

READABILITY EDITION

Heidi Portfolio Page Planning

User / Assistant chat log for implementation context

Prepared for implementation handoff • 5 September 2026

Reading guide. User messages are blue-tinted and assistant messages are grey-tinted. This transcript contains only the visible user and assistant conversation for the portfolio landing-page planning thread. Hidden system, developer, tool and internal reasoning messages are excluded.

Note: attached research files and the earlier clinical-trials prototype-planning chat log are referenced in the conversation but are not reproduced in full here.

USER • TURN 1

Written below is a task that I want you to carry out. Before we start, I want you to grill me on the details and minutiae of the plan. Be extremely comprehensive, and challenge any assumptions that I may have included. Ask me one question at a time until we can agree on a plan.

I have submitted an application to a job: <https://www.heidihealth.com/en-nz/careers/0d330297-625c-4e72-981c-860d21d3a25a>

After submitting, I found that there was a post 4 months ago from their head of product saying:

Woke up to the news that we're the 24th fastest hiring company in the world. Help me get us to 1:

1. find your 3 most cracked, detail-oriented friends (no formal exp necessary) - unusual hobbies, good at games (any), a little intense
2. send them to [heidihealth.com/careers](https://www.heidihealth.com/careers)
3. make them toggle to "Product & Design"
4. apply
5. [bonus] send something you've made to product@heidihealth.com to cut through the noise

I have decided I want to build a prototype to send it in and differentiate myself. But more than a prototype, I want to show my thought process. Because this clinical trials feature is new at Heidi, I wanted to show my thought process from the ground up. I want this project to be accessible via a link in the email.

I have an existing portfolio website called 'yasha.melbourne'. It is currently hosted on wix but I'd like to break free from the limitation for this project by using a subdomain.

I have created the following assets I would like to display on that page (assuming that the user gets an option between all three and I'd like to design how this page looks myself rather than this part being automated).

- I have the research that I used deep research for summarised
 - Potentially will change this asset to be more visual
- I have the thought process
 - via a Procreate timelapse
- I have the Prototype
 - and a document thought process/prompting via the chatgpt log

I would like to build a prototype hosted on chatgpt for me to review.

This is a design from my website previously. I'd like to keep the aesthetic and use this as a landing page.

- Under my name I want a large and wide card with a clear CTA of 'VIEW PROTOTYPE' with the subtext of 'How I challenged this job listing and created this in 2 days' - I'd like to have some options for alternative copy that better suit this CTA
- Under the large cta card, I'd like to have three hand drawn arrows coloured in grey leading to 3 different buttons, the design is to be the same as the large one but a suitable percentage smaller on the same hierarchy as each other (so on the same row but spaced apart)
 - Button 1 says: Data and Research
 - Button 1 subtext says: What data was used to make decisions?
 - Pressing the button opens up an overlay where I can show the data and research in a summarised visual way with links to the deep research PDF
 - Intent is to show how I use data to make product decisions
 - Button 2 says: Product Process
 - Button 2 subtext says: What was the problem? What defines success?
 - Pressing the button opens up an overlay where the video plays
 - Intent is to show how I worked through the problem, what success metrics and guardrails and why
 - Button 3 says: AI Prompting
 - Button 3 subtext says: How did I use AI through this process
 - Pressing the button opens up an overlay with summaries of when I used it why (will only provide 3 examples for brevity) with links to the chatlogs
 - Intent is to show how I work with AI and what I consider
- The above copy needs to be refined, I'd like a few options but I'd like the tone to be similar to the large CTA card
- I'd like to have an 'About me' inline hyperlink at the bottom the intent is that it should not interfere with the hierarchy but should still be visible so that the user does not need to scroll.

Colour code of the green is #88C5A5 and colour code of the grey is #555555. Do not create any design assets other than general UX/UI and arrows. Do not use AI generated art. Ask if specific icons are needed and I will source them myself

ASSISTANT • TURN 2

I've read the current listing as well. One thing I want to pressure-test throughout this process is whether the project demonstrates the exact traits Heidi is screening for, rather than merely looking polished. The role explicitly emphasizes working from ambiguity, finding the constraint that matters, building prototypes yourself, defining measurable outcomes, designing trust and human review into regulated workflows, and “default[ing] to building over requesting.”

