Traverse

Final Report

CS147 Fall 2025
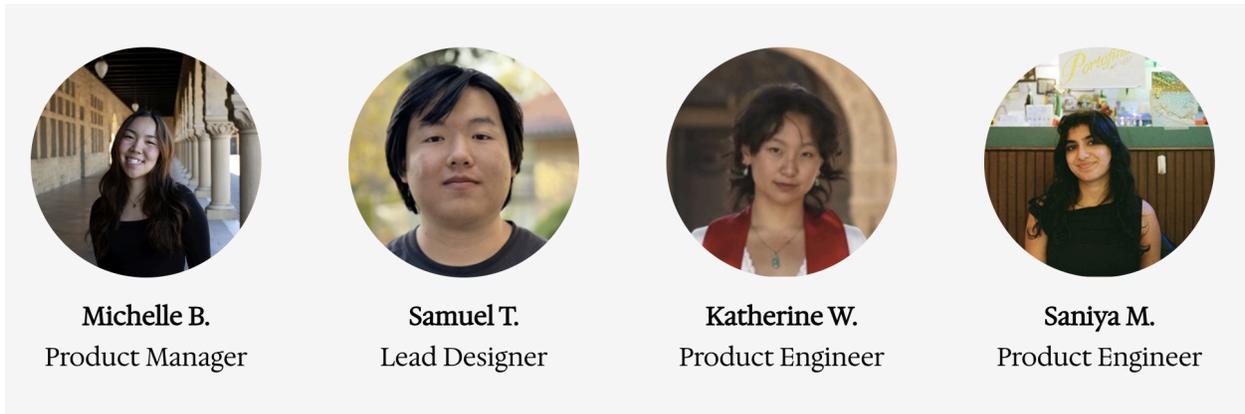Agents of Change
Katherine W., Michelle B., Samuel T., Saniya M.

Table of Contents

Project Name & Value Proposition

Project Name: Traverse
Value Proposition: Explore. Connect. Create.

_____

Team

| Michelle B. | Samuel T. | Katherine W. | Saniya M. |
| Product Manager | Lead Designer | Product Engineer | Product Engineer |

_____

Problem and Solution Overview

Problem: As digital spaces for arts engagement prioritize consumption over emotional depth, casual creatives struggle to find trust in new communities and wonder when exploring the artistic world

Solution: An AI-powered platform for users to explore the connections between artworks and add their reflections – deepening artistic engagement and jumpstarting curiosity

_____

Needfinding

a. Interviews

We spoke with a mix of casual creatives, hobbyists, and art professionals to understand how people discover art, build meaning, and connect with others. Our interviewees included:

- **Jennie** – creative but not formally trained; sees art as emotional release
- **Andi** – maker and fabric-store teacher; active online but values in-person connection
- **Mary** – museum coordinator; loves helping others engage with art
- **Grace** – curatorial fellow; cares about access, context, and validation
- **Skye** – artist and gallery assistant; craves community but hates shallow platforms
- **Mila** – software engineer and writer; selective about sharing and wants safe, trusted spaces
- **Daisy** – independent artist; relies on Instagram but feels burned out by it
- **Jay** – successful artist/professor; values dialogue and context over visibility

*All pseudonymised for privacy.

Recruitment happened through on-site visits at Cantor and personal networks. We offered a conversational and relaxed interview format with two researchers – one who asked questions and one who took notes. Everyone signed a consent form giving permission for us to use their sentiments in our project. Interviews lasted 25-60 minutes long and in some cases, we compensated with coffee. We used audio recordings, handwritten notes, and follow-up questions based on each person's background.

b. Synthesis

For each interview, we constructed an "empathy map," organized by what the user "does" and "says" and then made insights into what they "think" and "feel." This helped us detect patterns across our two weeks of interviews. Three themes kept repeating:

1. **Art helps people express emotion and identity.** Users want a personal connection, not just information.

2. **Digital platforms are useful but mentally draining.** People like access and inspiration online but dislike algorithms, performative posting, and shallow engagement.

3. **People want curiosity and trust.** Users crave spaces that feel safer, slower, and more reflective—without pressure for exposure or perfection.
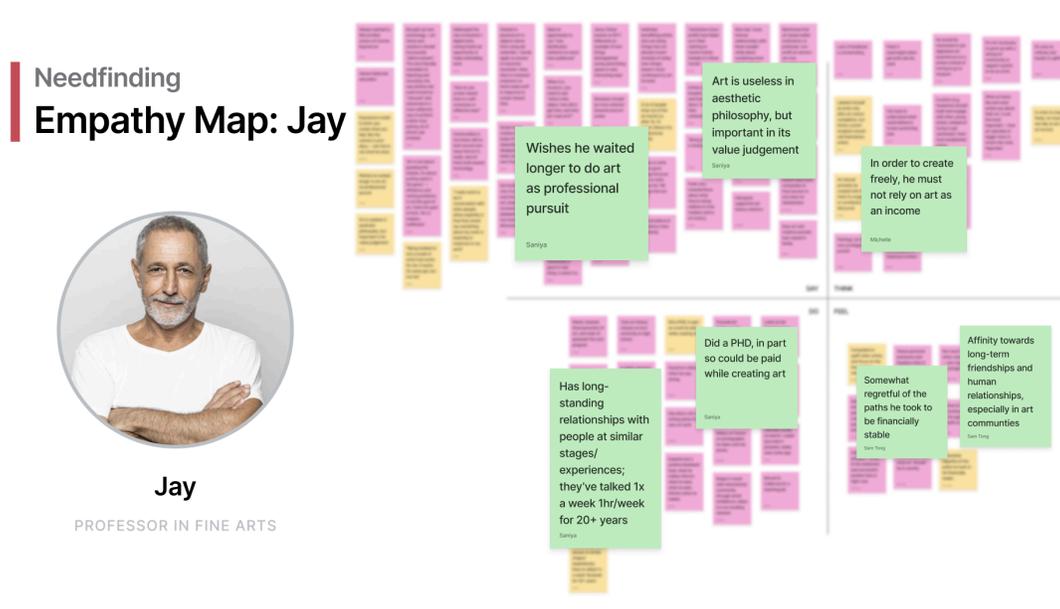
Figure 1: The empathy map for Jay, a successful artist, art critic, and professor who has valued being in art circles all his life.

