

Team F3 – Week 6 consolidation: problem statement, Innovation Goal and user needs → solutions

The Role of AI in Digital Innovation (INOV0028), Week 6 (Consolidation Week – no class). For the start of Week 7 (building solutions and prototyping).

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

This week's task, from the end of the Week 5 seminar: arrive at Week 7 with a defined Problem Statement, a defined Innovation Goal, and user needs mapped to potential solutions (the "Mapping User Needs to Solutions" homework). The seminar also asked us to answer the five Midpoint Reflection questions, and Activity 1 (affinity clusters and the persona templates) still needed finishing.

Slides that go with these notes: F3_028_W6_Innovation_Goal_and_Needs_Map.pptx (12 slides, in the tutor's templates).

Version 2 (22 October): updated with quotes and numbers from our interview transcript. Quotes are copied word for word and tagged with the question number (Maya Q1–5, Sabba Q1–3).

Working notes – not a finished report. Anything marked "?" still needs checking.

1. Where we're starting from

Our whiteboard is the version we go by. By the end of Week 5 it had:

- Three problem areas from Week 4, each with Desirability / Feasibility / Viability evidence: time-consuming video editing (Wistia 2026 State of Video Report, Descript), inconsistent AI output quality (NIST AI Risk Management Framework and Playbook) and high AI operating costs (VEED case study, competitor pricing as context).
- Two AI-simulated interviews with the course Custom GPT personas: Sabba Keynejad, the VEED CEO (3 questions), and Maya Chen, a Video Strategist who uses VEED (5 questions). Questions were written with the Mom Test – ask about things that happened, not opinions.
- Two personas built from those interviews: Sabba Keynejad (client) and Maya Chen (end user). These are final; we haven't changed them.
- A problem statement on the Who / What / Barriers / Cause / Impact template, and a consolidated version of it.

The consolidated problem statement we agreed on the whiteboard (unchanged):

"Marketing professionals using VEED need to create high-quality videos efficiently and obtain timely approval from multiple stakeholders. However, feedback scattered across different communication channels, conflicting revision requests and unclear approval responsibilities can make the review process difficult to manage. This leads to repeated work, additional coordination and unpredictable

delays. In our AI-simulated interview, Maya estimated that these issues affected approximately 60–70% of her projects, with one example involving 3–4 extra hours of work and an additional business day of waiting.”

The turn we already made

- We went into Week 4 assuming the user's main problem was that editing takes too long (A1). The Wistia evidence pointed that way too.
- Maya's answers pointed elsewhere. Her first cut took about an hour (Q1); the time went on collecting feedback, sorting out conflicting requests and waiting for approval. In her words: “Editing is quick; getting five people to agree on the final video is difficult.” (Q2)
- We disagreed about this for a while. Each of us wrote a version, we merged them, and we chose fragmented feedback and approval as the core problem. Editing speed stays on the list as a secondary issue.

The transcript

- The full transcript is now in our folder (028 / 3_我们的作业 / W4-5_访谈记录 / VEED_Simulated_Interview.docx). The quotes on the slides and in these notes come from it.
- Verbatim evidence for the problem statement: barrier – “Editing is quick; getting five people to agree on the final video is difficult.” (Maya Q2); cause – “VEED makes content creation faster, but organisational coordination hasn't caught up.” (Q2); impact – “A project that could realistically have taken one working day ended up closer to two.” (Q3)
- The numbers in the problem statement (60–70%, 3–4 hours, one extra business day) are in Maya Q3. She says herself “These are my estimates.” (Q3) They come from one AI persona, not measured data, and we say so wherever we use them.

2. Affinity clusters and personas (Week 5 Activity 1)

Activity 1 asks us to tag quotes as goal, frustration, behaviour, workaround, trigger or metric, cluster the recurring ones, attach 2–3 verbatim quotes to each cluster and put the clusters linked to measurable outcomes first. We did this for Maya's five answers.

Cluster (tags)	Verbatim quotes	Measurable outcome
1. Approvals, not editing, eat the time (frustration, metric)	“I spend more time chasing approvals than actually creating.” (Q1) · “Editing is quick; getting five people to agree on the final video is difficult.” (Q2) · “I'd rather save three hours coordinating feedback than another fifteen minutes trimming a video.” (Q5)	First cut ≈ 1 h, then +3–4 h hands-on and +1 business day in one example (Q3)
2. Feedback scattered and out of date (trigger, behaviour, frustration)	“I shared a VEED review link, but feedback came through multiple channels.” (Q2) · “A person could review the first cut and reply two days later after the second version was already updated, producing comments about scenes that no longer existed.” (Q2)	≈ 1 h consolidating conflicting feedback in one project (Q3)
3. Home-made workarounds to keep control (workaround, behaviour, goal)	“I consolidated everything in a tracker with timestamps, requested change,	She predicts a shared feedback location, clear decision authority and

