CS 147, FALL 2025 – EMPOWERING LEARNERS WITH AI

FINAL REPORT

theo

Gentle nudges. Real progress.

ANANYA N. | AYANA G. | ANTHONY C. | FELIX Z.

Contents

Project Name & Value Proposition

PROJECT NAME

Theo

VALUE PROPOSITION

Gentle nudges. Real progress.

Team Members & Roles

ANANYA N.
Designer/
Developer

AYANA G.
Designer/
Developer

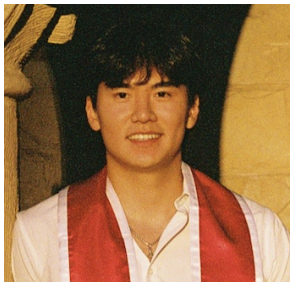

ANTHONY C.
Designer/
Developer

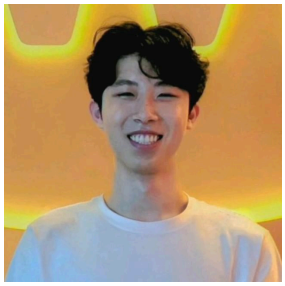

FELIX Z.
Designer/
Developer

Problem & Solution Overview

Many learners with ADHD (Attention Deficit/Hyperactivity Disorder) struggle with a range of focus-related challenges regarding their learning habits, including hyperfocus, hypofocus, approaching large goals, and finding optimal work environments. **Theo** works to support these learners by providing an artificial source of *body-doubling* (the practice of having another individual nearby for passive support) through the form of a virtual study companion. By offering an AI-supported organizational tool to aid in goal breakdown as well as a reflection chatbot for users to share their thoughts with while they work, **Theo** endeavors to create anchor points for learners to direct their focus towards and strive for personal success and growth.

Needfinding

INTERVIEWS

In our domain ideation stage prior to interviews, our team attempted to brainstorm learning communities that are largely underserved in our communities. Among our top focus groups were ESL learners, senior citizen learners, and neurodivergent students (our final pick). Finding the third option to be the most accessible, both in terms of participants and anticipated level of tech literacy, we chose to focus our attention on neurodivergent learners, specifically those with ADHD.

Our interviewing process consisted of meetings with seven individuals, split between two groups: (1) learners with ADHD of varying ages, and (2) field experts. Four interviewees (three of which were high school students and the fourth was a special education director) were based in San Francisco and three were based in and around the Palo Alto/Stanford area (a Stanford RF and early career graphic designer—both of whom had ADHD—and a director at the Stanford CTL).

With all interviews, we aimed to uncover the role of technology in achieving educational goals of any magnitude for ADHD learners, and how it could be improved. Some questions we asked our interviewees with ADHD included:

- Walk us through a regular weekday – when do you usually work/study?
- Describe your ideal/most productive work environment.
- Does technology play a role in your daily life? If yes, can you elaborate on how and where you use it?
- Do you use any tools to improve your productivity and efficiency?
- What is the hardest part of a task or assignment?
- How would you describe your work habits?
- What are things that bother you? What excites or encourages you?
- If you encountered someone who didn't know what the term ADHD meant, how would you explain it to them?

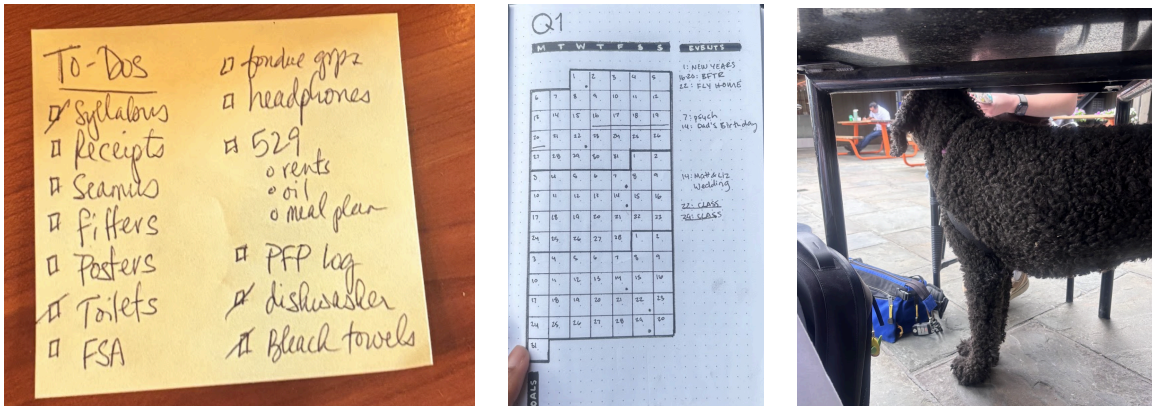

Figure 1. Artifacts collected during interviews. (From left to right) (1) Geoff B.'s daily to-do checklist, (2) Kayla M.'s Q1 planner, and (3) Kayla M.'s service dog, Percy.

For later review, we recorded our interviews with the help of the *HeyMarvin* transcription tool and video recordings of our interviewees (with their prior consent, of course), along with Google Docs for written notes.

SYNTHESIS

To synthesize our findings, we created empathy maps with *Miro* that laid out each interviewee's key verbal responses, actions, thoughts, and feelings. These maps helped us visualize and create connections between similar aspects both within a single interview and across multiple perspectives as well. The main trends we pulled from our interviews are as follows:

- Hyperfocus → Spending an excessive amount of time on smaller details increases overall frustration.
- External reminders and cues aid in keeping track of necessary tasks. This could come in many forms (digital reminders, body doubling, pets).
- Active engagement and relatability makes work feel more accessible and goals seem achievable.

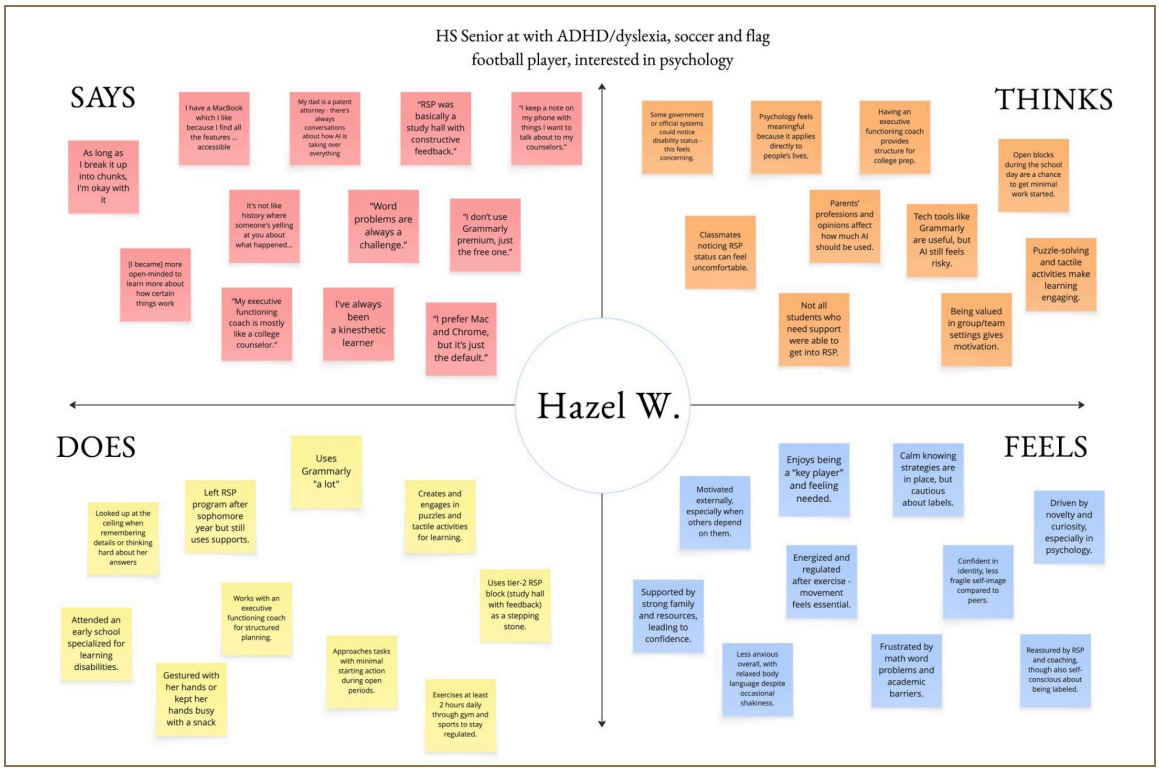

Figure 2. Example empathy map for Hazel W.: a high school senior with ADHD and dyslexia, competitive soccer and flag football player, interested in pursuing psychology in college.

We also learned about some security and social concerns, specifically those relating to being identified or officially documented as neurodivergent, that would give us some values to consider when developing our future prototypes and final product.

POVs & Experience Prototypes

FINAL POV STATEMENTS

We created three POV statements from three of our seven interviewees:

1. GEOFF B.
    a. **We met:** a 56-year-old husband with ADHD who spends most of his workday on his laptop. We were surprised to learn that even though he recognizes when he's hyperfocused on trivial work tasks, he can't redirect himself without his wife stepping in.
    b. **We were surprised to learn that:** even though he recognizes when he's hyperfocused on trivial work tasks, he can't redirect himself without his wife stepping in.
    c. **We wonder if this means:** the repeated experience of losing control leaves him feeling frustrated or dependent, undermining his confidence in managing his own attention.
    d. **It would be game-changing to:** help him reclaim control in the moment to recognize when his attention drifts and consciously realign with what matters most.

2. KAYLA M.
    a. **We met:** A freelance graphic designer with ADHD who mainly interacts with her husband and her dog most days.
    b. **We were surprised to learn that:** she uses her service dog's needs to trigger her own self-care (eating, going to the bathroom etc.).
    c. **We wonder if this means:** she struggles to balance her need for external stimuli/structure with her need to feel independent.
    d. **It would be game-changing to:** create supportive rhythms that help her feel anchored and cared for while still giving her agency, even in solitude.

3. CELINE P.
    a. **We met:** a CMU student with ADHD and anxiety who studies best and experiences significantly less burnout when other people are around.
    b. **We were surprised to learn that:** she focuses better when there's a small group of people nearby having an unrelated conversation than when she's watching or listening to a TV show. Quiet feels too empty, and TV feels too scripted.
    c. **We wonder if this means:** she craves the subtle comfort of human presence; ambient signs of life that make her feel safe, grounded, and less alone with her thoughts.
    d. **It would be game-changing to:** offer ways for her to tap into that shared energy of "working alongside," wherever she is.

We used these POVs to identify areas of improvement and target specific aspects of each individual's experience to alter or enhance.