I'll ask one question at a time, and I'm going to challenge scope, evidence, narrative, UX, technical choices, and what you should deliberately leave out.

Question 1

What is the single thing you want the recipient to believe about you after spending 60–90 seconds on this page?

Not what you want them to see—what conclusion should they walk away with?

For example, these are meaningfully different:

“Yasha can independently take an ambiguous clinical product problem from research → reasoning → prototype very quickly.”

“Yasha has unusually strong product judgment despite not coming from a traditional clinical-trials background.”

“Yasha understands the clinical-trials problem deeply enough that we should interview her.”

“Yasha is exactly the kind of intense, hands-on builder Heidi hires.”

“Yasha knows how to use AI as leverage without outsourcing her judgment to it.”

You can combine qualities eventually, but I want you to force-rank one primary belief first. That decision should govern the landing-page hierarchy, CTA copy, what the prototype proves, and how much of the supporting process we expose.

USER • TURN 3

I really want to express 1, 4 and 5 in that order.

ASSISTANT • TURN 4

Good — that gives us a clear hierarchy:

1. Primary: you can take an ambiguous clinical product problem from research → reasoning → prototype, independently and quickly.
2. Secondary: you're the kind of intense, hands-on builder Heidi seems to value.
3. Tertiary: you use AI as leverage without outsourcing judgment.

That means the page should not feel like a portfolio case study with three equally weighted artifacts. It should feel like one coherent proof of execution, with the research/process/AI pieces supporting the prototype rather than competing with it.

Question 2

What exactly is the prototype solving?

I want the sharpest possible articulation of the user/problem space before we touch structure or copy. Please answer in this format:

User:

Problem:

Why it matters:

What the prototype does:

What it deliberately does not do:

Be specific about the clinical-trials feature you're proposing. I'm particularly interested in whether you're solving for clinicians, patients, trial coordinators, or some combination — because if that's fuzzy, the whole project risks looking broad rather than incisive.

USER • TURN 5

User: Primary user and target is the research nurse / clinical research coordinator

Problem: Clinical trial teams experience high eligibility screen-failure rates because determining whether a patient meets complex trial criteria is difficult and relies on incomplete or manually interpreted clinical information. This creates wasted screening effort, delays enrolment, and risks suitable patients being missed.

Product hypothesis: If Heidi uses existing clinical information to identify potential trial matches and support eligibility screening, research teams will screen patients faster and with less manual effort, while more suitable patients progress towards enrolment.

Commercial hypothesis: If Heidi can demonstrate measurable improvements in screening efficiency and patient recruitment, research sites, sponsors and CROs will see sufficient value to pay for it as an additional clinical trials product.

What the Prototype does: The prototype takes the user through a workflow where recommendations are provided leveraging existing clinical information within their Heidi platform

What it deliberately does not do: real integrations with EHR/EMR, CTMS, EDC, eConsent, ePRO or registries; real patient data; automated patient outreach; clinician in-session alerts; full referral tracking; pricing/checkout; permissions management; notification preferences; mobile/tablet; production-grade compliance/audit tooling; and any true AI matching logic beyond the scripted dummy-data states.

ASSISTANT • TURN 6

This is strong enough to build from, but there's one important tension I want to resolve before we go further:

Your problem statement is about eligibility screening, while your prototype description says "recommendations are provided leveraging existing clinical information." That wording could imply either:

- Heidi is recommending trials for a patient, or
- Heidi is helping a coordinator determine eligibility for a specific trial.

Those are related, but they are different product problems with different workflows, risks, and success metrics.

Question 3

Which of these is the prototype primarily demonstrating?

A. Patient → trial matching

The coordinator starts with a patient and Heidi recommends potentially relevant trials.

B. Trial → patient screening

The coordinator starts with a trial and Heidi helps identify/screen potentially eligible patients.

C. Eligibility review for a patient already associated with a trial

The coordinator is already considering a specific patient for a specific trial, and Heidi helps interpret criteria against the patient record.