_____

POVs & Experience Prototypes

a. POV statements

We developed Point of View (POV) statements for Mia, Daisy, and Jay, which are structured summaries of a user's perspective, intended to highlight surprises or tensions found in the interview.

POV 1 – Mila
- **We met:** Mila, a full-time software engineer and long-time creative writer who recently moved from Chicago to San Francisco.
- **We were surprised to realize:** that even though Mila creates art to connect with her creative community, she's selective about who she shares her writing and art experiences with and has avoided finding other creatives in SF since moving here.
- **We wondered if that means:** that Michelle seeks deeper, trusting connections over just visibility and exposure due to the vulnerable and personal nature of her writing, and the tech culture of SF has prevented her from feeling confident in finding these relationships.

- **It would be game-changing to:** make processes for sharing art experiences feel safe, genuine, and accessible for Mila to find more immediate motivation and fulfillment from her community.

POV 2 – Daisy
- **We met:** Daisy, an independent artist and recent college grad who shares and sells her art online
- **We were surprised to realize:** that even though Daisy relies on Instagram to share her work and connect with other artists for inspiration and platform, she mostly engages on Instagram with in-person friends, struggles to put in the extra work to build new digital connections, and experiences burnout from posting daily.
- **We wondered if that means:** that for Daisy, trying to get visibility via the Instagram algorithm actually hinders the more personal connections that motivate her and sustain her creative growth.
- **It would be game-changing to:** if creative social engagement felt like reciprocal and authentic exchange rather than being performative and draining.

POV 3 – Jay
- **We met:** Jay, a successful artist, art critic, and professor who has valued being in art circles all his life.
- **We were surprised to realize:** that despite his success, Jay regrets pursuing art professionally early on and feels that it forced him to create for recognition and market appeal rather than authentic creative exploration and the productive dialogue at the heart of art.
- **We wondered if that means:** that for Jay, creating for the market detracts from conversation around art, which unlocks the true value.
- **It would be game-changing to:** make the context and conversations around art more accessible to inspire curiosity and personal engagement.

b. HMW Questions

The POVs led to the creation of How Might We (HMW) questions, asking how we could transform insights from our POVs into solutions. We selected our top three HMWs for each POV, as well as our top three HMWs overall.

From Mila:
1. How might we create intimate creative spaces?

2. How might we make people feel more safe about sharing their art without being put in a vulnerable state?
3. **How might we help creatives like Mila find trust and emotional fulfillment in unfamiliar creative spaces, rather than just exposure?**

From Daisy:
1. How might we encourage people to share art daily without being trapped by algorithms?
2. How might we connect artists who want to be mentored, and ones who want to act as mentors?
3. **How might we redesign digital spaces to inspire wonder and emotional fulfillment through the discovery of human art and the natural world?**

From Jay:
1. How might we change how art is presented (like at museums)?
2. How might we encourage dialogue and conversations among people and spark new connections?
3. **How might we make exploring art contexts and ecosystems more childlike, like solving a puzzle?**

c. Solutions

We brainstormed numerous solutions from our top HMWs and selected our top three solutions.

1. A mobile app where users take a photo of artwork, add reflection, and see what others thought about the same piece; location-based map.
2. A mobile app where users group art together in personal ways that can be shared and commented on by others.
3. A mobile app with an art database indexed by AI; placed on large graphs that users can traverse through to find connected artwork.

Figure 2: Solution brainstorming for our top 3 HMWs questions

d. Experience Prototypes

From each solution, we are making an assumption about our user's behavior or preference. In order to test the validity of these assumptions, we generated an experience prototype for each solution.

Prototype 1: Reflection Jar at Cantor
- Assumption tested: People will reflect on art if the process feels simple, low-pressure, and situated in the moment.
- Key aspects of setup: We placed a jar, sticky notes, and pens at the Gates of Hell sculpture at Cantor and invited visitors to share a reflection.
- What worked / didn't, implications: People paused, wrote notes, and showed clear interest. Over four hours, seven visitors added reflections. But users could not see past reflections later or build ongoing meaning. This validated that reflection is natural and welcomed but needs a more sustainable digital home.
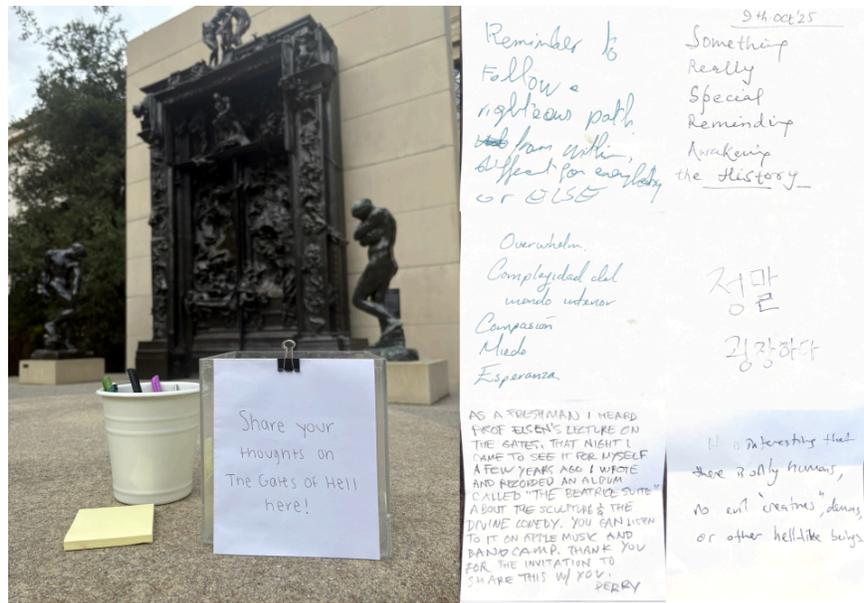
Figure 3: Experience Prototype 1 with our reflections jar and some of the reflections left.