SAMPLE HMW STATEMENTS

After completing our POVs, we utilized these to formulate How-Might-We statements that could channel our design thinking towards more specific needs and goals. Here are some examples of our collected statements and respective POVs:

1. Geoff POV → HMWs:
    a. HMW *give* him tools to ***self-intervene*** without disrupting his flow?
    b. HMW *support* him in ***aligning attention*** with his goals *in real time*?
    c. HMW *allow* him to design his own ***"intervention triggers"***?
2. Kayla POV → HMWs:
    a. HMW *offer* ***body-doubling*** when she's alone?
    b. HMW *place nudges* where she commonly lacks the ***initiation energy***?
    c. HMW *make* breaks feel short, ***valid***, and finished?

3. Celine POV → HMWs:
    a. HMW *provide* low-effort access to ***social presence*** during independent work?
    b. HMW *embed* or *simulate* external cues that ***prompt refocusing*** or ***breaking hyperfocus***?
    c. HMW *regulate* her ***work rhythms*** through something similar to human presence?

From the above HMWs, we extracted clear values and ideation points that would become the core of our future work, and narrowed down our list to the following three statements, which became our springboards for solution brainstorming:

1. HMW *enable* individuals to ***connect what they're doing*** in the moment to their larger goals?
2. HMW *replicate* the ***accountability*** of an external reminder through ***self-generated systems***?
3. HMW *recreate* the benefits of ***background conversation*** and/or ***body doubling***?

TOP 3 SOLUTIONS

From the above top three HMW statements, we brainstormed potential solutions in short 5-10 minute sprints. After evaluating all ideas, we decided to pursue the following three unique options as our experience prototypes:

1. **Check-in moments for reflection and reminders for breaks:** An AI study-buddy that aids in setting the structure of a productive work session: identifying goals and steps to achieve the goal, determining the time spent on the task(s) at hand, including and enforcing periodic breaks for resetting and re-energizing
2. **AI soundscapes for focus:** Curating sound environments to reduce stress using grounding techniques while working

3. **AI animal companion for body-doubling:** An animal character that uses body-doubling techniques to encourage users to take breaks and take care of biological needs

To test the projected effectiveness of each solution, we next developed experience prototypes mimicking their respective behaviors and UX.

EXPERIENCE PROTOTYPES

For each experience prototype, we established an assumption to test, derived from our needfinding and POVs. We then found participants in our focus audience who would interact with these prototypes and give us their initial feedback, which would help inform our decision regarding which solution(s) to proceed with.

1. **PROTOTYPE #1: Check-in moments for reflection & break reminders**
    a. **Critical Assumption:** Users struggle with ensuring that they are satisfied with the quality of a work session. (Hyperfocus, distracted easily, or lack of attention to personal/natural needs)
    b. **Execution:** One team member simulated an AI chatbot that timed a 1-hr work session, during which the participant set a goal and worked towards it remotely (to simulate the lack of physical human involvement). At intervals, the team member would send a text to the participant, reminding them to get a drink of water, eat a snack, or reflect on the work they had completed so far.
    c. **Implications:** The participant enjoyed the motivation to stay focused, and was pleasantly surprised by the post-work reflection. However, they would have preferred a stronger sense of personality and visuals to supplement the chatbot's lack of physical presence. Our next step would be to create some notion of personality, catering directly to the user's own persona to seem more like a friend than a generic AI assistant.

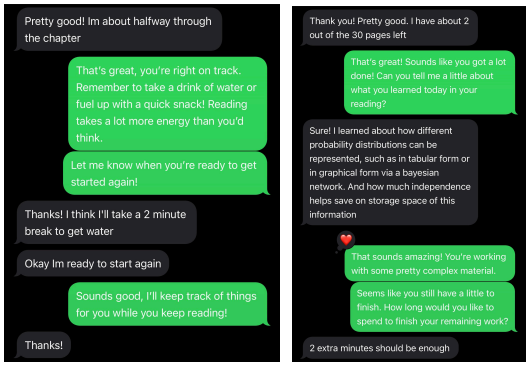

Figure 3. Experience prototype #1. Text conversation simulating an AI chatbot.

2. **PROTOTYPE #2: AI soundscapes for focus**

    a. **Critical Assumption:** Users will feel comfortable relinquishing control over their sound environment to something or someone else.

    b. **Execution:** The participant was given a short Google form to complete, detailing their listening preferences. Following this, the tester from our team selected a video with audio to play in the background while the participant worked for 30 min on a Zoom call. Once the session was completed, the participant shared their feedback regarding the ease of the experience.

    c. **Implications:** While the participant was initially focused and engaged (they enjoyed the chaotic, conversational energy and preferred dynamic noise over silence), they eventually got distracted by the dialogue in the video instead of their task(s), repeating phrases and mimicking sounds. This would pose the challenge of curating sound that could tread the fine line between white noise and distracting noise.

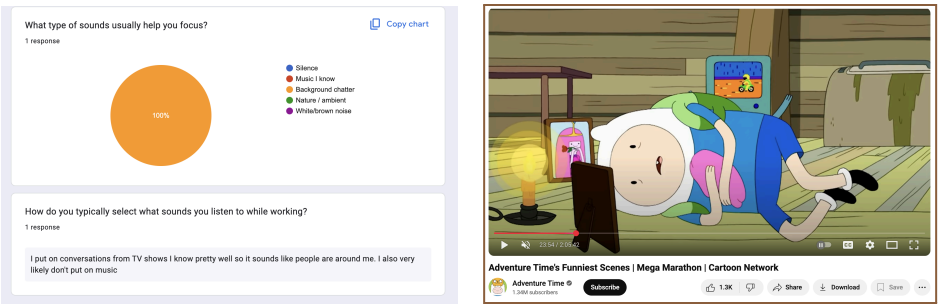

Figure 4. Experience prototype #2. (From left to right) (1) Google form requesting listening preferences, (2) Selected video for user to listen to based on results.

3. **PROTOTYPE #3: AI animal companion for body-doubling**
    a. **Critical Assumption:** Individuals with ADHD prefer to take cues from others to remind themselves to complete simple but necessary tasks (snacking, hydrating, going to the bathroom) when hyperfocused.
    b. **Execution:** The participant was observed and assessed in three different working situations: (1) alone, (2) while with another person, and (3) with the AI companion. In the third scenario, a team member texted AI-generated images of an animal fulfilling biological needs (getting food, going to the bathroom, etc.). At the end of the experiment, the participant reported their thoughts about each experience.
    c. **Implications:** The participant did find the AI companion engaging; however, they were ultimately distracted by the AI generated images and lost focus. This left questions for how to provide more relatable AI animal companions that provide just enough cues for the user but not so much novelty to require the user's full attention.

Figure 5. Experience prototype #3. Google Gemini-generated image of a sloth drinking water. The sloth was chosen based on the participant's preferred animal.

Design Evolution

FINAL SOLUTION

We decided to combine elements from Solution #1 (**Check-in moments for reflection and reminders for breaks**) and Solution #3 (**AI animal companion for body-doubling**).

Based on the findings from the experience prototypes, we found that both check-ins and an animal companion provided promise, and these two can be woven together to create one cohesive experience where the animal companion is scaffolding and guiding those check-in moments. In addition, going back to our needfinding interviews and POVs, we felt that the combination of these solutions best addressed the pain points of our users: reminders for self care, ways to connect what they are doing in a given moment to their broader goals, and ways to feel a sense of presence even in solitude work.

We ultimately decided to not incorporate elements of Solution #2 (**AI soundscapes for focus**) as we felt this relied a lot on sound design and leveraging current sound-based generative AI systems, which is an emerging technology that is not currently robust. We wanted to focus our efforts on user interaction elements, not sound design.

TASKS

SIMPLE TASK: Start a basic work session

Why? All of our target users are focused on doing work sessions, although in different contexts (work, school, personal planning, etc.). So, our simple task–the one that all users would want to do–is simply beginning a work session with a certain goal in mind.

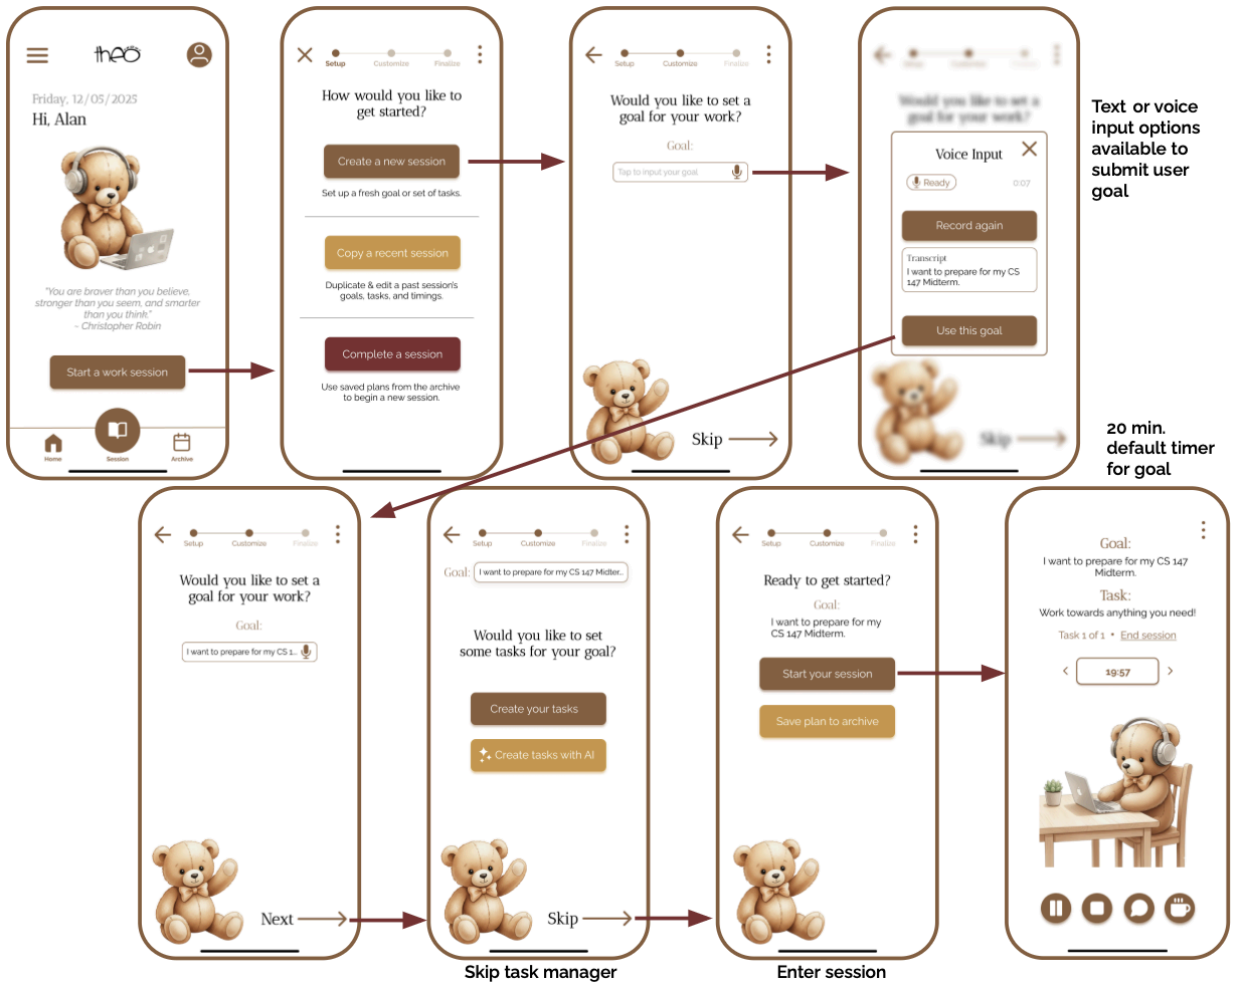

Figure 6. Simple task flow shown on our final interface.