D. A combined workflow

For example: Heidi identifies a potential match, then guides the coordinator through eligibility review.

I want you to choose the single primary workflow even if the prototype includes pieces of another. My instinct is that D may be too broad for a two-day prototype unless the interaction is extremely disciplined, so I want to test that assumption with you.

USER • TURN 7

The prototype is already built and currently focuses on:

B. Trial → patient screening

The coordinator starts with a trial and Heidi helps identify/screen potentially eligible patients. They are then able to review the list including the recommended eligibility criteria from there.

Note: For each of the following questions, show an indication of how many questions are left

ASSISTANT • TURN 8

That sharpens the story considerably. Your prototype is not “AI for clinical trials” in the abstract; it demonstrates a specific operating loop:

Trial selected → candidate patients surfaced → eligibility evidence reviewed → coordinator decides what deserves further screening.

That’s useful because it lets the page communicate a very concrete product judgment: Heidi should reduce the amount of manual chart review required to get from a trial to a defensible screening shortlist, without pretending the system can make the final eligibility decision.

Question 4 of 12 — 8 questions after this

What did you define as “success” for this product, and what are the explicit guardrails?

I want the measures you actually used while designing it, not metrics we invent afterwards.

In particular, tell me what you currently believe for:

Primary outcome: What one measurable thing should improve if this works?

Secondary outcomes: What else should improve?

Guardrails: What must not get worse in pursuit of those outcomes?

Human decision boundary: What is Heidi allowed to recommend or infer, and what must remain the coordinator's responsibility?

I'm pushing on this because "faster screening" by itself is dangerously easy to optimise. A product could cut screening time by confidently excluding people it shouldn't, or create more work by surfacing a huge number of weak matches. Your definition of success is one of the clearest places to demonstrate the judgment behind the prototype rather than just the interface.

USER • TURN 9

For context, here are the considerations. In summary:

Primary outcome: What one measurable thing should improve if this works?

Improve the number of truly eligible patients in clinical trials.

Note, I noted that guard rails are very important for this.

Secondary outcomes: What else should improve?

- Reduce Time to Value - Reduce time it takes for clinical trial teams to verify and screen patients.
- Reduce rejection % of patients going through whilst increasing number of eligible applicants
- Auditability of the trail

Guardrails: What must not get worse in pursuit of those outcomes?

- Trust is the biggest - Evidence traceability, current system performance, usability, measured through repeated usage
- Workflow burden - are we creating more work and administrative effort? This will be measured by number of clicks and time taken to complete the task
- Accuracy - It is for naught if the system allows a large number of false positives through. There needs to be feedback of if the patient ends up being eligible for the final trial. It must also have human in the loop oversight at checkpoints to ensure accuracy.

Human decision boundary: What is Heidi allowed to recommend or infer, and what must remain the coordinator's responsibility? Heidi is allowed to recommend however it will need to provide reasoning, evidence traceability with links and the final review must sit with the coordinator.

ASSISTANT • TURN 10

That gives us a solid measurement hierarchy, but I want to tighten one thing before we move on.

Your primary product outcome — increasing the number of truly eligible patients progressing into trials — is directionally right, but it is a lagging outcome with several causes outside Heidi's control: trial capacity, patient

consent, investigator judgment, competing studies, recruitment windows, etc.

Your earlier prototype work already made a useful distinction: the product can ultimately target patient progression, while an early prototype should primarily validate task completion, comprehension and trust, with speed as a secondary measure. I think we should preserve that distinction in the portfolio story. Otherwise a reviewer may reasonably ask, “How could you possibly claim this prototype increases enrolment?”

So I'd frame the hierarchy roughly as:

North-star product outcome: More truly eligible patients are identified and progress toward enrolment.

Leading product signals: Screening quality/accuracy, fewer unsuitable patients progressed, fewer eligible patients missed, lower screening effort.

Prototype validation: Can a CRC understand the recommendation, verify its evidence, recognise uncertainty, and confidently make the appropriate human decision?

