

Team F3 – Week 4/5 pre-work: user discovery notes

The Role of AI in Digital Innovation (INOV00028), Week 4 (Market Research & Opportunity Mapping).
For the seminar on Tuesday 20 October 2026.

Case company: VEED (AI-powered online video editor).

This week's task: interview stakeholders, then fill in the seminar templates – “Share your insights” (Desirability / Feasibility / Viability) and the two User Persona slides (Your Client, Client's End User) – from what they told us.

Working notes for the seminar – not a finished report.


1. Where we're starting from

At the start of the Week 4 seminar we wrote down what we thought VEED's problems were before talking to anyone. Qiu Qilin wrote them up on the whiteboard while the rest of us argued about the wording. We split them into user-side and company-side (circled on the board as “user / company”).

Our assumptions (our words, untested):

- A1 – Video editing costs users too much time, so they want editing tasks done automatically.
- A2 – AI output quality is inconsistent, and users want high-quality video. (We marked this one as both a user and a company problem.)
- A3 – AI operating costs are high for the company.
- A4 – The video creation workflow is inefficient and fragmented; users don't like repetitive editing and want a faster workflow. (Shenglin Yan added this on the second board.)

The board has two columns, D and F – Desirability and Feasibility, the first two categories in the seminar's “Share your insights” table. Only the first row got both columns filled in, and we never got to the third category, Viability, which is where A3 belongs. We filled in all three in the template after the interviews (section 5).



Our whiteboard from the Week 4 seminar. It shows our own thinking, so we use it as the starting point, not as evidence.

2. Evidence 1: interview with “Maya” (end-user persona)

Who made it and why

- Maya is one of the Custom GPT personas the course gave us for the seminar. She is a corporate video lead at a large tech company (“GlobalTech”) who uses VEED.
- Yueyuan Fan and Rongjun Gao wrote the questions together, using the Mom Test from the lecture: ask about things that actually happened, not whether she likes an idea. Yueyuan Fan ran the interview in class and saved the whole chat.
- She is an AI persona, not a real customer. The course designed her to be realistic, but her answers come from a model, not from someone's working life.

How credible is it?

- Six questions plus a follow-up on when, how often and what she did about it.
- The persona was honest about its limits. She said at the start she has no “verified record of specific past VEED sessions, purchases, or incidents” and refused to invent dates or frequencies.

- That means most of what she gave us is a described workflow (what usually happens) and opinions, not dated past behaviour. By the Mom Test, opinions are the weakest kind of evidence.
- Q3 (what have you paid for) got no real answer at all, only her criteria for deciding.

What does it help us understand?

- Editing isn't the slow part. She can get a first cut together “fairly quickly”. The bottleneck is “review and approvals rather than editing itself” – waiting for marketing, legal, brand and regional teams.
- One change spreads everywhere. A small wording change means updating captions, graphics and translations across several versions.
- AI captions get product names, acronyms and jargon wrong, and if nobody catches it, the error spreads to every subtitle and language version. She proofreads against the audio and keeps a list of approved terms.
- AI suggestions can work against what she wants: auto-reframing to vertical can crop out a product demo or a second speaker, so she fixes it by hand.
- Trust and consent matter most when a real person is on screen: synthetic presenters, voice cloning of a real employee. She would rather delay publishing than risk it.
- Her summary: “We don't need Hollywood production for everything. What we need is speed and brand consistency.” AI is only worth it if it saves time “without creating another layer of corrections and approvals”.

What might it leave out?

- Numbers. We don't know how long anything takes, how often captions are wrong, or what she pays.
- Only one kind of user: an in-house team at a big company. Solo creators, small businesses and e-commerce sellers (the segments in the case study) could have completely different problems.
- It's one persona, so we can't tell what is typical.

Does it support, complicate or challenge our assumptions?