MODERATE TASK: Reflect during or after a work session

Why? From our needfinding, we found that our interviewees often only reflected on what they were doing if this reflection was triggered by an outside source, such as a friend or partner. From our experience prototype, we found that our participant responded well to reflection prompts and reported feeling motivated and engaged. This takes more effort on the user's end, as they need to input their thoughts.

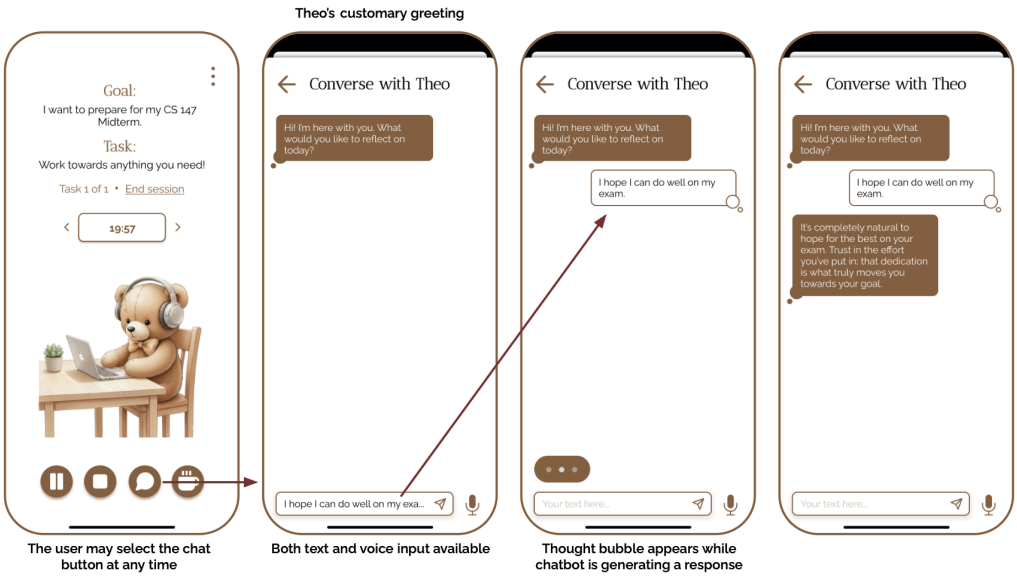

Figure 7. Moderate task flow shown on our final interface.

COMPLEX TASK: Break down a larger goal into smaller tasks

Why? One of the key pain points of our target user base is task initiation (getting started on something that they need to do). A key strategy for simply getting started is to make a larger goal less daunting by breaking it down into smaller, more manageable pieces. This is more complex as it involves the user manually inputting tasks, or using our AI-feature to do it for them, but this still requires them to edit and review the output.

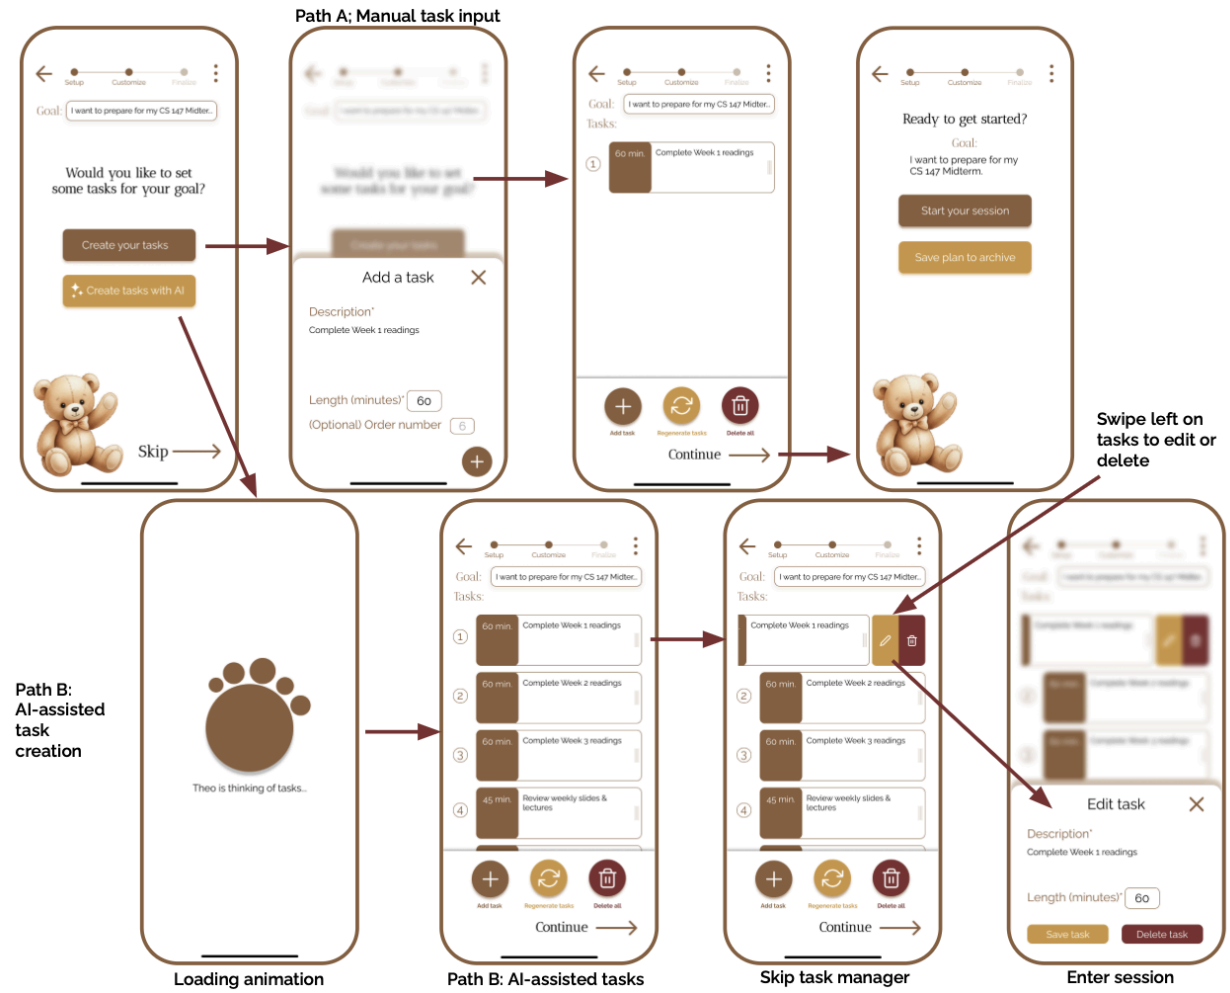

Figure 8. Complex task flow shown on our final interface. This image outlines two simple paths for (A) manual goal breakdown and (B) AI-assisted goal breakdown.

DESIGN EVOLUTION VISUALIZATION & RATIONALE

INITIAL SKETCHES

During the initial sketching and brainstorming phases, we tried to get as creative as possible and focus on quantity. Through sketching exercises, we came up with a breadth of ideas across various realizations, including a physical teddy bear, a speech assistant, and a Tamagatchi pet. After discussing and narrowing down our ideas, we went deeper into two realizations.

Web application

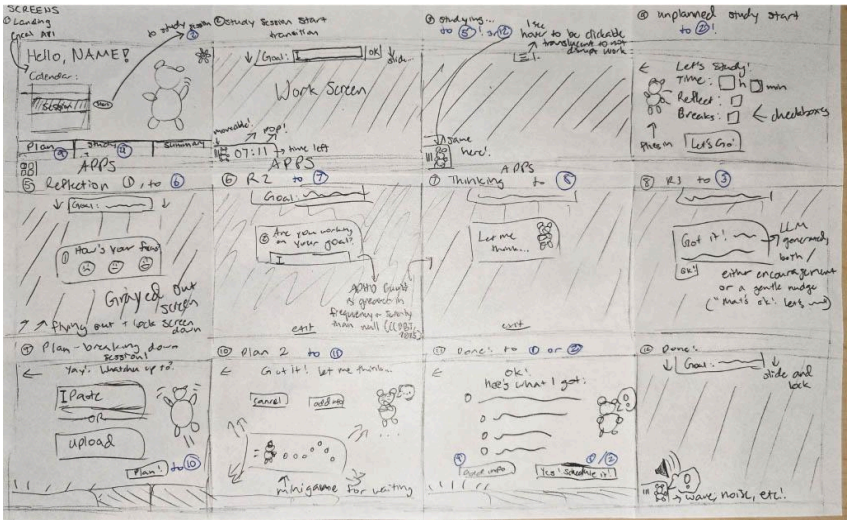

Figure 9. Initial sketches of our web app realization.

We liked the idea of a web application because the interface is more consolidated with a typical work platform and having a larger interface leaves room for more features and more long form responses. However, it becomes less convenient for all-time use, due to its reduced portability, restricted access to audio input, and limited integration capabilities with other common work apps, and overall potentially lowered effectiveness to remove users from zones of hyperfocus.