	requester and status, and distinguished factual errors from subjective preferences." (Q2) · "I established one contact in each stakeholder group, sent a revised video with a feedback deadline and requested consolidated comments." (Q2)	an approval cutoff "could have removed at least half the extra work" (Q3 – her prediction, not an observed saving)
4. The same AI caption fixes, again (frustration, workaround, metric)	"The main frustration was repeating the same corrections across projects." (Q4) · "At present I often transfer their corrections myself." (Q4)	20–30 min fixing English captions; 1.5–2.25 h in total with translation (Q4). She ranks this below approvals (Q5)

Clusters 1–3 are the ones behind our problem statement, and they all tie to time and delay we can measure. Cluster 4 is real but secondary – Maya says approvals have "the biggest impact" (Q5).

Personas on the tutor's templates

- We rebuilt both persona templates (Client and Client's End User, 10 boxes each) on slides 3–4. The content is the whiteboard persona text, with transcript evidence added and a question number on each point.
- Quote boxes now hold real quotes. Sabba: "The ultimate challenge is keeping VEED the easiest way to turn ideas into compelling video stories." (Q1). Maya: "I spend more time chasing approvals than actually creating." (Q1)
- Age, location and salary are "?" for both. None of the interviews or the case study gives them, and we don't want to make them up. Sabba is a real person, so we'd leave personal details out anyway.
- Sabba's \$500,000 and \$1 million monthly API figures (Q1, Q2) are his own hypothetical examples, not VEED's actual spending. We only use them labelled that way.

3. From the problem statement to the Innovation Goal

Reading the problem statement again

The statement has three barriers and one root cause. Before writing the goal we asked which of them AI can actually help with:

- Feedback scattered across email, Slack and other channels → a good fit. Collecting, summarising and grouping text comments are standard AI tasks (summarisation, clustering).
- Conflicting revision requests → partly. AI can spot that two comments pull in opposite directions (conflict detection). It should not decide which one wins. Maya resolves conflicts "based on audience needs rather than seniority" (Q2) – that's a judgement between colleagues and it belongs to people.
- Unclear approval responsibilities → mostly not an AI problem. Who approves is a workflow and people question. AI can only help around the edges (e.g. sorting comments into must-fix / should-fix / nice-to-have, which Maya already does by hand, Q2).
- Root cause (no coordinated review process) → AI alone won't fix this. Maya says she can't solve stakeholder responsiveness on her own; it needs "clearer processes, ownership and tools" (Q5). Whatever we build has to give the review a shape: one place, named approvers, clear status.

That is why the goal says "consolidate and flag", not "approve" or "decide".

The goal (W5 lecture formula: “Harness AI to [purpose] for [persona] so they can [benefit]”)

Our project aims to harness AI to consolidate scattered review feedback and flag conflicting requests for marketing professionals like Maya Chen who need several stakeholders to approve their VEED videos so that they can reach a clear, human-approved final version with fewer revision rounds and less waiting.

- Purpose (a recognisable AI task): pull stakeholder comments into one list pinned to the video timeline; summarise and group them by issue; flag requests that contradict each other; show who still has to approve.
- Persona (someone we have evidence from): Maya Chen. She coordinates product marketing, internal communications and brand reviewers (Q2) and consolidates their comments from email, Slack and other channels.
- Benefit (observable and measurable): fewer repeated revisions, less coordination work, more predictable approval. It also links to her goal: “I want my team to focus on storytelling and strategy rather than resizing, subtitles and chasing minor edits.” (Q1)

Other wordings on the table (to pick at the Week 7 meeting)

- Shorter: “Harness AI to summarise and cross-check stakeholder feedback for Maya Chen so she can get sign-off with fewer revision rounds.” Easier to say out loud; loses the “human-approved” part, which we think matters.
- Dropped: “Harness AI to speed up video editing for marketers ...” – this is our old assumption A1. Maya: “Faster editing delivers diminishing returns if videos still get stuck in approvals.” (Q5)
- Dropped: “Harness AI to manage approvals for ...” – too vague, and it sounds as if AI does the approving, which is exactly what we don't want.