That actually strengthens your story because it demonstrates that you know the difference between what the business ultimately wants and what this prototype can honestly test.

Question 5 of 12 — 7 questions after this

The next thing I want to challenge is the “2 days” claim, because it is prominent enough that it could either become a major strength or accidentally undermine the work.

When you say:

“How I challenged this job listing and created this in 2 days”

What exactly happened within those two days?

I want the real boundary, including whether those two days covered research, synthesis, defining the problem/hypotheses, the extensive product interrogation in the chatlog, sketching/Procreate, prototyping, testing/refinement, and building this portfolio presentation.

And just as importantly: what existed before the clock started?

The goal here is not to diminish the speed. It's to make sure that when we eventually use “2 days” as part of the hook, it is impressive and defensible, rather than sounding like we're claiming two days of work for something that actually incorporates substantial earlier preparation.

What exactly happened within those two days?

I started this project at 8:30am (saturday) and it is 11:15pm now. In that time, I used Procreate to draw out my thought process from research, problem statement, and inbetween used ChatGPT for deep research, using it to prioritise and focus specifically on Eligibility (screen) failure. I then mapped out my thoughts around further problem definition on procreate, ideation and early engineering engagement to assess effort against value (combined score of company value and customer/user value) to assess potential prototypes. I then used ChatGPT (as per prompts before) to walk through and generate a prototype, noting that I haven't included where I further refined the prompts and flow in another chat until I was satisfied.

I'm including tomorrow since I want to leave time to create the portfolio page and I have a social family commitment. I intend to submit this linked in a scheduled email that will send at 8:56am Monday.

What existed before the clock started?

I started from scratch. I applied to the role yesterday and saw the post, wasn't able to work on this yesterday and started this morning.

ASSISTANT • TURN 12

That makes the "2 days" claim much stronger, because the clock genuinely starts from zero: application seen Friday, project started Saturday morning, then research → synthesis → prioritisation → product framing → ideation → effort/value assessment → prototype generation/refinement → portfolio packaging, with submission planned Monday morning.

I would still avoid wording that implies the entire clinical-trials problem was solved in two days. The stronger framing is that you took an ambiguous problem from zero to a reasoned, testable prototype in two days. That is both more credible and more aligned with the signal you want to send.

One nuance from your research supports that framing well: the evidence base itself is patchy. The research summary explicitly says Australian quantitative evidence on specific failure modes is limited, confidence is low, and several estimates are inference rather than proven fact. The broader research also identifies eligibility complexity and manual screening as a high-confidence problem area while noting important evidence gaps and the need for primary research. So the story should be: you made a disciplined decision under uncertainty, not that you discovered a definitive truth.

Question 6 of 12 — 6 questions after this

What do you want the recipient to encounter first after clicking "VIEW PROTOTYPE"?

There are three materially different choices:

A. Drop them straight into the prototype.

Fastest and most confident. It says "judge the work." Downside: they may miss the problem framing and not understand why specific decisions matter.

B. Give them a 15–30 second framing screen first.

For example: user, problem, hypothesis, prototype scope, "built in 2 days," then Enter prototype. Stronger narrative,

but adds friction.

C. Let “VIEW PROTOTYPE” open the prototype immediately, but include a very compact contextual strip/header inside it explaining what they are looking at and the assumptions. This keeps momentum while preserving context.

My current instinct is C, because your landing page already tells the story around the work. I don’t want the hiring manager to click the primary CTA and hit another “presentation about the presentation.”

Which do you prefer?

USER • TURN 13

Drop them straight into the prototype, there is a landing page built into the prototype

ASSISTANT • TURN 14

Agreed. Then the portfolio page should not create a second layer of explanation before the prototype. The primary CTA should behave like a direct handoff: click → prototype landing screen.

That also means the portfolio page's job is slightly different from what it first appeared to be. It does not need to explain the clinical-trials concept in detail. Its job is to establish credibility, signal the speed and depth of the process, and let the reviewer choose whether they want to inspect the evidence behind the prototype.

Question 7 of 12 — 5 questions after this

What do you want the three supporting cards to do strategically?