Prototype 2: Art Grouping Activity

- Assumption tested: People enjoy making personal connections between artworks and will share their reasoning if invited.
- Key aspects of setup: We printed 16 artworks and asked participants to group them in any way they felt was meaningful.
- What worked / didn't, implications: Participants engaged deeply and explained their choices emotionally. One participant said they wouldn't make connections without the prompt, which showed the idea is engaging but needs built-in guidance. It reinforced that prompting curiosity works.

Figure 4: Experience Prototype 2 with one of our participants selected groupings.

Prototype 3: Traversable Art Graph
- Assumption tested: People will explore related artworks in an open-ended, curiosity-driven way.
- Key aspects of setup: Participants chose an artwork from a board, then uncovered related images in a branching pattern, navigating several steps deep. We tested both open exploration and goal-oriented exploration.
- What worked / didn't, implications: Engagement was high. Participants described the experience as feeling like a "game." They kept exploring even when images repeated. One preferred no goal; another preferred a specific goal. This validated the central concept: the graph supports curiosity and sustained attention. It led directly to choosing this as our final direction.
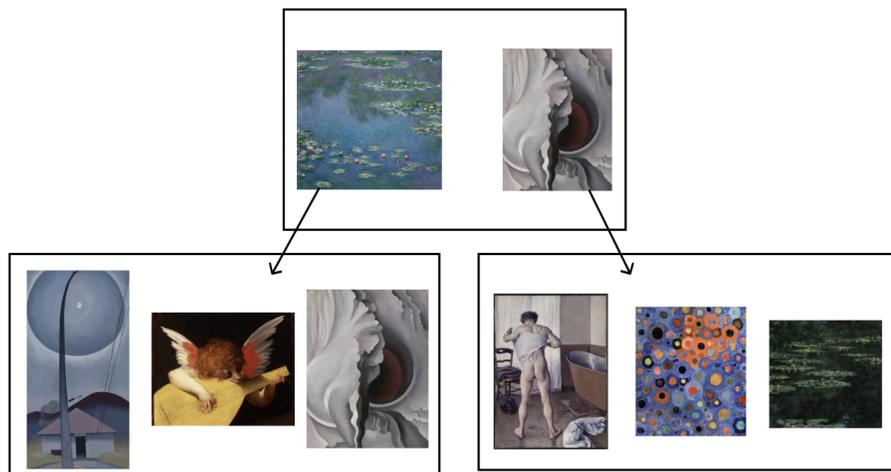
Figure 5: Experience Prototype 3 with an example of the screens participants were shown when traversing art.

_____

Design Evolution

a. Final Solution

Our final solution is Traverse, a mobile app for exploring interconnected artworks and leaving reflections. The app centers on a dynamic graph of artworks that users can move through at their own pace, opening pieces that interest them and adding reflections by

either selecting a region of the artwork or writing a simple note. Users can also upload their own work and tag related pieces, contributing back to the graph.

The decision to move forward with Traverse comes from what we learned from our experience prototypes. Our first prototype tested whether people would reflect when given simple materials and a prompt. This was validated as visitors stopped, read the prompt, and left thoughtful notes, but the main limitation was that they couldn't revisit their own reflections or engage with others' thoughts. This showed us that reflection matters, but it needs a space where people can return to what they wrote and see what others noticed. This inspired our reflection features on Traverse. The third prototype tested whether exploration is motivating when users can move through a set of related images. Participants were highly engaged, spent several minutes traversing, and compared the activity to a game. This was our strongest evidence that open-ended exploration can sustain attention and spark curiosity when users feel guided by relationships between artworks.

Together, these results pointed us toward a solution that combines a place to reflect, a structure that surfaces connections, and an exploratory environment that encourages wandering.

b. Tasks

Simple Task: Explore the artwork graph.

