PORTFOLIO BRIEF · 2026

HCD AI Framework

Human-Centered UX Research for an AI Feature Initiative

Tariq Cilione · UX Researcher & Product Strategist · tariqcilione.com

PROJECT SNAPSHOT

What this work is

A two-week agile HCD sprint inside an early-stage AI product team. The goal was to determine whether an AI feature could provide measurable value to users, and to produce a research-grounded framework that development partners could build from.

This brief documents the research process, the pivotal behavioral finding, the deliverables produced, and the researcher's working philosophy for applying HCD methods to AI product decisions.

2	9	100%	9
week sprint	research participants	shared core pain point	persona framework

SPRINT METHODOLOGY

Four phases. Two weeks.

The sprint was scoped deliberately: mixed-methods survey to surface pain points at scale, then qualitative focus group synthesis to understand behavioral context, then experience mapping to locate the intervention point, then a persona framework as the primary design deliverable.

01
Mixed-methods survey
Qual + quant questions to surface pain points without leading participants.
QUANTITATIVE

02
Focus group synthesis
Rose, Thorn, Bud with affinity clustering to conceptualize behavioral themes.
QUALITATIVE

03
Experience mapping
Job-search journey mapped across time and mood to locate motivation collapse.
SYNTHESIS

04
Persona framework
Nine differentiated personas with AI perceptions and two edge-case segments.
DELIVERABLE

PIVOTAL FINDING

The research changed the design direction.

100% of survey participants shared the same core pain point. That level of consensus is unusual. The follow-up focus group added the critical behavioral layer: participants could not articulate what support they needed, which itself became a design input.

BEHAVIORAL INSIGHT

Motivation collapse, not skill gaps or tool access, is the primary barrier in the job-search experience.

The experience mapping session confirmed the pattern: the arc moves from frustration to fatalism without structured intervention. This became the behavioral rationale for AI-assisted coaching as a product direction.

RESEARCH DELIVERABLES

What was produced.

DELIVERABLE 1: 9-PERSONA AI FRAMEWORK

Nine differentiated user personas spanning individual job seekers, career advancers, career pivoters, and business clients. Each includes motivations, goals, pain points, tech comfort, and an AI Perceptions section that maps how each user type responds to AI interactions, where trust breaks, and what the feature needs to handle correctly. Two edge personas designed specifically to prevent exclusionary product decisions: a late-career user facing age discrimination in automated screening systems, and an over-credentialed academic exiting to industry.

DELIVERABLE 2: BEHAVIORAL EXPERIENCE MAP

The job-search journey mapped across time and mood, with standout moments flagged by participants as highest priority for intervention. The experience map anchors the motivation collapse finding in a specific arc of the user journey, and identifies the point at which AI-assisted coaching has the highest potential to interrupt fatalism before it becomes disengagement.

DELIVERABLE 3: HCD SPRINT DOCUMENTATION

Full LUMA System application presentation: methods selected, why each was chosen, how they were facilitated, Rose/Thorn/Bud reflections on what worked, lessons learned, and recommended next methods for continued research. This documentation demonstrates an applied and reflective HCD practice, not just method execution.

RESEARCH CONFIDENTIALITY NOTE

Client details and proprietary feature specifications are protected. The research process, methodology, behavioral findings, and deliverable framework shown in this brief represent my work as lead UX researcher. Full documentation is available on request for relevant opportunities.

MY APPROACH TO AI

How I work with AI.

WORKING PHILOSOPHY

"I study how people relate to AI. That means I have to be rigorous about how I relate to it myself."

I use AI across the full research lifecycle, from artifact creation and gap analysis to synthesis and strategic framing. Every output is treated as a draft that needs editorial judgment. My writing rules, my voice, and my research standards apply to everything AI produces for me.

SIX WORKING PRINCIPLES

Artifacts in, not prompts alone

Real research materials come first: survey data, Mural exports, persona docs. AI analysis of actual artifacts produces sharper insight than generation from description.

Iterative by default

Each request builds on the last. Sprint methodology, gap analysis, persona enrichment, writing rules: each layer is a deliberate decision about what the work needs next.

Voice discipline throughout

Documented writing rules enforced on every AI output: no corporate filler, empathy before credentials, short punchy sentences. AI that does not match my voice is a draft.

Deliberate about disclosure

Methodology stays visible. Product specifics stay private. That is a conscious research ethics decision, not a legal reflex.

Edge cases are not optional

Edge personas are designed in from the start: late-career ageism, over-credentialed pivots, AI-displaced workers. Research that only covers the median user harms those who need help most.

Trust is a design input

Every persona includes an AI perceptions section. How a user relates to AI is a product requirement, researched the same way as any behavioral variable.

BOTH SIDES OF THE SAME PROBLEM

AS A RESEARCHER

Studying how users relate to AI

- Trust calibration across user segments
- Automation anxiety as a behavioral variable
- AI perceptions mapped to design requirements
- Edge cases that surface ethical design gaps
- HCD frameworks applied to AI product decisions

AS A PRACTITIONER

Using AI to do the research

- AI as research collaborator, not oracle
- Editorial judgment applied to every output
- Sprint-paced iteration across a session
- Artifact-first inputs over prompt-only requests
- Voice and standards enforced throughout

HOW A SESSION ACTUALLY RUNS

01	02	03	04	05
Bring the artifact	**Ask for analysis**	**Apply a layer**	**Enforce standards**	**Ship the deliverable**

Real data, not description	Gap-finding, not generation	Each iteration has a purpose	Voice, ethics, disclosure	Partner-ready, not draft-ready

Tariq Cilione

Human-Centered UX Researcher & Product Strategist

tariqcilione.com | AI Trust & Ethics | Learning Science x HCD