Mobile application

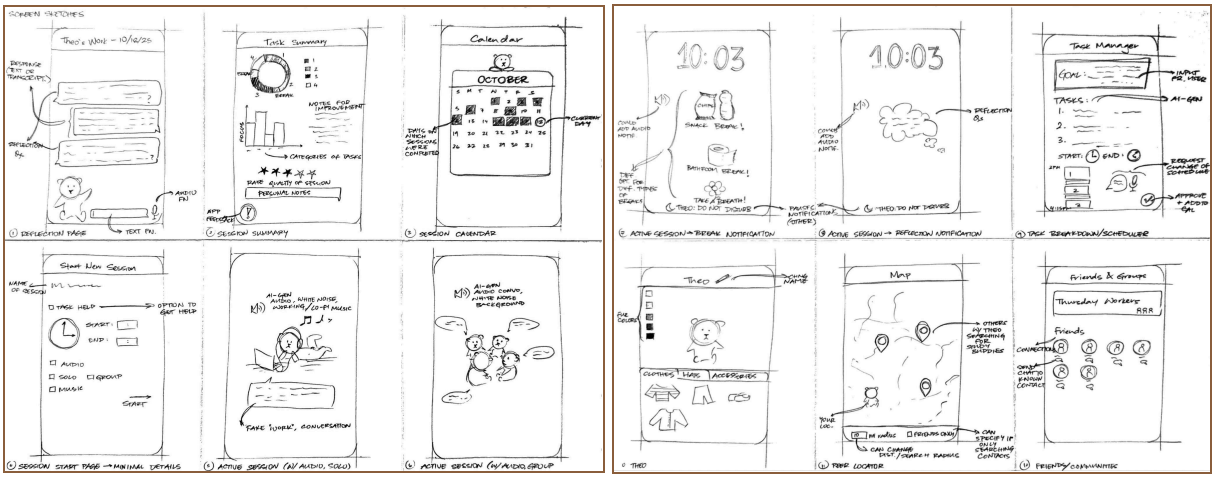

Figure 10. Initial sketches of our mobile app realization.

We were drawn to a mobile application for its greater accessibility, portability, and more universal usability. It becomes easier to utilize voice and camera features built into a mobile device and notifications/reminders are more natural and noticeable. On the other hand, a much smaller interface inherently impedes more detailed interaction/reflection by nature of needing heightened finger dexterity. Additionally, switching between a phone and regular work device (typically a desktop) may be frustrating and could lead to further distraction – a problem our solution aims to curb.

Ultimately, while a web application has many benefits and lends itself better to specific features that would be effective to include, **the mobile application is a better embodiment of our values, physically and metaphorically.**

We hoped to produce a product that offered a more passive than active presence, feels organic to use, and can be utilized as a companion. To implement the desired effects, we found strengths in the **portability, size, simplicity,** and **access** to the features that a mobile device offers, specifically voice control, emulating conversations with friends and family.

LOW-FIDELITY PROTOTYPE

With our realization decided, we began with a low-fi paper prototype, made both on paper as well as online in Excalidraw and then printed for user testing. This aimed to help us get a basic understanding of how users interacted with our task flows.

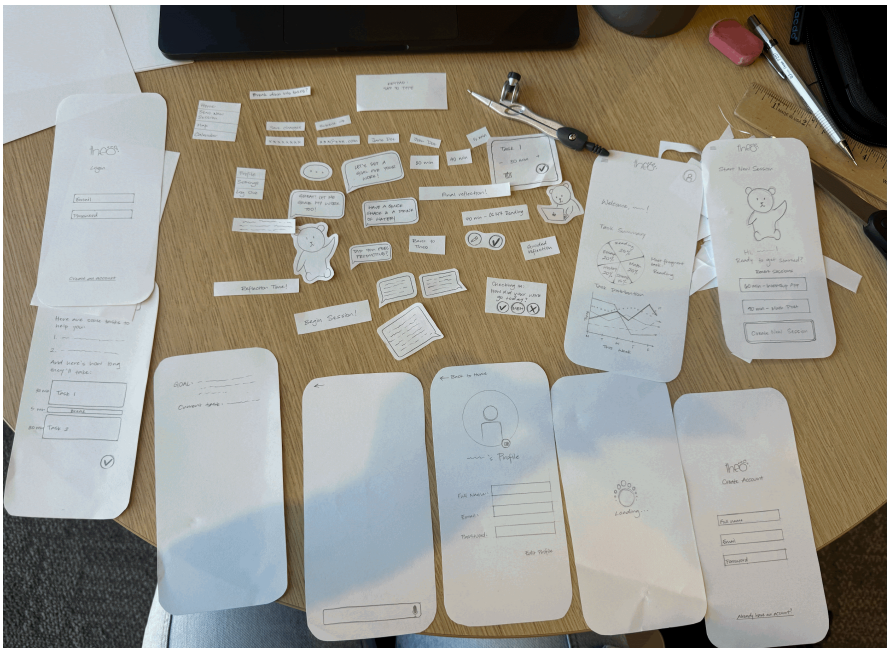

Figure 11. Our initial handmade paper prototype.

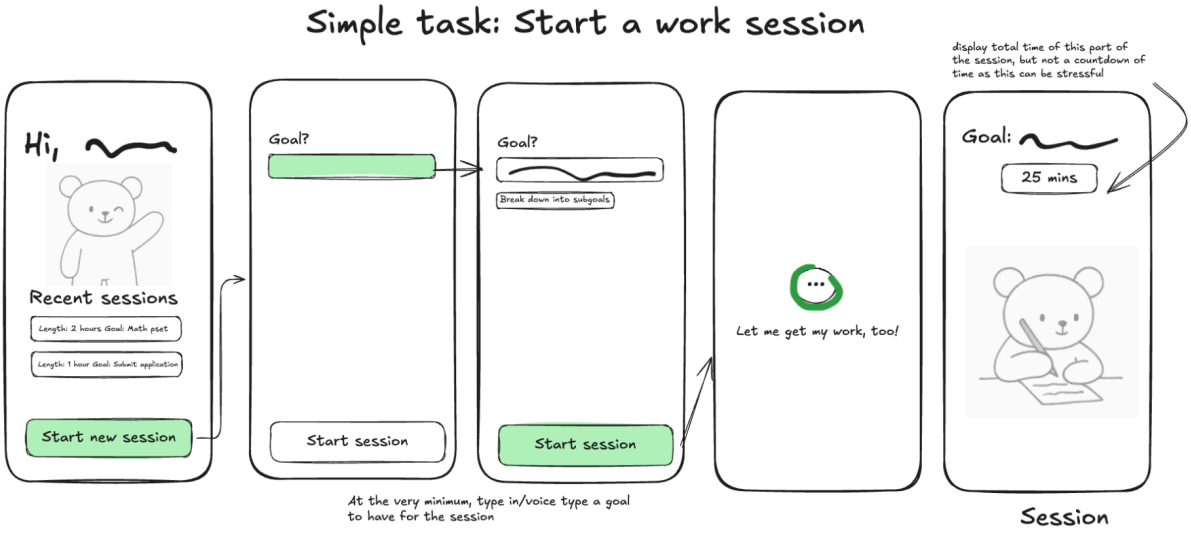

Figure 12. Simple task flow in our digital low-fi prototype.

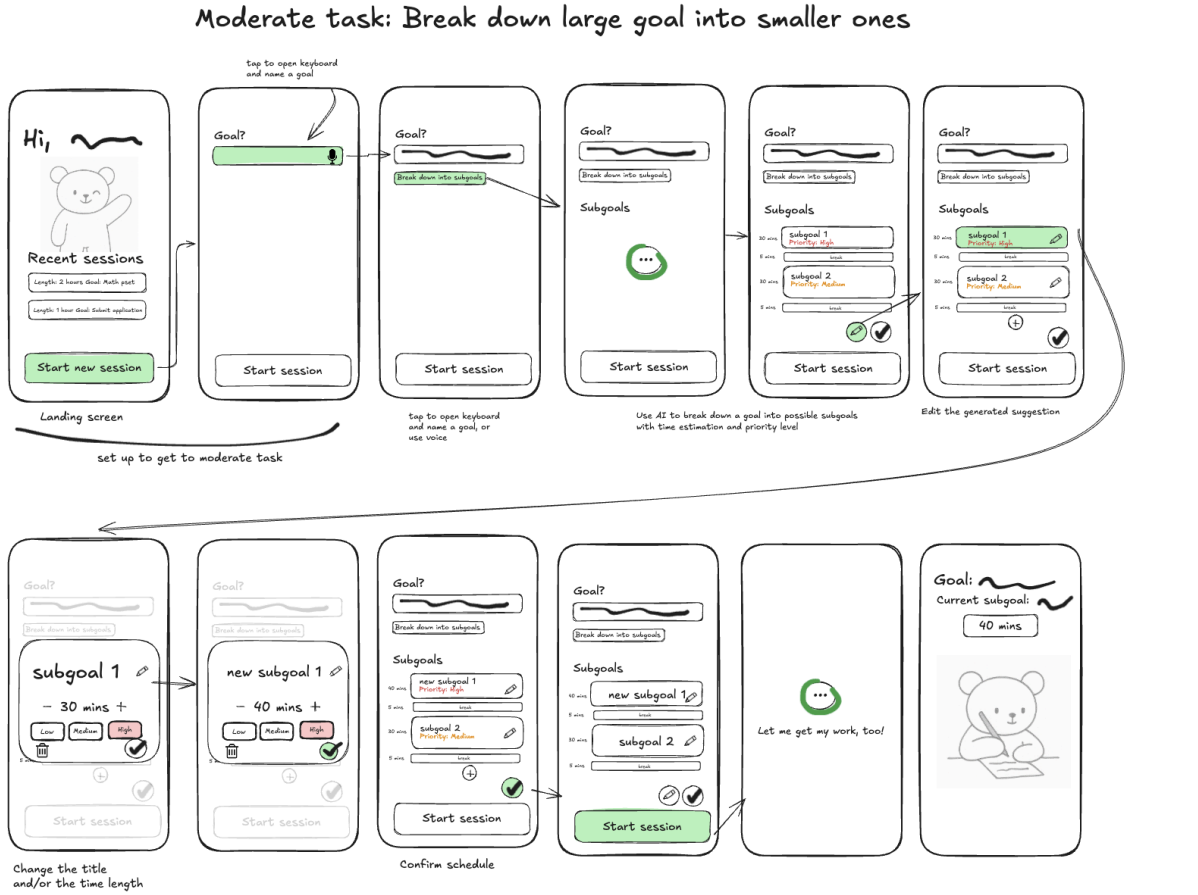

Figure 13. Moderate task flow in our digital low-fi prototype.