This task establishes the foundation of why people use Traverse: to discover art in a way that feels open, intuitive, and curiosity-driven. The graph allows them to zoom, tap, and move through connected works, by giving them a clear starting point and a natural path without needing instructions or prompts.
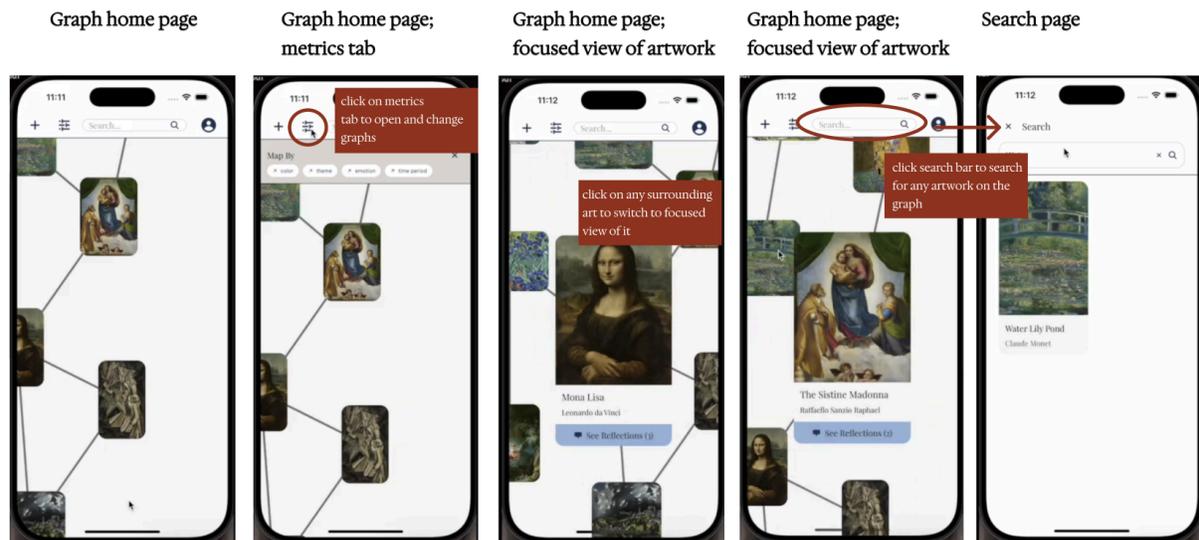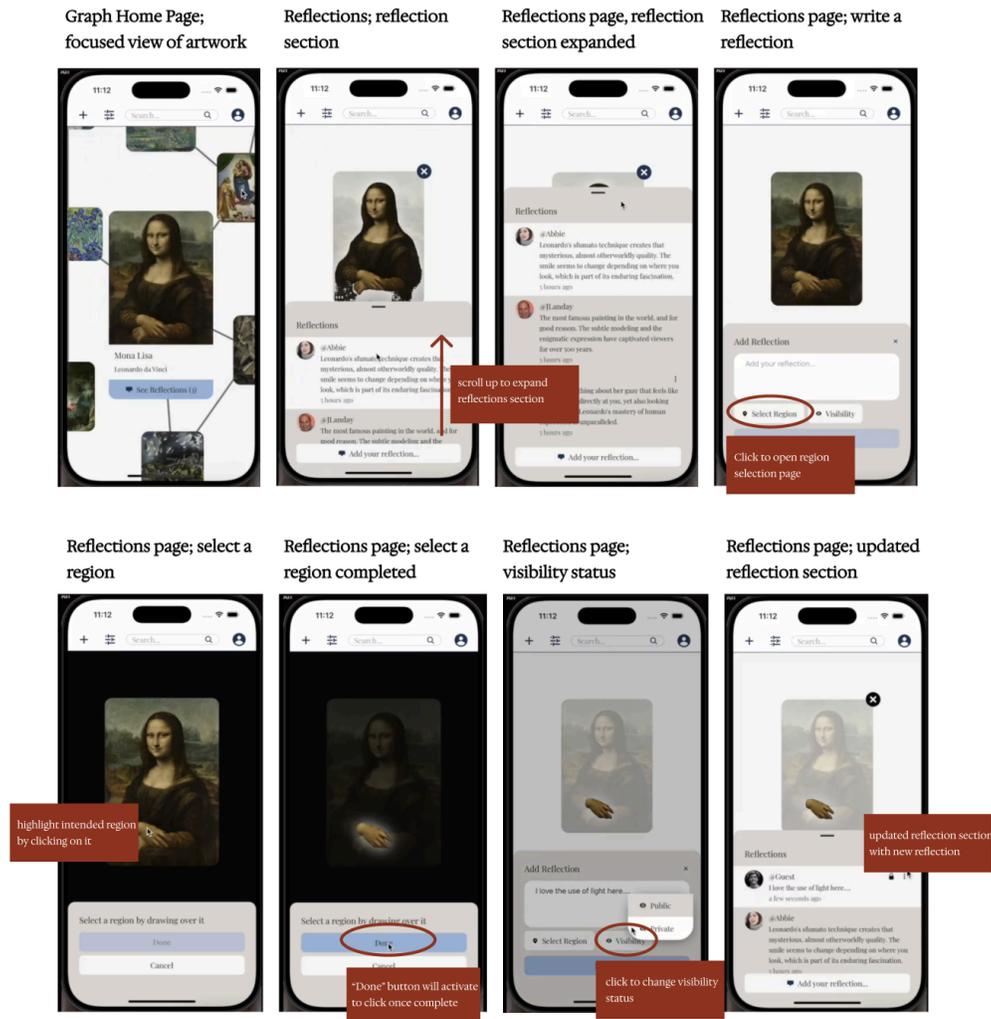
Figure 6: Annotated simple task flow from high-fidelity prototype

**Moderate Task: Leave a reflection on an artwork.**

This task allows users to move from exploring the graph into engaging with a specific artwork more deeply. When users tap into a piece, they can read reflections that others have shared, making their artwork feel richer by showing how different viewers experience the same piece. While writing their own reflection, they can select a region of the artwork to show exactly what part they are responding to. Users can also choose between posting publicly or keeping their reflection private, depending on whether they want a community-oriented space or a more personal, diary-like one.

Figure 7: Annotated moderate task flow from high-fidelity prototype

Complex Task: Upload personal artwork and tag related works.

This task lets users contribute their own creative work to the platform and place it within the larger network of art. This transforms Traverse from a viewing tool with a fixed collection into a shared, community-driven creative space. When uploading an artwork, users add basic details and then tag related pieces that inspired or connect to their work. Tagging helps the system understand how their piece fits into the graph and creates meaningful links for future users to explore.
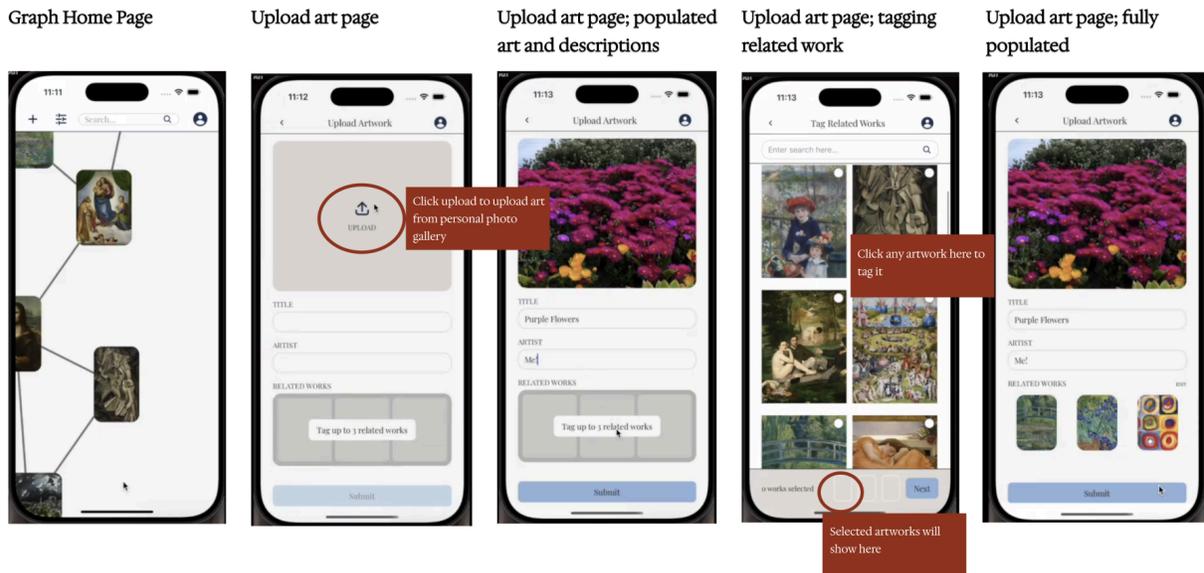
Figure 8: Annotated complex task flow from high-fidelity prototype