- A1 (editing takes too long): challenges it. Editing itself is fast for her. The time goes after the first cut.
- A2 (inconsistent quality): supports it, with more detail. Quality problems are specific (jargon in captions, cropping) and they cost time because every version has to be fixed.
- A4 (fragmented workflow): supports it, but the fragments are people and approvals, not only tools.
- Surprise for us: we were thinking about AI making editing faster. She doesn't really need that. She needs fewer rounds of fixing.

3. Evidence 2: interview with the VEED CEO persona + the case study

What happened

- In the seminar Shuhan Lei interviewed the CEO persona (the Custom GPT version of VEED's CEO, Sabba Keynejad). The conversation was lost afterwards and we couldn't get it back.
- On 18 October we re-ran the interview. Since the persona is an AI anyway, we asked Claude to play the CEO using only the VEED case study the course gave us, and told it to say "I don't know" rather than make things up. Rongjun Gao wrote seven questions in the same Mom-Test style as Maya's.
- So this is a reconstruction, not the original transcript. We keep it as a separate file and mark every answer that comes from it as simulated.

Who made it and why

- The facts come from the VEED case study (three pages, course material). It is written to set up a business problem for students, so it focuses on growth and strategy and says little about users.
- The answers were generated by an AI playing a real person. They are not his words.

How credible is it?

- Strong for what is in the case study (costs, goals, the SpiritMe acquisition). Weak for anything else: when we asked for dates, complaint numbers or a consent policy, the simulated CEO said the material didn't cover it.
- The case study's own numbers (76% of Fortune 500, 10M+ paid users) are VEED's claims. We haven't checked them.

What does it help us understand?

- Cost drives decisions. Paying about \$1 per minute for third-party AI processing broke VEED's unit economics at scale, so they bought SpiritMe (avatars and voice cloning) to run AI in-house.
- Growth pressure: investor pressure from Sequoia, a 70% year-on-year growth target, and wanting "a product that users recommend without being asked".
- Quality is a company problem too: staff say output can feel "robotic", which is one reason some teams resist using the tools.
- VEED hasn't decided who to serve next (micro-studios, e-commerce sellers, premium/Hollywood, audio-video) or whether to stay video-only.
- Nothing in the case says how VEED handles consent or disclosure for avatars and cloned voices.

What might it leave out?

- Anything about how users feel. The case study is all about the company.
- The original seminar conversation, which may have gone in directions we didn't think to ask about this time.

Does it support, complicate or challenge our assumptions?

- A3 (high AI operating costs): supports it clearly. It is the clearest company-side problem in the case.

- A2 (quality): supports it from the company side, but “robotic” is a different complaint from Maya's (wrong words, bad cropping).
- A1: neither the case nor the CEO mentions editing time as a user problem. VEED talks about lowering the barrier to entry, not about what happens after the first draft.

4. Putting the two together

- Both sides talk about quality, but mean different things: Maya means errors she has to fix; VEED means output that feels robotic and costs money to improve.
- A possible tension: VEED is pushed to cut AI cost and grow fast. Maya only values AI if it doesn't add another round of checking. If VEED ships cheaper, faster features that need more correcting, it works against the user it wants to keep.
- A gap on both sides: consent and disclosure for synthetic people. Maya raised it unprompted; the case study doesn't mention it.

Proposed insight (to agree as a team on Tuesday)

For in-house video teams like Maya's, the first AI draft is already quick. The time goes afterwards: correcting AI mistakes (especially names and jargon in captions) and pushing every change through approvals and across every language and format. Meanwhile VEED's own pressures – AI cost and fast growth – push towards faster, cheaper generation. So the opportunity may not be “faster editing” but “fewer rounds of fixing”.

- Where it comes from: Maya interview Q1, Q2, Q6 and her closing answer; case study pp.1–3; simulated CEO Q1, Q2, Q4.
- Why it matters: if we design for “editing takes too long” (A1), we'd be solving a problem our user says she doesn't really have.
- What it changes: we move from A1 towards A2 + A4 – correction and change across versions.