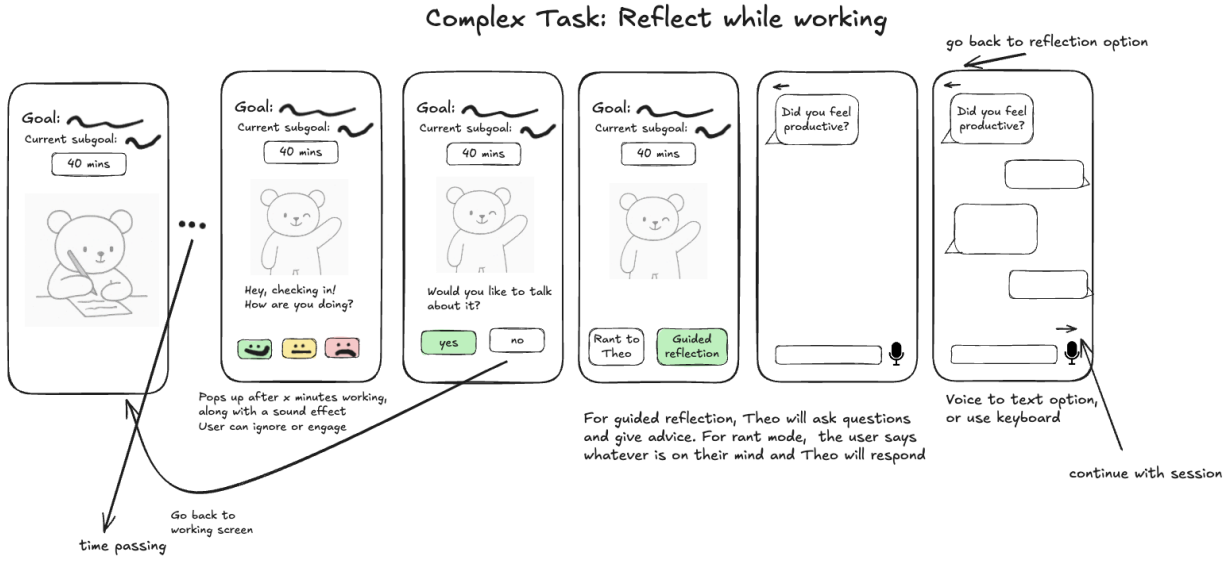

Figure 14. Complex task flow in our digital low-fi prototype.

Evaluation Technique: For this iteration of testing, we printed our prototype, walked participants through how to interact with the prototype, and asked them to complete each of our three tasks. A facilitator introduced and described the prototype. The participant then tapped on the paper screens, and the "computer" placed the next appropriate screen/modal. The observer recorded sessions for notes and review.

Figure 15. The printed version of the prototype used for testing, featuring modals, a keyboard, and a voice input button for users to interact with.

As our assessment metrics, we recorded bottom-line data: the number of taps a participant utilized for each task, the time taken to complete each task, and the number of incorrect taps. We also collected process data (verbal feedback) to get a better understanding of points of confusion as well as points of enjoyment.

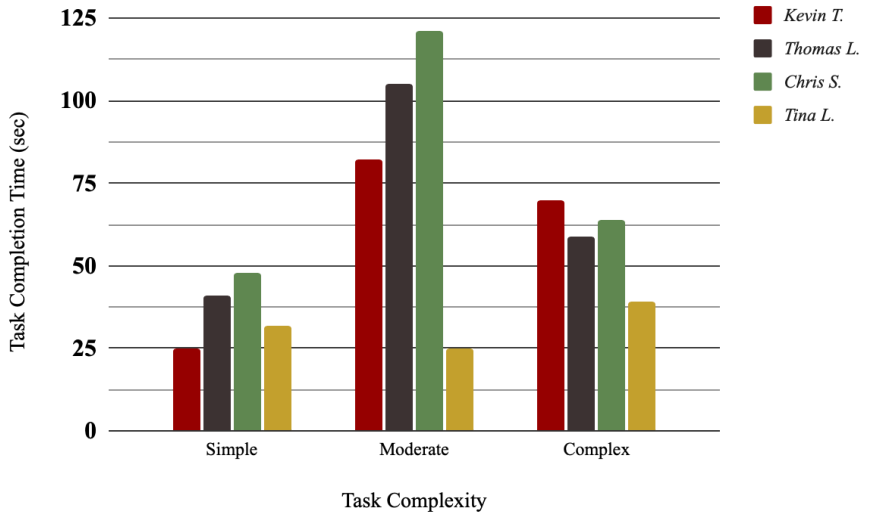

Figure 16: A graph representing our bottom-line data (time taken to complete each task for all participants), showing the need to modify our moderate task.

Learnings and Implications: Through analyzing both process data and bottom-line data, we found the body-doubling aspect through a teddy bear to be a key hook for all participants! Some users enjoyed the prospect of having AI-generated suggestions for larger goals, while others felt restricted by the generated tasks and wanted more customization power. Excessive time and focus investment was needed from the participant to first ensure task breakdown was up to par before beginning a work session. The current interface proved overwhelming and required additional explanations and clarifications for most participants.

Our findings revealed several focus areas for improvement:

- **Emphasize companionship** further to draw users in and provide motivation *(as evidenced by the enjoyment participants reported by having the teddy bear companion; this was our strong suit!)*
- **Reconfigure UI for goal breakdown**: simplify while still providing maximum customizability *(as evidenced by the number of taps and time taken to complete this task, and verbal confusion about this flow and its level of customization)*
- **Ensure reflection reminders are subtle but effective,** not distracting, more supportive, and customizable *(as evidenced by the limitation of being able to test this via this method; concerns raised by evaluators brought this to our attention)*

MEDIUM-FIDELITY PROTOTYPE

Using the learnings from the low-fi, we created a medium-fidelity prototype on Figma, which focused further on stylistic elements as we introduced colors, fonts, and our teddy bear mascot, as well as switched the previous moderate and complex tasks to better reflect user feedback. Instead of relying on a human computer, this prototype was operable on its own, with built-in screen changing functionality and backend variables tracking user preferences. This prototype underwent two iterations after receiving expert feedback.

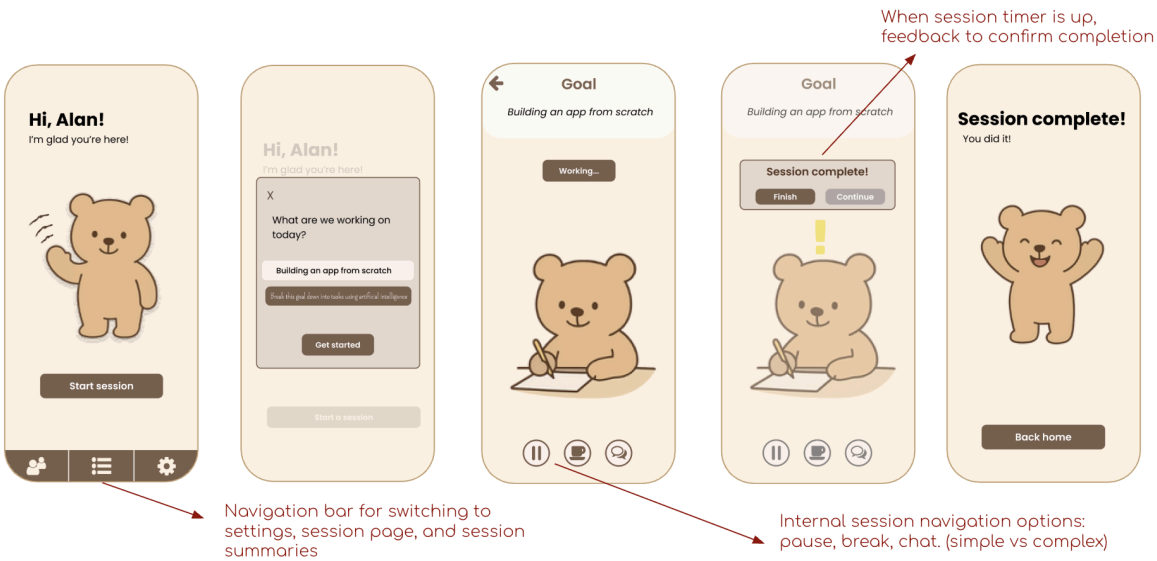

Figure 17. Our simple task flow in our med-fi prototype (v1)

Figure 18. Our adjusted moderate task flow in our med-fi prototype (v1)

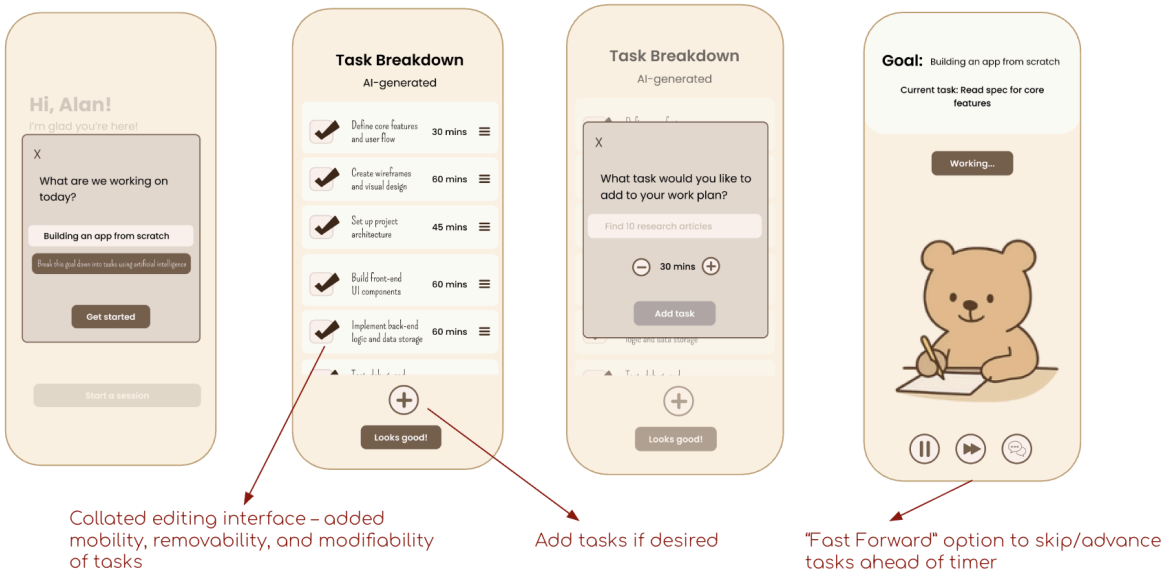

Figure 19. Our adjusted complex task flow in our med-fi prototype (v1)

After this initial implementation, we wanted to further enhance our overall theming (such as the color palette and style of the teddy bear). This second version explored a more fleshed-out style while expanding the customizability of our complex task.