c. Design Iterations

Low-Fidelity Prototype

Our first implementation of Traverse was a paper prototype that captured the core screens and interactions of the app at a very early stage. The prototype was drawn digitally on an iPad and printed out to simulate a mobile interface. This format let us test whether the foundational ideas behind Traverse—graph exploration, reflections, and artwork uploads—were understandable and usable without the influence of polished visuals. It also helped us evaluate early task flows and navigation patterns before building anything.
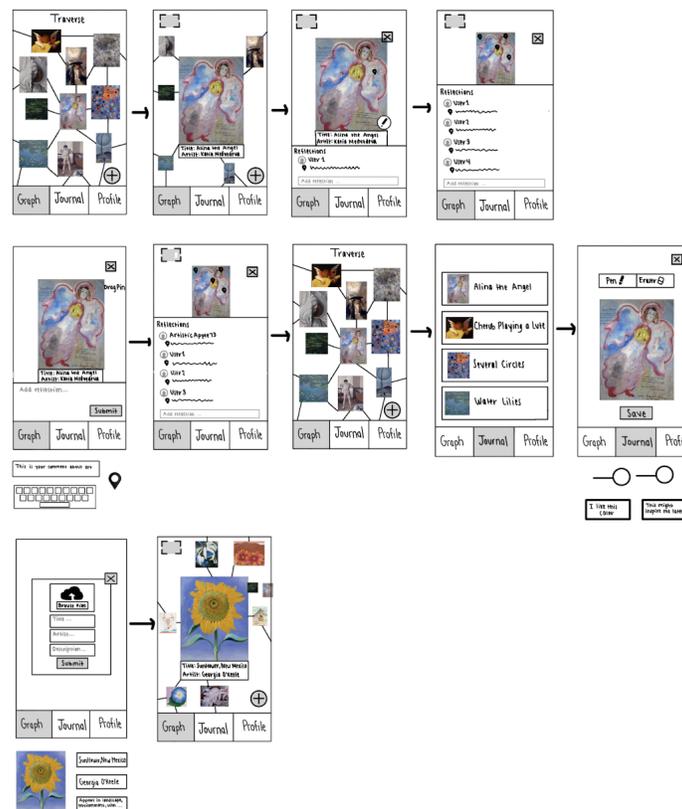
Figure 9: Screens from the low-fidelity prototype

Low-Fidelity Prototype Usability Testing

We conducted in-person usability testing at the Cantor Arts Center with four participants who varied in their engagement with art. Each participant completed five sequential tasks: exploring the graph, opening a specific artwork, viewing comments, leaving a comment, creating an annotation on the artwork and saving it to the journal, and finally uploading their own piece. After completing all tasks, we conducted a short verbal survey on ease of navigation, satisfaction, and overall impressions.

We tracked misclicks, time on task, ease of use, and likelihood of installing the app. Overall:
- Ease of use: 2.375/7 (1 = easiest), suggesting users generally understood the interface.
- Likelihood to install: 5.25/7, indicating strong interest in the idea.
- Efficiency: Task 1 took the longest (avg. 2m:23s), while all other tasks were under ~1m:15s.

- Misclicks: Highest for Task 1 (avg. 3.75), while later tasks had far fewer mistakes. These findings showed that initial navigation through the graph needed improvement, while the upload and reflection tasks were clearer to users.

Findings from the usability tests shaped the direction of our next iteration:
- Unify comments and annotations. Users didn't understand the difference. We consolidated them into a single "Reflections" system.
- Rework graph interactions. We adjusted how clicking vs. zooming behaved to make traversal clearer and more predictable.
- Simplify navigation. Early feedback showed the need for fewer tabs and more intuitive hierarchy, leading us to streamline the home screen and reduce pathways.
- Support both private and public meaning-making. Confusion around journals pushed us toward a more flexible reflection flow that allows both modes without separate pages.

Medium-Fidelity Prototype

Our next major iteration was a medium-fidelity prototype created in Figma. This version focused on refining the core screens, layout, and interactions of Traverse so we could better understand how the graph, reflections, and upload flows should feel in a real app environment. Figma allowed us to build more polished interfaces and use transitions to simulate movement.
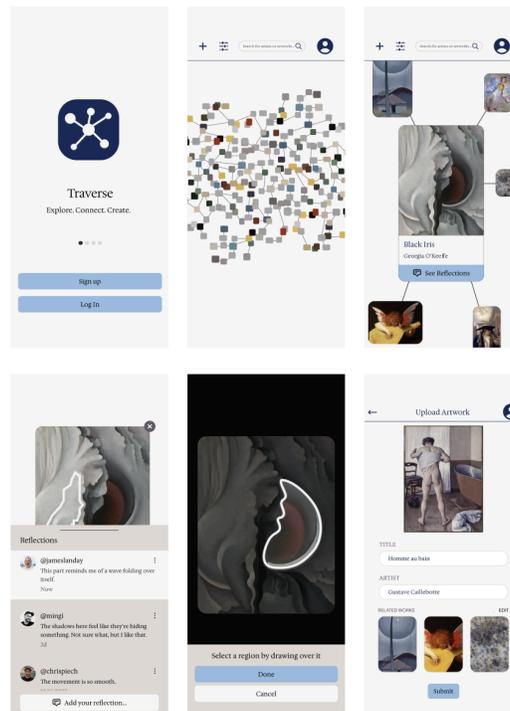
Figure 10: Some screens from the medium-fidelity prototype

Medium-Fidelity Heuristic Evaluation

After completing our medium-fidelity prototype in Figma, we exchanged prototypes with another group to conduct a heuristic evaluation using Nielsen's 10 Usability Heuristics and the two CS147-specific heuristics.

Reviewers documented 53 total violations, with 24 rated at severity 3–4, indicating several issues that could interrupt or confuse user interactions. The most categories for violations were Visibility of System Status, Match between System and World, and Consistency and Standards.