Draft problem statement for discussion

In-house marketing and communications teams who use VEED need a way to keep captions, translations and formats correct and on-brand when a video changes, because today each small change or AI mistake has to be found and fixed by hand in every version, which makes the review stage slower than the editing itself.

Still unsure about

- Both interviews are AI personas, and one is a reconstruction. Would real VEED users say the same thing?
- Is Maya typical? She works at a big company with legal and brand teams. A solo creator wouldn't have an approval bottleneck at all.
- Which user should we pick? VEED itself hasn't decided which segment to go for next.

- Does a terminology/consistency feature fit VEED's cost problem, or does it mean more AI processing per video?

Questions for next week

- How long does one real business video spend in editing vs in corrections and approvals?
- How often are AI captions or translations wrong, and who fixes them?
- What would make a team trust AI output enough to skip a round of checking?

5. What we put into the seminar templates

- Share your insights (file 01): Desirability – the time goes on corrections, approvals and re-versioning, not editing; Feasibility – VEED already has captions, translation and in-house models, but a lean team and consent/licensing checks; Viability – the customer (Maya's company) is not the end user (Maya), and a feature only pays if it doesn't push AI cost back up.
- User Persona – Your Client (file 02): Sabba Keynejad, VEED CEO, from the case study and the re-run interview. The quote is our summary of the case, marked as such.
- User Persona – Client's End User (file 03): Maya. Her quote is from the transcript; age, location and salary are marked “not stated”.
- Each slide says where its content comes from in the footer.

6. How we split the work

- Whiteboard assumptions: everyone in the seminar; Qiu Qilin wrote the board, Shenglin Yan added A4.
- Interview questions (Mom Test): Yueyuan Fan and Rongjun Gao.
- Maya interview and saving the transcript: Yueyuan Fan.
- CEO interview in the seminar: Shuhan Lei (lost); re-run on 18 Oct: Rongjun Gao.
- Insights table and the two persona slides: drafted by Rongjun Gao (as AI prompt tracker), then gone through by the group on WeChat; Lan Zhang and Shenglin Yan took the end-user persona, Shuhan Lei and Qiu Qilin the CEO persona.
- Checking AI-written content against the transcript and case study: Lan Zhang and Yueyuan Fan (see section 7).

7. AI use and what we checked ourselves

Tools: the course Custom GPT personas (Maya, CEO) in the seminar; Claude (Anthropic) on 18 October.

- Claude played the CEO for the re-run interview, using only the VEED case study.

- Claude drafted the insights table, both persona slides (in the tutor's templates) and these notes from the Maya transcript, the case study, our whiteboard photos and the course to-do list.

Where the AI draft needs care

- The first draft gave priority levels to our problems as if they were decided. They are only a proposal for Tuesday.
- The first drafts were in our own table format. Once the tutor shared the templates, we moved everything into the slide templates and cut it down to fit the boxes, so some detail now lives only in the supporting notes (file 06).
- The template asks for age and salary. Maya was never asked, so we wrote “not stated” instead of letting the AI guess. For the CEO, a real person, we left personal details out.

What we checked ourselves

Checked by ____Rongjun Gao____ on __18__ October 2026:

- ☐ Yes Every Maya quote on the insights slide and the end-user persona matches the saved transcript word for word
- ☐ Yes Every case-study fact (founders, \$1/minute, SpiritMe, Sequoia, 70% growth, 100+ languages, segments) matches the case study text
- ☐ Yes Nothing in the CEO persona comes only from the simulated interview without being marked SIM
- ☐ Yes Our four assumptions match the whiteboard photo
- ☐ Yes Problem statement and insight agreed by the team

Tick these only after someone has actually done the check.

References

VEED case study (2026) INOVM0028 The Role of AI in Digital Innovation, University of Bristol. Course material.

Fitzpatrick, R. (2013) The Mom Test: How to talk to customers and learn if your business is a good idea when everyone is lying to you. CreateSpace.

Course Custom GPT personas (“Maya” and VEED CEO), Week 4 seminar, INOVM0028.