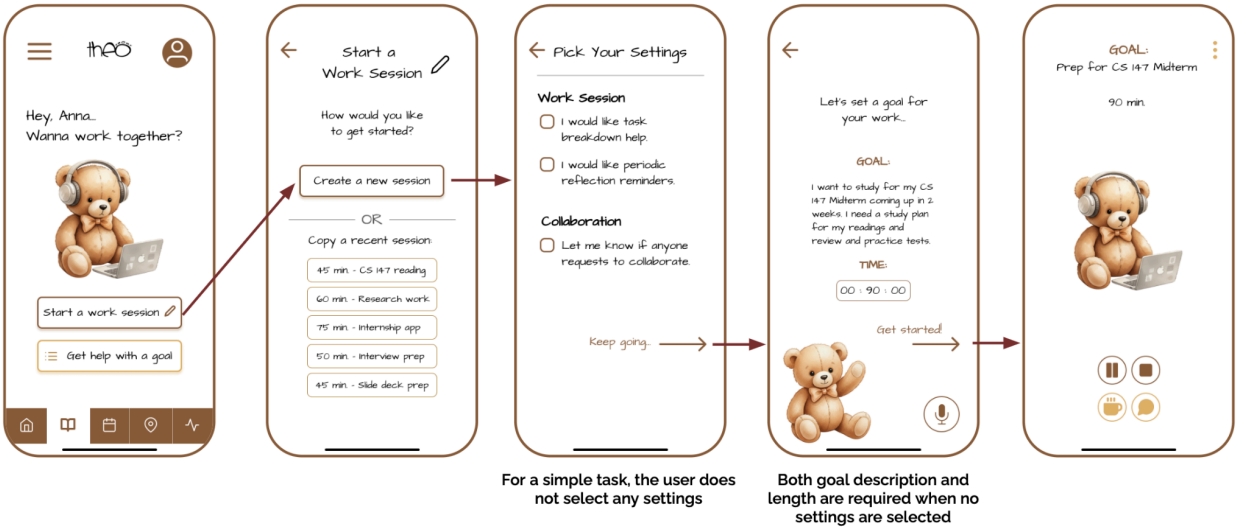

Figure 20. Our simple task flow in our med-fi prototype (v2)

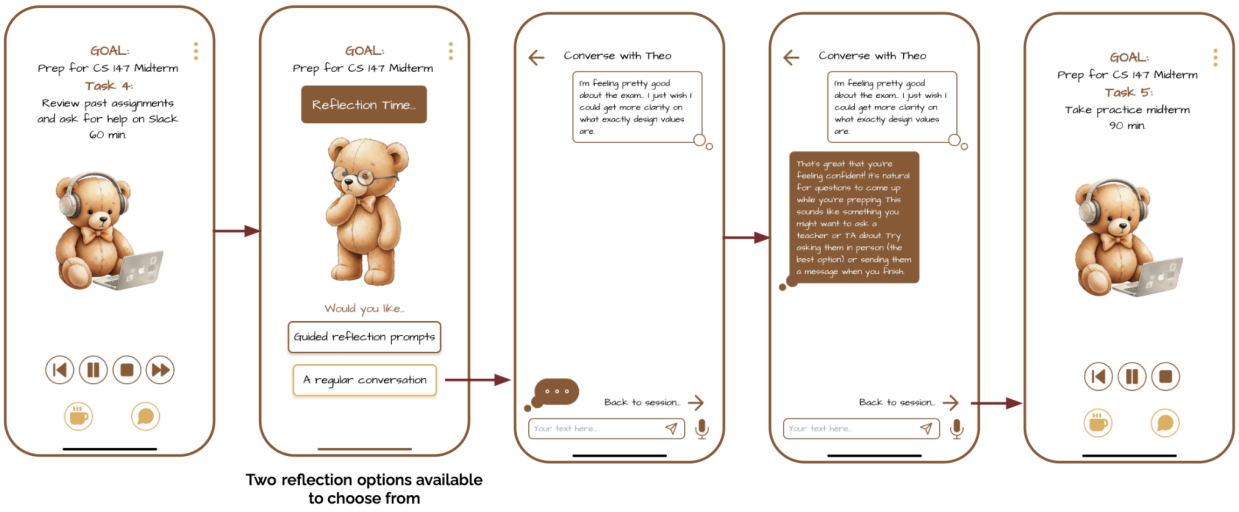

Figure 21. Our moderate task flow in our med-fi prototype (v2)

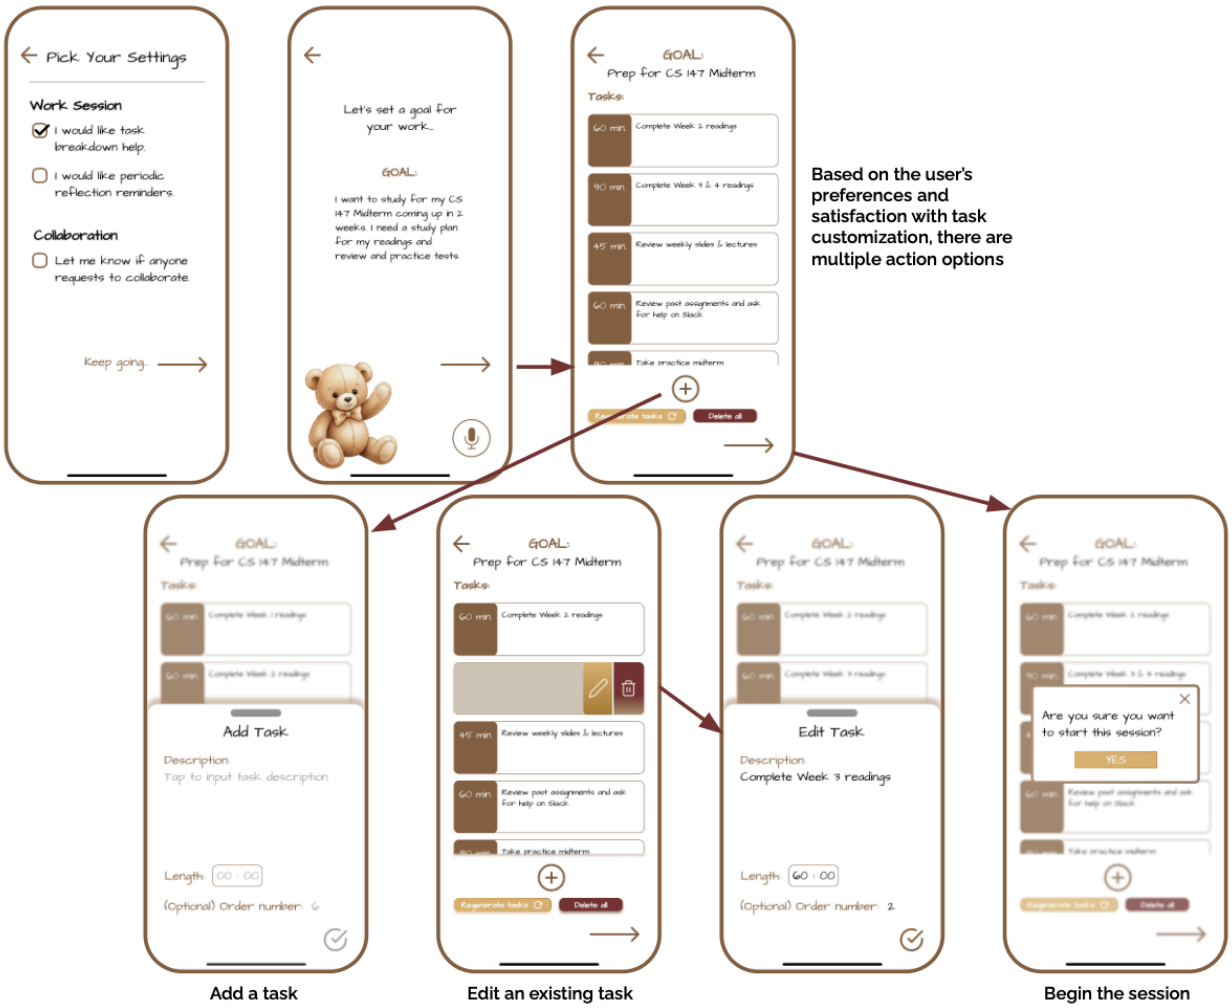

Figure 22. Our complex task flow in our med-fi prototype (v2)

Evaluation Technique: We relied on a heuristic evaluation compiled by our peers, following guidelines established by Nielsen's heuristics (further enhanced by two additional CS147 heuristics) to identify aspects of our prototype that violated standards and could improve user experience.

Learnings and Implications: Based on the results of the evaluation, we determined points of confusion, ambiguity, and inconsistency, as well as received recommendations for potential improvements. Out of an initial 11 violations with a severity degree of 3, we opted to focus on 4 of these that were directly relevant to our primary three task flows (the remaining 7 violations were either tied to our stretch workflows or prototype-based flaws rather than conceptual oversights):

1. **Issue #1:** There was a lack of confirmation after a user selected "Delete all tasks" in the Task Manager. As a result, any accidental selection would result in irreversible damage. *(H5: Error Prevention)*
   Fix #1: A confirmation modal was added as a buffer between the action selection and followthrough to give the user a chance to confirm their decision before deleting all tasks. This technique was later implemented throughout the app to provide safeguards against all large impact actions.

Figure 23. Updated workflow including a confirmation modal

2. **Issue #2:** Although the draw of the goal breakdown feature would be its AI integration, there was no provision for initial manual task input.
   (H3: User Control and Freedom)
   Fix #2: We revised the flow of the goal breakdown feature, providing the user with the option to generate tasks with the assistance of AI, manually input their own tasks, or skip the task creation process entirely. This created a greater sense of flexibility in the customization.

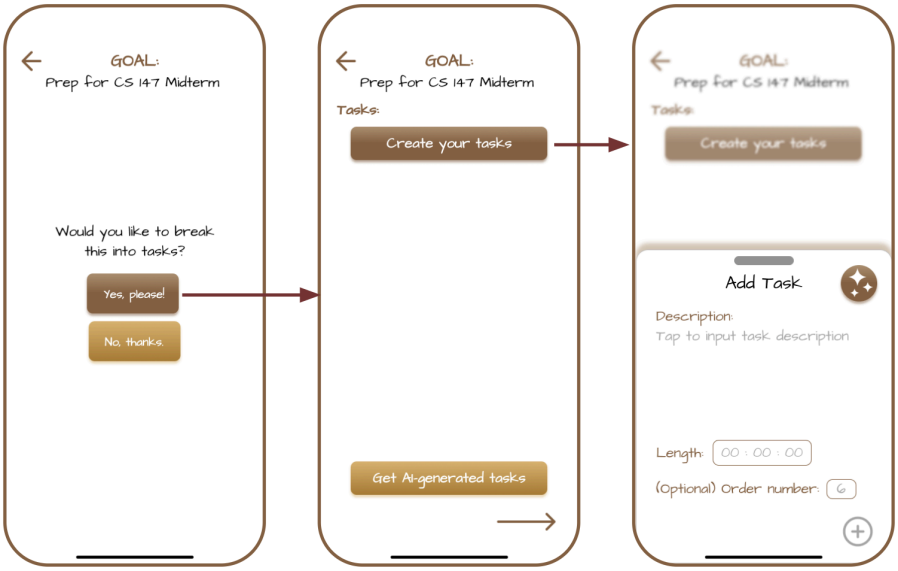

Figure 24. Updated workflow that offers users additional branches of decision-making.