Revisions

We used this feedback to revise key parts of Traverse's design, strengthening the usability and helping prepare the design for its high-fidelity implementation. Here, we outline the higher severity violations highlighted organized by each task and the revisions we made to address them.

Onboarding Flow

Problems	Fixes
Zooming and dragging interactions are unclear	Added onboarding screens introducing the user to the app's core tasks
The fully zoomed-out starting view is overwhelming and unclear	Added a short personalization question to guide starting point
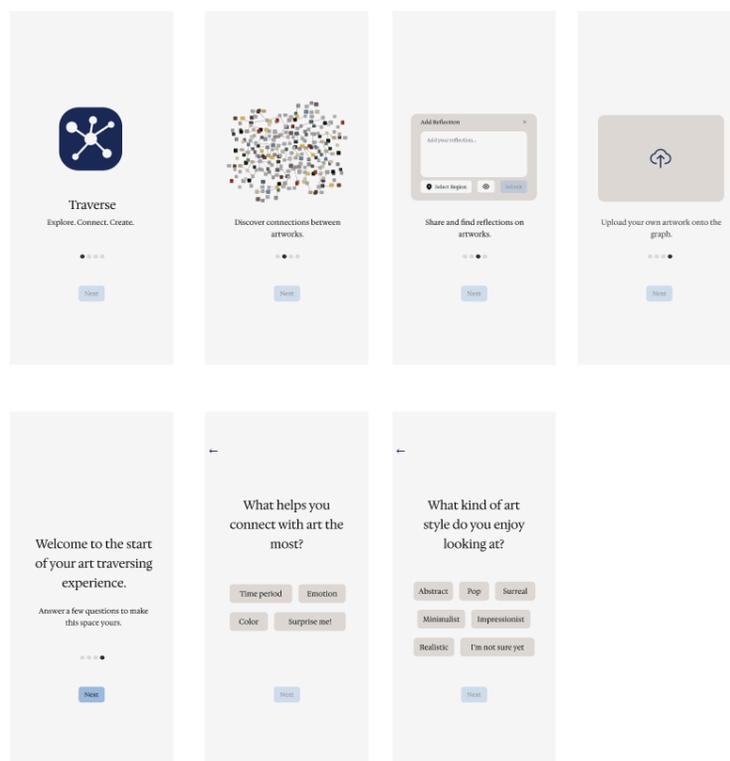
Figure 10: Onboarding screens from the medium-fidelity prototype

Simple Task

Problems	Fixes
The graph icon (used as a back button) does not match user expectations	Removed the mini-map as users will be able to pinch to zoom out

Artwork images do not clearly behave as tappable elements	Redesigned the main artwork card so it reads as one clear tappable element
Confusion about the function of the the "Metrics" tab	Renamed feature to "Map by"
"Metrics" tab lacks options for users who rely on non-visual patterns	Changed graph options from "color, object, style, and time period" to "color, theme, emotion, and time period"

Before   After
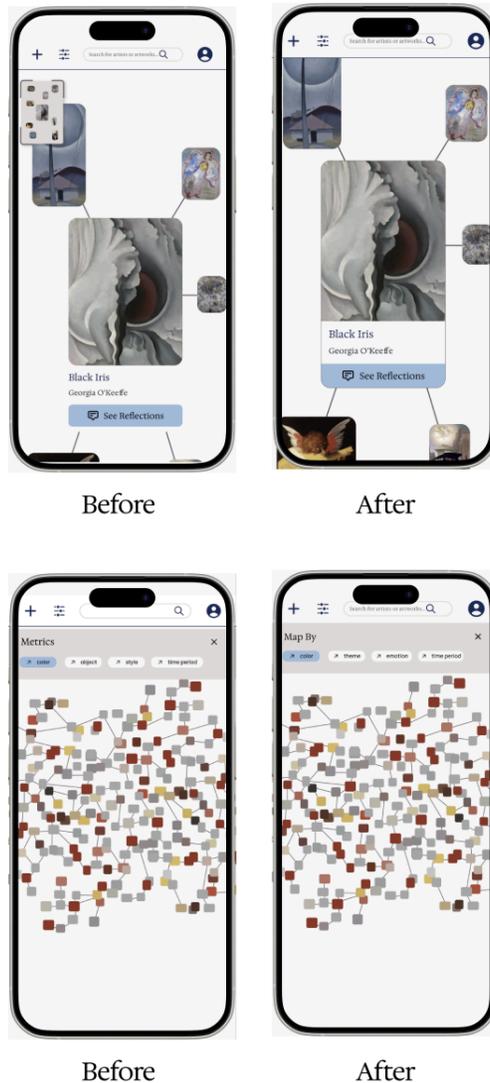
Before   After

Figure 11: Before and after screens of the simple task revisions

Moderate Task Revisions

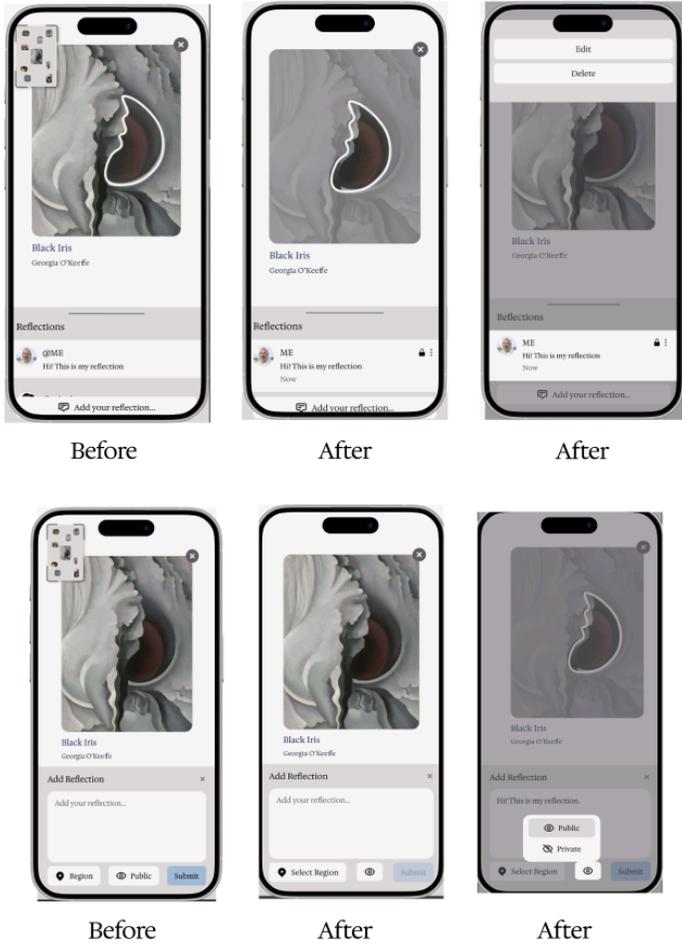| Problems	Fixes
Users cannot edit or delete their reflections once posted	Added timestamps so users know when each reflection was posted
No clear indication of whether a reflection is public or private	Added a lock icon to clearly mark private reflections
Highlighted regions lack visual contrast Reflections have no timestamps	Improved region highlighting with higher-contrast strokes
The "Submit" button is active even when no reflection is is typed	Added edit and delete options for personal reflections
"Region" button label does not match user expectations	Disabled the Submit button until text is entered
Public/Private toggle lacks visual clarity	Renamed "Region" to "Select Region" to set clearer expectations
	Replaced the public/private toggle with a drop-up menu

