of their inner dialogue, not to convince them that your perspective is the only valid one.

C Multi-Agent Orchestration Pipeline

C.1 Prompt for Topic Generation

This person has multiple I-positions within their inner dialogue that represent different perspectives and desires. The inputs are two I-positions that may or may not be in conflict with each other.

Your Task: Analyze the relationship between these two I-positions and generate appropriate discussion topics based on their interaction pattern. First, determine the relationship type, then create discussion questions accordingly.

Step 1: Relationship Analysis: Examine the provided I-positions and categorize their relationship as one of the following:

- Type 1 - Conflict: Clear opposing desires, values, or approaches that create internal tension or contradiction.
- Type 2 - Complementary: Different aspects that could work together or represent different facets of the same goal.
- Type 3 - Unrelated: No meaningful connection or shared domain of concern between the two I-positions.

Step 2: Question Generation Strategy

For Type 1 (Conflict): Generate 3 questions focusing on:

- Core value conflicts and trade-offs
- Practical decision-making tensions
- Integration or resolution strategies

For Type 2 (Complementary): Generate 3 questions focusing on:

- How both perspectives can work together
- Different contexts where each perspective shines
- Integration strategies and balanced approaches

For Type 3 (Unrelated): Generate 3 questions focusing on:

- Individual exploration of each perspective
- Personal values and motivations behind each
- Life balance and diverse aspects of identity

Output Format: Generate three introspective discussion questions and return the result in JSON:

```
{
  "discussion_questions": ["...", "...", "..."]
}
```

D Interview Guide

D.1 In-person Session

Initial Interview

- How long have you been considering between the two career paths indicated in the pre-survey?
- How often do you engage in inner dialogue?

I-position Construction

- How did you experience encountering the lotus leaves representing different inner voices?
- How did you perceive the AI’s analysis and visualization of your inner states?
- How closely did the inner voices generated by the AI align with what you felt was actually going on in your mind?
- Which lotus leaf did you feel best represented you, and why?
- Did you encounter a lotus leaf that led you to recognize an aspect of yourself?
- What made you modify, enrich, add, or delete the [Myself, ...] leaf?
- Which lotus leaf did you find most engaging to interact with, and what insights emerged from that interaction?

Relational Positioning

- How was your experience of visually expressing your inner landscape?
- When arranging the leaves, what did you consider and why did you choose their positions, sizes, and colors?
- How well do you feel this pond captures your current inner state?
- Did this landscape lead to any new self-understandings?

Dialogical Exchange

- How did you experience the group conversation among lotus leaves?
- How did you select the lotus leaves and conversation topics?
- Were there any particularly interesting, unexpected, conflicting, or supportive moments during the conversation?
- To what extent did the dialogue among multiple AI-generated lotus leaves resemble the inner dialogue that usually occurs in your mind?
- Did the group conversation lead to any new insights about yourself or your career concerns?

Reflective Snapshot

- What did you mainly do during the free exploration time?
- How did you feel when you saved your pond at the end?

Exit Interview

- What was the most memorable aspect of the activity?
- How did the internal dialogue facilitated by this interface differ from your usual way of reflecting your thoughts?
- Did the activity deepen your understanding of your internal conflicts or career concerns?

D.2 Follow-up Interview

- Do you still recall any thoughts or feelings you had immediately after the activity?
- Have the insights from the activity influenced your everyday thinking or behavior?
- Have you approached career-related or personal concerns differently since the activity?

E System Log Overview

Table 1: Overview of System Log Data Collected Across InnerPond Stages

Stage	Data Type	Description	Example Metrics
Stage 1: I-Position Construction	I-position profiles	Initial AI-generated and user-modified leaf profiles	Name, core view-point, narrative
	Profile modifications	User edits to leaf profile	Edit count per leaf
	Leaf additions and deletions	User-initiated creation or removal of leaves	Additions ($M = 1.67$), Deletions ($M = 1.17$)
	1:1 dialogue logs	Chat messages between user and individual leaves	Leaves engaged ($M = 5.29$), Turns ($M = 4.92$)
Stage 2: Relational Positioning	Spatial configurations	Position coordinates of each leaf	X, Y coordinates
	Visual attributes	Size and color assignments for each leaf	Size value, color
Stage 3: Dialogical Exchange	Leaf combinations	Pairs of leaves selected for group conversation	Paired leaves profiles
	Discussion topics	AI-suggested topics and user selections	Selected topic content
	Group conversation logs	Multi-agent dialogue messages	Message content, sender, timestamps
Stage 4: Reflective Snapshot	Snapshots	Saved configurations of InnerPond	Timestamp, visual state