3. **Issue #3:** During a session, users did not have access to the timer running behind the scenes, tracking the amount of time spent on a given task.
   (H1: Visibility of System Status)
   Fix #3: We learned from our needfinding that having access to a real-time countdown timer vs. an overall allotted time label could affect users in different ways: some found it motivating, while others found it intimidating. To cater to as many perspectives as possible, we adjusted the time display to offer 2 alternative display versions to give each user more control: (1) time remaining as a countdown, (2) overall allotted time as a banner, or (3) a neutral word to detach entirely from time. Users could tap the display component to toggle between the available options.

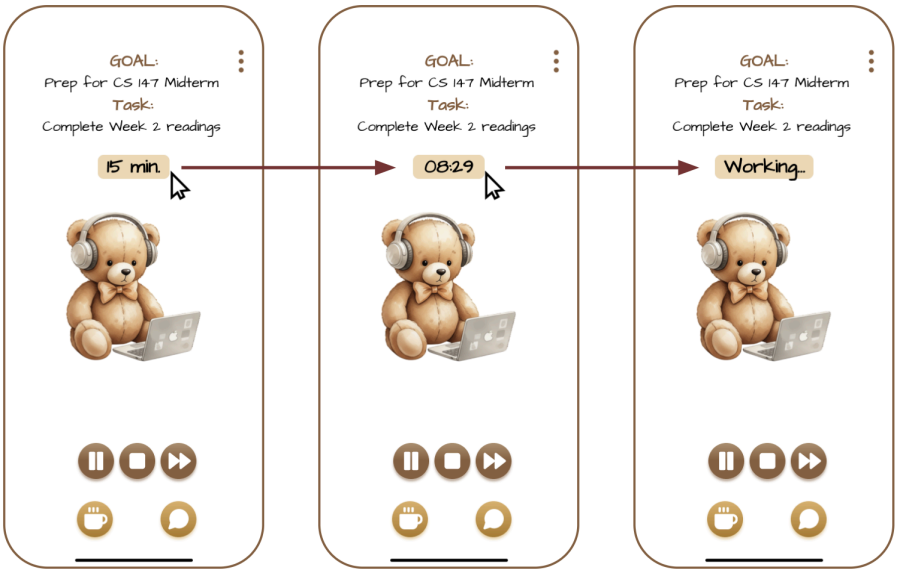

Figure 25. Updated timer component that toggles between three display versions on tap.

4. **Issue #4:** The visual depiction of pressables across the app lacked consistency in fill, stroke color, and shape. *(H4: Consistency & Standards)*
   Fix #4: We standardized our pressable components to increase salience and clearly indicate clickability. We created a block fill and color palette hierarchy (primary/default: **brown**, secondary/AI: **yellow**, tertiary/large-impact: **red**).

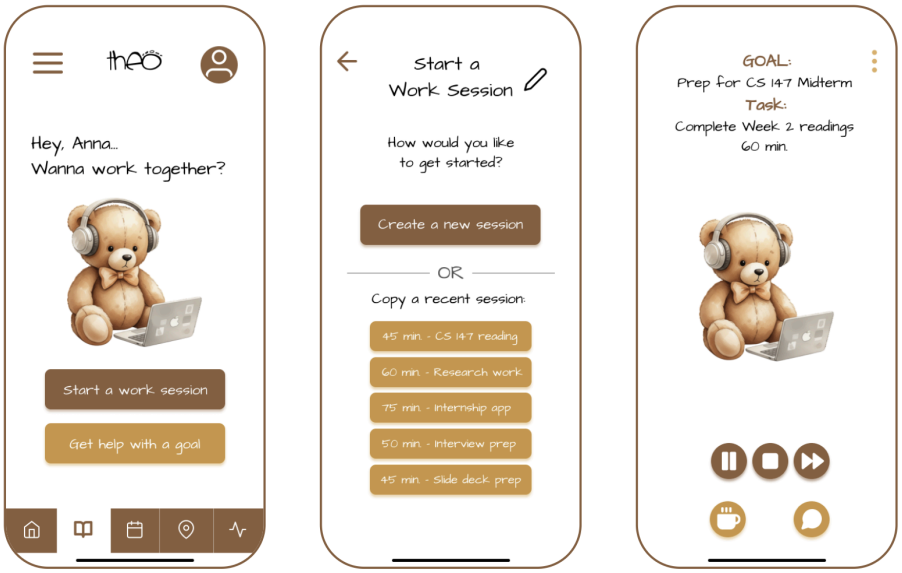

Figure 26. Updated styling with consistent clickables.

Beyond these changes, other key revisions we incorporated in our final high-fidelity prototype included:

- Goal breakdown styling & flow updates
- A simplified home screen with our central, featured action
- A larger variety of in-session flexibility

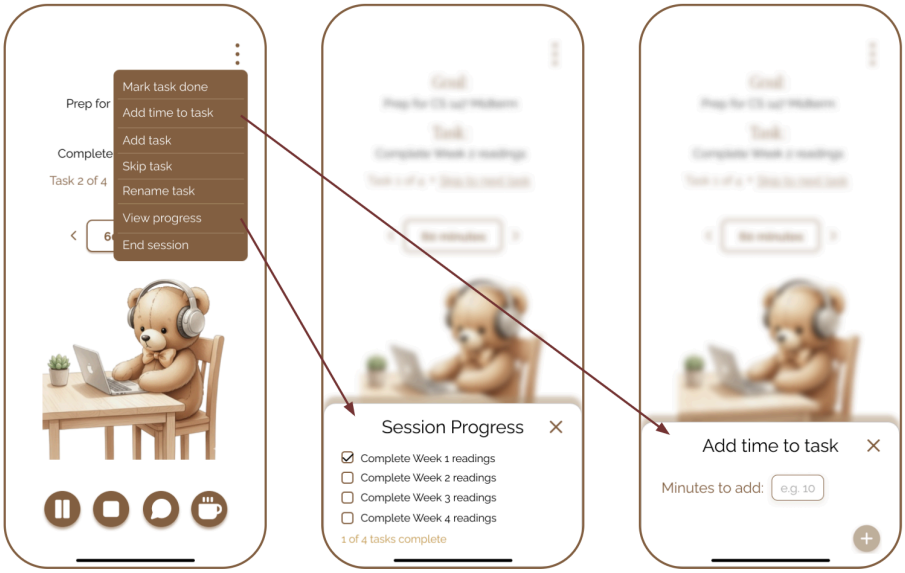

Figure 27. Updated in-session experience: (1) A menu of actions to navigate a session, including: (2) viewing a session's progress and (3) increasing the time allocated to the current task.

- An animation of Theo on screen during a session to further emphasize the body-doubling experience

VALUES IN DESIGN

Our main values centered around the accessibility of all features to as broad a spectrum of users as possible, extending a sense of empathy to all users, and providing maximum flexibility to personalize each user's experience.

ACCESSIBILITY

The more discernible an element, feature, or system is, the more likely users will be able to utilize it effectively.

Our designs take advantage of whitespace, Gestalt principles, and Fitts' Law by displaying minimal elements on any given screen, grouping clickable items, and minimizing the distance between navigational elements for ease.

All clickable components in our design system are visually coded through color: **brown**, **yellow**, or **red** fill with white text/strokes. We additionally revised our use of icons based on feedback to ensure their meaning is consistent with expectations. For example, a fast-forward icon that was originally intended to skip tasks in a session, was replaced by an explicitly-labeled "Skip to the next task" link and positioned near the task label for proximity.
The addition of a progress indicator when completing the session set up flow provides concrete visual cues regarding the current stage of progress.

Tradeoffs: A monotonous palette could reduce visual engagement or overwhelm a user who is more sensitive to color if a single screen contains multiple clickable components. Conversely, an excess of color would require the user to actively remember the significance of each tone, defeating our goal of simplicity. To address this, we maintained the previously mentioned color hierarchy to differentiate between clickables with distinct results and implications.

Furthermore, positioning similar actions adjacent to each other (with respect to the device screen size) risked an increased error margin in selecting an unintended action. To counteract this, all action clickable dimensions were increased.

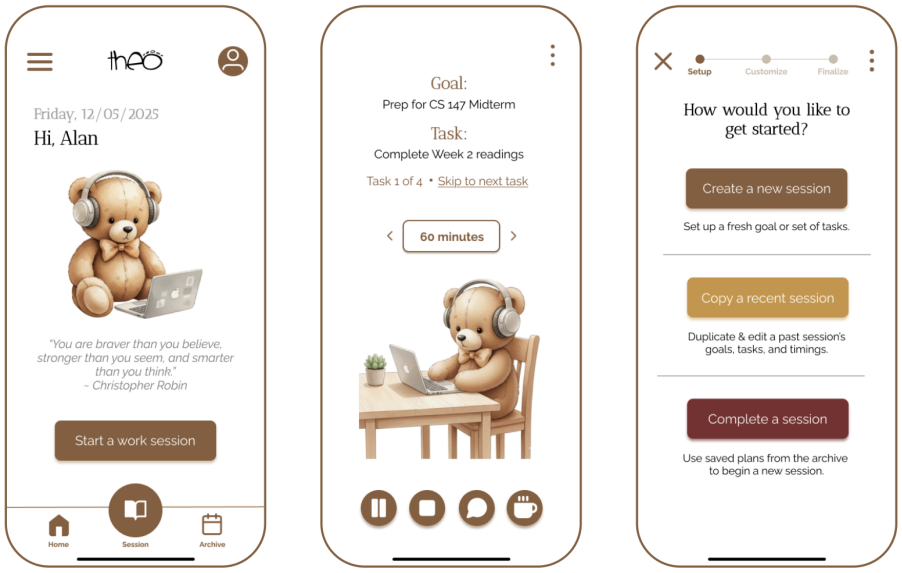

Figure 28. (From left to right) Minimal clickable components per screen, consolidated similar buttons, tricolor clickable hierarchy.

EMPATHY

Presenting personality lowers the barrier for entry and induces a feeling of support.

Using a warm color palette and welcoming language, we highlight our desire for users to feel comfortable and safe in the app space. Inspired by a cozy autumn fireplace ambience, our analogous color scheme creates an intimate, nonsterile environment. We frame actions as questions to give users a sense of control without being directive and include positive affirmations and images that accentuate our goal to provide support. By using a teddy bear mascot and figure throughout all processes within the app, we hope to evoke a sense of friendliness and passive encouragement, as physical teddy bears are expected to give to children.

Tradeoffs: Our color palette and use of a teddy bear appeals more to a younger, predominantly female demographic; this risks limiting the inclusivity and broader appeal of our app. Additionally, these two aspects that are the core of our product may hold larger significance for Western audiences, due to their greater prominence and significance in these regions. This was a difficult tension to resolve due to the large variety of connotations of color and character based on culture. To maintain the novelty of our solution and differentiate our approach from alternatives we researched through a competitive market analysis, we decided to retain our original mascot as the most globally relatable and acceptable figure for the purpose of artificial body-doubling.