Right now they are:

- Data and Research
- Product Process
- AI Prompting

My concern is that this structure may accidentally make the page feel like three parallel case-study sections, when your intended hierarchy is:

prototype first → process proof second → AI usage third.

So I want you to choose the role of those cards:

A. Equal supporting evidence

All three matter equally and should have similar visual prominence.

B. Ranked supporting evidence

Product Process is strongest, Data and Research second, AI Prompting third.

C. Different purposes

For example:

- Data and Research = "Why I chose this problem"
- Product Process = "How I made the decisions"
- AI Prompting = "How I accelerated the work"

I think C is probably strongest, because it makes each card answer a distinct reviewer question instead of simply naming an artifact.

Do you agree, or do you want a different hierarchy?

USER • TURN 15

C, although I want to keep the hierarchy visually of them being on the same row

ASSISTANT • TURN 16

Agreed. Keep the three cards on the same row and at the same visual tier, but make their semantic roles distinct through copy rather than size.

That gives you a cleaner story:

- Data & Research → Why I chose this problem
- Product Process → How I made the decisions
- AI Prompting → How I accelerated the work

So visually they read as peers, while cognitively they answer three different questions.

Question 8 of 12 — 4 questions after this

How much of the research uncertainty and caveats do you want visible on the portfolio page itself?

Your source material is careful about this. The research summary repeatedly notes sparse Australian evidence, low confidence in several estimates, and areas where findings are inference rather than fact.

There are two viable approaches:

A. Keep caveats inside the research overlay only.

The landing page stays punchy and confident. The deeper evidence layer shows where the data is strong vs weak.

B. Signal uncertainty directly on the landing page.

For example, subtle wording like “Desk research under uncertainty” or “What the evidence did — and didn’t — support.”

My recommendation is A, because the landing page should sell the quality of your judgment, not foreground limitations. But inside the research overlay, I would make your evidence discipline very visible — especially the fact that you did not treat weak estimates as proven facts.

Do you agree?

USER • TURN 17

A

ASSISTANT • TURN 18

Agreed. We'll keep the landing page crisp and confident, and reserve the evidence-strength caveats for the Data & Research overlay where they can actually demonstrate rigor instead of weakening the first impression.

Question 9 of 12 — 3 questions after this

How much copy density are you comfortable with on the landing page?

Your reference aesthetic is extremely minimal, which means every extra line carries a lot of visual weight. We need to decide whether the page should feel more like a teaser or a mini case-study index.

My recommendation is:

- Name + Product Manager
- One large primary CTA card
- Three smaller supporting cards
- One subtle “About me” link
- No additional explanatory paragraph on the main page

That would preserve the original aesthetic and force the copy inside the cards to do the work.

Do you want to keep it that minimal, or are you open to one short framing sentence somewhere on the page?

USER • TURN 19

I want to keep it that minimal

ASSISTANT • TURN 20

Agreed. Then the landing page should stay almost ruthlessly minimal: identity → primary action → three supporting proof points → About me. No explanatory body copy.

Question 10 of 12 — 2 questions after this

For the three overlays, do you want each one to feel like a lightweight “peek behind the curtain,” or like a more substantial mini case-study?

Concretely:

A. Lightweight overlay

- 1 short headline
- 2–4 concise visual/content blocks
- one artifact embedded or linked
- one clear close/return action

B. Deeper overlay

- stronger narrative sequence
- more explanation of decisions
- multiple sections
- more like a compact case study without leaving the page

My recommendation is A. The primary story should remain the prototype, and the supporting material should reward curiosity without turning into homework.

Do you agree?

USER • TURN 21

A

ASSISTANT • TURN 22

Agreed. Each overlay should feel like a quick proof layer, not a second portfolio page.

Question 11 of 12 — 1 question after this

For the primary CTA copy, do you want the emphasis to be on:

A. Speed

VIEW PROTOTYPE

How I took an ambiguous product problem from zero to prototype in 2 days

B. Challenge

VIEW PROTOTYPE

How I challenged this job listing and built the answer in 2 days

C. Process

VIEW PROTOTYPE

From research to product decision to prototype in 2 days

D. Builder identity

VIEW PROTOTYPE

How I went from “interesting role” to working prototype in 2 days

My preference is C, because it directly reinforces your primary desired signal: research → reasoning → prototype, quickly and independently.