Figure 12: Before and after screens of the moderate task revisions

Complex Task Revisions

Problems	Fixes

The "Submit" button is active even when no file is uploaded	Disabled the Submit button until all required fields are complete
No loading or progress feedback during uploads	Added a loading indicator to show progress after selecting an image
Tagging UI does not show a maximum number of allowed related works	Updated the tag section label to "Tag up to 3 related works" to set the limit clearly
No option to edit or remove tagged related artworks	Added an Edit option for related works so users can revise or remove selections
Tagging UI does not show a maximum number of allowed related works	Added placeholder slots to show the 3-tag limit so users know how many related works they can select
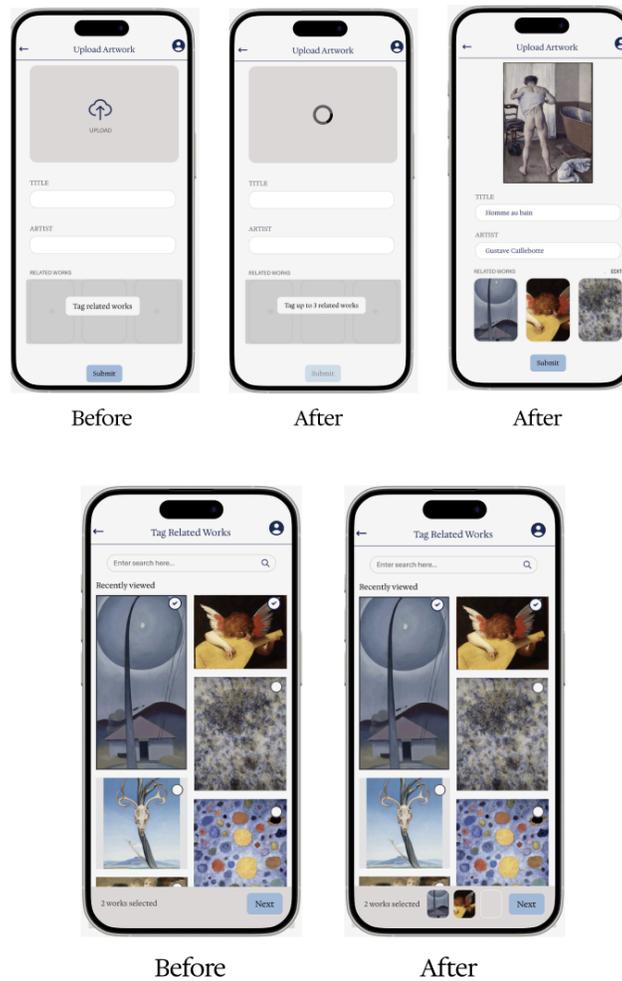
Figure 13: Before and after screens of the complex task revisions

d. Values in Design

Throughout the design of Traverse, we were guided by three core values: curiosity, authenticity, and accessibility. These values shaped how we envisioned art exploration and how each feature supports a meaningful, low-pressure creative experience.

Curiosity
-   A central value of Traverse is encouraging open-ended exploration. We wanted users to feel free to wander, follow visual threads, and discover connections without pressure or expectations. This value is embedded directly in the interactive graph, which presents artworks as interconnected nodes rather than a linear feed. The graph rewards exploration for its own sake—there are no likes, popularity metrics, or

predefined paths. Its design invites people to tap into their natural sense of curiosity and move through art in a self-guided, intuitive way

Authenticity
- We aimed to build a space where people feel comfortable engaging with art and expressing their own thoughts honestly. Traverse supports authenticity through features like public or private reflections, giving users autonomy over how they want to share. Public reflections allow users to connect with a broader community of art lovers, while private reflections create a personal, diary-like space for more intimate responses.

Accessibility
- Traverse is designed to be usable for both casual creatives and art enthusiasts. Our mobile-first interface supports quick browsing and familiar gestures, lowering barriers to entry for users who may not have deep experience with art tools. Additionally, the ability to explore through visual patterns, themes, or emotions makes the graph approachable to people with different backgrounds and levels of knowledge.

Value Tensions

Curiosity: Curation vs. Openness
- We wanted Traverse to offer meaningful pathways through the graph, but too much curation could create a narrow or gatekept experience. At the same time, too little structure risked overwhelming users with an unfiltered map of artworks. We addressed this tension by offering lightweight organizational cues—such as theme, color, or emotional mappings—without dictating a specific route. This preserves openness while giving users optional guidance when they want it.

Authenticity: Visibility vs. Privacy
- Authenticity often relies on vulnerability, yet visibility can discourage honest expression. Some users want to share their reflections publicly to connect with others, while others prefer private, personal engagement. Traverse addresses this tension by embedding public/private controls directly into the reflection flow, letting users choose visibility for each individual reflection. This approach supports authentic engagement while avoiding the judgment or pressure that often emerges in social platforms.