FLEXIBILITY

Providing multiple methods of accomplishing tasks enables users to work in a way that best fits their preferences and circumstances.

Perhaps the most significant and underscored value in our design, flexibility emerges through our design's attention to maintaining human control over the course of a workflow. We understand that our users come from a variety of backgrounds, levels of education, and abilities; we hope that each individual may

discover their optimal work environment with Theo through its customizability. Both our goal breakdown flow and in-session experience are structured to offer a user all the necessary accommodations for a circumstance in which they would set up and conduct all actions manually without the app.

Tradeoffs: While we aim to provide enough options to minimize the possibility of a user feeling constrained by the capabilities of the app, we realize that an excess of choice can be overwhelming. Especially for our target audience, this abundance has the potential to create additional anxiety or distribution of attention, detracting from our primary goal of focus. To counteract this, we simulated the flow of a process funnel for session setup, handling each aspect of the creation process step by step, while discretely providing options for emergency exits or editing capabilities in reach of users without appearing at the forefront. We similarly used strategic placement of customization options for our in-session experience, using commonly recognized icons to indicate navigation and a vertical 3-dot menu to store all other actions.

Final Prototype Implementations

TOOLS USED

Our high-fidelity prototype was built as an Expo project using a React Native/Typescript framework. Backend implementation involved usage of a Supabase database to store user, session, and task information, as well as a Gemini and Whisper API for AI assistance in goal breakdown/chatbot reflection and voice input, respectively. The project's version control was managed through Git/Github.

REACT NATIVE/EXPO

Building our project with Expo allowed us to work with quick turnaround and view changes to the code immediately. The provided Expo project templates made the setup process for our tab navigation system simple and straightforward. These tools also easily accommodated our team's varied access to device operating systems (ranging from iOS to Android), making our product cross-platform compatible. The large library of publicly available packages made visual rendering seamless, especially regarding our session archive calendar view.

While very accessible and beginner-friendly to use, some aspects of React Native and Expo proved restrictive, specifically routing options. Workarounds and reconfigurations of navigators were necessary at times to ensure that all routing was operating as a user would expect.

Overall, this external framework more than adequately equipped us to produce an app of very close likeness to the expectations set by our medium-fidelity prototype in Figma.

SUPABASE

Supabase's developer-friendly browser interface made table setup a very streamlined process, which subsequently simplified the integration of selection, insertion, and update into the main codebase. The built-in authentication table (auth)

reduced the amount of coding from scratch necessary to set up initial tables for users and account information. Real-time updates were quick, and the free provisions offered all the necessary tools to sufficiently run the application.

However, the Row Level Security (RLS) policies proved a steep learning curve to understand and tackle. Ultimately, we found that temporarily disabling the RLS was easier to manage. We recognize that this is not conventional practice, and the need for disabling the RLS arose primarily due to the modifications to the backend structure that were made. Debugging issues related to the database functionality was additionally challenging, due to the relatively unhelpful nature of the returned error messages.

GEMINI

Two of our three main tasks depended on AI usage: the goal breakdown and the reflection chat. For goal breakdown, we used Gemini to transform a single, high-level user goal into a structured list of smaller, time-bound subtasks that could be scheduled and tracked within our session flow. For reflection, we prompted Gemini to act as a gentle, nonjudgmental companion, asking follow-up questions and mirroring the tone of our teddy bear mascot to maintain emotional consistency across the app experience. Carefully engineered prompts, including explicit role instructions, few-shot examples, and desired output formats (e.g., JSON-like structures), were essential to obtain responses that were both parseable and aligned with our values of empathy and flexibility.

Because we were working with a general-purpose LLM, safety and robustness became major considerations. We added guardrails in our prompts to discourage harmful or overly prescriptive advice and designed our UI so that AI suggestions were always framed as optional, editable recommendations rather than directives. On the implementation side, Gemini's response structure and occasional malformed outputs required us to build multiple layers of fallback logic, including automatic retries, simplified backup prompts, and manual-edit states when parsing failed.

Although this increased development complexity, we ultimately chose to continue with Gemini due to its cost effectiveness, TypeScript-friendly SDK, and sufficient quality for our use cases.

WHISPER

To support more natural interaction and richer data collection around user goals, we integrated OpenAI's API, powered by OpenAI's Whisper, for voice input. Adding voice control offered an alternative for input in situations where typing might be difficult, and enhanced the simulation of conversing with a peer. At present, all recognized speech is funneled into the goal input, enabling users to "brain dump" detailed descriptions of what they want to accomplish. This opens the door for future enhancements, such as synthesizing spoken notes into a concise goal statement.

From an interaction design standpoint, we wrapped voice input in a dedicated modal that prioritizes accuracy and user control. The modal includes clear affordances to start and stop recording, a "record again" option for quickly discarding unsatisfactory attempts, and a final confirmation step before text is committed. These safeguards help mitigate common issues with speech recognition, such as background noise or misheard phrases. However, relying on a cloud-based transcription service introduces tradeoffs in latency, dependence on network connectivity, and the need to transmit short snippets of audio to a third party. We addressed these concerns by keeping recordings brief, surfacing editable text before submission, and positioning voice as optional rather than a requirement.

WIZARD OF OZ TECHNIQUES & HARD-CODED ITEMS

Our app is a fully-functioning, standalone product with a completely implemented backend that uses APIs and a database. When users create accounts for themselves, the app provides a fresh version of itself with an empty session archive and default settings. Only after creating their own sessions will a user's archive begin to populate, retrieving data from our Supabase database.

Thankfully, we have been able to replace all Wizard of Oz techniques with full backend functionality, specifically in the goal breakdown feature and reflection chatbot. Both use the Gemini API to provide real-time AI-generated answers to user input and requests.

AI INTEGRATION

A combination of AI tools were used to maximize the learning experience for users. We used the Gemini API to implement an empathetic AI chatbot and help break down complicated goals. OpenAI's Whisper API was used to implement voice features for a more natural and seamless user experience. Additionally, all teddy bear-themed images were generated using Gemini's image creation feature.

Reflection & Next Steps

MAIN LEARNINGS

THE DESIGN THINKING PROCESS

The importance of quantity: Generating a large quantity of ideas in the early stages of design was far more useful than polishing a few initial concepts. The "wilder" options pushed our creativity and surfaced issues we needed to consider. For example, while creating a physical teddy bear companion embedded with AI or projection capabilities was beyond our reach, imagining a toy that functioned as a true body-doubling presence triggered reflection on the privacy risks of voice input.

Maintaining purpose across iterations: The constant iterations indisputably improved our project; however, we occasionally found ourselves losing sight of our value proposition and purpose when focusing on small details. We learned that at each iteration, taking steps back to refocus on our core goals drove our ideas to higher heights.

Design inspiration can come from anywhere: As we worked through the process, we discovered that meaningful inspiration often emerged from unexpected sources. Casual conversations, unrelated apps, even small personal frustrations and moments of silence and reflection revealed ideas that would augment the visual and functional quality of our work. These moments reminded us that design insight isn't limited to formal research—it can surface at any time if we stay observant.

AI FOR EMPOWERING LEARNERS

The range of learners: Learners are ubiquitous, and their needs run the gamut from completing a school assignment to developing new skills. Early in the process, we considered designing a tool for older adults to understand technology as well as for children developing early digital habits. Each group required a completely different approach to pacing, scaffolding and interaction. This thinking encouraged us to view

“learners” as a global population of diverse individuals whose backgrounds shape their responses to different forms of education, especially through technology.

AI as support: The most effective designs we explored in the studio treated AI as a tool that scaffolds a user's agency instead of taking over their tasks. When we considered moments where AI should guide or clarify rather than act, the experience became trustworthy and adaptable. This shift clarified that empowerment comes from strengthening the user's capacity.

OUR PROJECT

The tension between personalization and a clean interface: Our audience isn't uniform, and ADHD individuals often desire control over their environment, which pushed us toward high levels of customization. The challenge became translating that desire into a product that felt coherent and approachable. We had to pare down options, establish a solid base experience and let users uncover additional controls at their own pace. Thoughtful use of menus, tooltips and settings let us offer flexibility without overwhelming the screen.

“Thinking outside the box” vs. finding a niche: This concept of reinvention drives new designers to attempt to produce groundbreaking work that discovers a novel approach to addressing a problem. While this first appeared to be a guaranteed path to success, it was sometimes fruitless to view ourselves as pioneers. We learned to draw from the successes and failures of existing tools and listen to the underlying needs of our users, who were already immersed in the current technological era.

THE FUTURE

More robust AI usage: It is an aspiration to make the AI component more reliable and secure. The current model still fails in predictable ways, and cost limits what models and functionality we have access to. Using a chatbot reflection feature introduces real risk, so strengthening accuracy, guardrails and monitoring would be a priority before scaling this feature.

Map feature: We would also explore a map feature that helps users find nearby friends or peers looking for human forms of body-doubling, creating opportunities for more direct co-working and accountability.

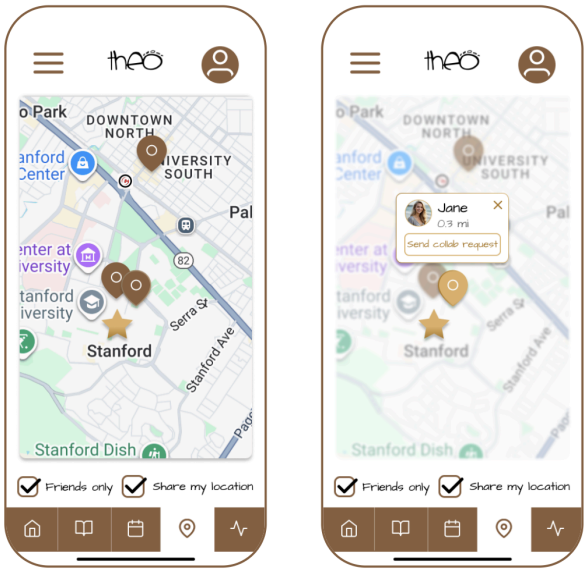

Figure 29. A medium-fi prototype of a potential map feature to discover nearby human collaborators.

Account verification: We would add proper verification, including email confirmation, multi-factor authentication, SSO, and GoogleAuth. The current flow accepts any email without validation, which leaves the system open to impersonation and weakens trust.

Figure 30. A medium-fi prototype of progress visualizations.

Visualize progress and show trends: We would also explore how we can depict progress and trends over time in a form that motivates rather than discourages users. The goal would be to highlight growth, patterns and small wins without framing lapses as failures. Below are our initial sketches for this feature; further research and feedback would be critical for confirming a format that would evoke the desired positive reactions and behavior from users.