Which direction feels most like you?

USER • TURN 23

Is it possible to have C but some copy options? I want it to have a bit of personality but still be short and show how I took it on as a challenge to show initiative

ASSISTANT • TURN 24

Yes. I'd keep “VIEW PROTOTYPE” fixed and use the subtext to carry the personality, initiative, and two-day constraint.

A few directions:

- From job listing to working prototype in 2 days — simplest and punchiest.
- How I turned this job listing into a product challenge in 2 days — strongest on initiative.
- How I took the brief from research to prototype in 2 days — cleanest process story.
- I gave myself 2 days to turn this brief into a prototype — more personal and slightly playful.
- How I went from “I should apply” to a working prototype in 2 days — most personality, but less product-focused.
- How I challenged the brief, made the calls, and built the prototype in 2 days — communicates judgment as well as execution.

- Research. Decisions. Prototype. 2 days. — very compact and visually strong, though less explanatory.
- How far could I take this brief in 2 days? — strongest challenge framing; creates curiosity.
- I treated the job listing like a product brief. Here's what I built in 2 days. — probably the clearest expression of initiative.
- I didn't just apply. I gave myself 2 days to build the idea. — highest personality, but I'd be cautious because "the idea" is vague.

My top three for your objective are:

VIEW PROTOTYPE

I treated the job listing like a product brief. Here's what I built in 2 days.

This is probably my favourite. It explains why this exists, shows initiative, and doesn't imply you've solved clinical trials in a weekend.

VIEW PROTOTYPE

How I challenged the brief, made the calls, and built the prototype in 2 days

More product-manager-ish. "Made the calls" subtly communicates prioritisation and judgment.

VIEW PROTOTYPE

How far could I take this brief in 2 days?

Most characterful and intriguing. It turns the project into a self-imposed challenge, but says less about your process.

I'd avoid the original "How I challenged this job listing and created this in 2 days". "Created this" is vague, while "challenged this job listing" can read slightly as though you disagreed with the listing rather than challenged yourself with it.

Question 12 of 12 — final question

When the recipient clicks About me, what should happen?

Do you want it to:

A. Link straight back to the existing yasha.melbourne portfolio/About page, making this project a focused satellite of your broader portfolio.

B. Open a tiny overlay on this page with just enough personal context—perhaps 2–3 lines + link to full portfolio.

C. Navigate to a dedicated About section within this new subdomain.

My recommendation is A. You've deliberately kept this page focused on one signal, so I don't think we should rebuild your general portfolio inside it.

USER • TURN 25

I'd like to go with:

VIEW PROTOTYPE

I challenged myself. How far could I take this brief in 2 days?

For About me, I'd like it to have a similar overlay and I want to include how I love learning, growing and challenging myself with a photo of me rock climbing which will further reinforce the message

ASSISTANT • TURN 26

That works well. I'd lock the primary CTA as:

VIEW PROTOTYPE

I challenged myself. How far could I take this brief in 2 days?

It has more personality than the process-led versions, but still communicates initiative, constraint, and curiosity. Most importantly, "I challenged myself" removes the ambiguity of "challenged the job listing."

The About me overlay also fits the narrative, provided it stays brief. I'd make it feel like a human proof point rather than a biography: a rock-climbing photo, 2–3 lines on loving learning, growth, and deliberately seeking hard problems, then a subtle link to your full portfolio. The climbing image reinforces the same theme without needing to explain it too heavily.

With that, the interrogation phase is complete. We have enough decisions to agree on a plan.

The next step should be to turn this into a single implementation brief for the landing page and overlays, covering hierarchy, exact card structure, copy options, overlay content rules, spacing/behaviour, the hand-drawn arrows, and the constraints from your existing aesthetic.

USER • TURN 27

Would you be able to provide an implementation brief for another agent to build and host this in chatgpt. I would also like a chat log to provide additional context, makes it clear who said what and is built for readability