_____

Final Prototype Implementation

a. Tools Used

1. React Native + Expo

Pros
- Enabled fast development and cross-platform support, allowing the prototype to run on both iOS and Android through Expo Go.
- Expo simplified device testing and sharing builds across teammates.

Cons
- Reloading the app within Expo Go was somewhat slow, impacting development speed and feedback loop time
- The lack of a purpose-built graph library in React Native required us to manually implement graph computation code (e.g. artwork placement, drawing connections, and animations during traversal)

2. Tailwind / Nativewind (Styling Library)

Pros
- Allowed for rapid UI iteration with utility classes.
- Provided consistency in spacing, sizing, and typography

Cons
- Not all group members had previously used Tailwind, so the increased consistency came with a tradeoff of increased learning curve.

3. Supabase (Backend)

Pros
- A straightforward, reliable backend for storing artworks, reflections, and user profiles
- Additionally supported file storage for artwork images and reflection region segmentation masks

Cons
- N/A

4. react-native-skia, react-native-reanimated, react-native-gesture-handler (graph rendering)

Pros
- Enabled custom computation/implementation of a dynamic artwork graph

- Enabled smooth animations and gesture handling for dynamic graph movement from user input

Cons
- There is no dedicated graph-layout library for React Native, so the graph had to be computed and positioned manually
- Required substantial custom code to simulate graph behavior
- Graph pan and zoom interactions were difficult to implement and were thus removed from the final hi-fidelity prototype

5. Segmentation Tools: Meta SAM 2 + Replicate.com
 Pros
- Enabled reflection region selection when adding reflections for the moderate task
- Provided high-quality segmentation masks with minimal user input

Cons
- Required nontrivial preprocessing, in particular translating user's tap on the screen to an x/y coordinate relative to the (scaled) image's pixels
- The model ran remotely on Replicate's servers, which introduced cost and latency tradeoffs
- For the prototype, adding a new reflection region only works while the server is manually enabled (disabled by default to avoid incurring costs).

b. Wizard of Oz Techniques

These features appeared functional to users but were simulated rather than fully implemented.

1. Graph Similarity Algorithm
- The UI suggests the graph can be reorganized by similarity metrics, but all "metric changes" simply display pre-made graph layouts instead of running ML models.
2. Artwork Upload → Dynamic Graph Update
- Uploading a new artwork appears to insert it dynamically into the graph. In reality, the upload always replaces the same placeholder node, and no new connections are computed.
3. Personalized Starting Graph
- The onboarding experience implies personalization, but all users begin on the same default graph.

c. Hard-coded Techniques

1. Pre-Populated Artwork & Reflection Dataset
    - The graph includes a fixed set of 24 artworks and manually written reflections. Nothing is fetched dynamically from an external database or API
2. Static User Profile
    - All testers use a single default "guest" profile with unchangeable profile information.
3. Simplified Upload Behavior
    - Uploads render correctly in the UI, but they are stored only on the local device, they always appear in the same graph location, and they do not persist across sessions.
4. Limited Interaction Model
    - Zooming was removed entirely due to gesture conflicts and complexity; panning operates only within a controlled region.

d. AI Tools

We used Copilot and Cursor to support development in targeted, lightweight ways. Copilot helped with debugging, boilerplate, and small utility functions, especially in areas where we had less React Native experience. Cursor was used mainly for speeding up UI layout and styling work, helping us iterate more efficiently without relying on AI for core logic or design decisions. Together, these tools helped speed up while allowing us to stay in control of the implementation.

_____

Reflection & Next Steps

a. Key Learnings

Design Thinking
This quarter shifted how we understand the design process. In past projects, we tended to start with a product idea and then run interviews to confirm it. With Traverse, we had to do the opposite: begin with needfinding, sit in the ambiguity, and let the problem space narrow over time. That meant listening more than speaking, asking open questions instead of pitching, and accepting that our first ideas were rarely the right ones. Moving from interviews to HMWs, to early concepts, to experience prototypes, and only then to lo-fi and hi-fi prototypes showed how each stage has its own purpose. We learned to focus on the

"bare bones" first, which is what problem we're solving and why, before worrying about aesthetics or implementation.

Studio
Our studio, Agents of Change, became a chance to practice framing and storytelling. Presenting in the studio and watching other teams present pushed us to be clearer about what Traverse is actually adding to the art world. We also learned how much the way you present matters such as pacing the slides, supporting whoever is speaking, and explaining technical decisions in language that still feels human. Those skills feel just as important as the prototype itself.

Our Project
Designing for the art community taught us that people's relationships with art are incredibly personal and varied. This reminded us that we didn't want to design for a certain level of expertise, but instead give people entry points and let them choose how far to go. We also designed for different comfort levels—those who want to observe, reflect privately, and contribute publicly.

Also, Traverse made the tradeoff between ideal design and technical feasibility very real. Early on, we let ourselves dream big—a dynamic art graph, AI-driven connections, region-based reflections—without thinking too hard about implementation. Once we started building the high-fidelity prototype, it became clear that custom graph rendering and segmentation were far more complex than a traditional app layout.

b. Next Steps

If we had more time, there are several directions we would explore to continue iterating:
1. Real graph algorithm and dynamic content
    - Instead of hard-coded artworks and connections, we would integrate with an artwork API and build an AI model to map relationships by color, theme, emotion, and time period. This would let the graph evolve as more works and reflections are added, and make the experience feel truly "alive" rather than fixed.
2. Richer region interactions for reflections
    - We would refine the segmentation pipeline and add more fluid animations when selecting or viewing regions—something closer to the subtle "glow" effects we admire in the Apple iOS photos app.

3. Full authentication and multi-user support
   - A real version of Traverse would need sign-up, log-in, profiles, and persistent data across devices. With more time, we would build a full auth flow and connect reflections and uploads to real user accounts so communities of artists and viewers could actually form.

c. Final Remarks

Overall, Traverse taught us how messy and rewarding it is to move from a feeling ("art should feel more curious and connective") to a working product. We're proud of where the prototype landed, and excited to use everything we learned in future projects.

We're so grateful to Professor Landay, our Course Assistant, Abbie Maemoto, and everyone in our Agents of Change studio, for all of their guidance and support throughout this process.

Figure 14: Our group at the Project Expo!