Where the client (Sabba) comes in

The goal is built around Maya, but it has to make sense to the client too. From the whiteboard Sabba persona, his interview and the case study:

- Who to serve: asked where VEED should place its next bet, he said he'd “lean towards small businesses and lean creative teams producing substantial volumes of marketing video” (Q3). He calls this “a strategic recommendation, not a confirmed VEED decision” (Q3), so we use it as support, not as VEED's plan. Maya's team fits that description.
- What matters most: “Owning technology is useful, but owning the customer experience is more important” (Q1). Our goal is about the customer's experience of getting a video signed off, not about a new model.
- AI cost: the case study's old third-party API cost of about \$1 per minute of video made VEED's unit economics unsustainable at scale. Sabba adds that external APIs “can be better for commoditised features” (Q2), while owned models need more upfront investment but can lower marginal cost at scale (Q2). Summarising and clustering comments is a fairly commoditised text task, so an external API might be fine to start with, and own models only make sense once volume is high (?).
- Focus: he notes that engineering time spent on avatars “cannot simultaneously improve editing, collaboration or reliability” (Q3). A review feature inside the editor improves collaboration and fits the focused route better than a separate tool.

- Case study: Sabba's success measures include “a product that users recommend without being asked”. Less painful approval rounds is the kind of thing a team would mention to others – but we have no evidence for that yet (?).

Success metrics (baseline → target)

The lecture says to pair the goal with metrics and guardrails. Baselines are Maya's estimates from the transcript; targets are our proposal.

Metric	Today (Maya's estimates)	Target (our proposal)	How we'd check it
Projects hit by approval delays or rework	3–4 of her last 5–6 projects had unexpected revisions or approval delays; ≈ 60–70% overall (Q3)	≤ 35% (halve it) *	Follow-up interviews; later, count projects in a pilot
Extra hands-on work per affected project	3–4 h = 1 h consolidating conflicting feedback + 1–1.5 h revising script and video + 30–45 min replacing a screenshot + 30–60 min another approval round (Q3)	≤ 2 h (halve it) *	Time log on a test project in the Week 9 prototype test
Extra waiting for final sign-off	≈ 1 business day; a one-day project “ended up closer to two” (Q3)	0 extra days	Time from “ready for review” to final approval
Secondary: fixing AI captions and translations	20–30 min English captions; 1.5–2.25 h in total with translation, plus waiting for regional review (Q4)	No repeat fix of a glossary term	Count repeated term corrections across projects
Accuracy of AI summaries	No baseline – new feature	≥ 9 in 10 points judged correct	User checks each summary point against the linked original comment

* Why “halve it”: Maya predicts that one shared feedback location, clear decision authority and a formal approval cutoff “could have removed at least half the extra work” (Q3). She calls it “a prediction rather than an observed saving”, so it supports our targets but doesn't prove them. In version 1 the hands-on target was ≤ 1 hour; we changed it to ≤ 2 hours so it matches her “at least half” instead of going beyond it.

Guardrails (red lines)

- 0 AI approvals: AI never marks a video as approved and never chooses between conflicting requests. The named approver decides and signs off.
- No summary without its source: every summarised point links to the original comment and its author, so a wrong summary (hallucination) can be caught.
- Client footage stays private: footage and comments stay in VEED's own pipeline; they are not used to train models or sent to third-party AI without the client's consent.

4. User needs → potential solutions

We took Maya's most urgent needs from the whiteboard persona and the problem statement, added the sub-needs she names in the transcript, then followed the W5 lecture steps: need → AI task type →

feasible method → what success looks like, plus risks (bias, hallucination, privacy) and accessibility. Approvals come first because Maya says “Managing stakeholder feedback and approvals has had the biggest impact, without question.” (Q5)

Need (priority)	Where it comes from	Feature	AI task type	Feasibility for VEED	Main risk → how we limit it
1. Clear sign-off: who approves, by when (top)	Pain point “unclear final approval responsibilities”; Q5 ranks approvals first; by hand she sets one contact per stakeholder group and a feedback deadline (Q2)	Approval Tracker – named sign-off owner and one contact per group, a formal approval cutoff, status per version; AI pre-sorts comments into must-fix / should-fix / nice-to-have for Maya to confirm	Classification	Extends the VEED review link Maya already uses (Q2). Mostly workflow, little AI – cheapest piece for a lean team.	A must-fix tagged as nice-to-have → Maya confirms tags; only named approvers can approve
2. All feedback in one place, on one version (urgent)	“I shared a VEED review link, but feedback came through multiple channels.” (Q2); persona pain point “fragmented feedback”	Feedback Inbox – email and Slack comments pulled into one list next to VEED's review comments, pinned to the same version's timeline; AI summarises and groups by issue	Summarisation; clustering; matching comments to timecodes	The review link exists but people bypass it, so the real work is bringing email/Slack in. Captions give timecoded text to pin comments to. Summarising is a commodity task – an external API may be cheaper at first (Sabba Q2); own models later at scale (?).	Summary drops or changes what someone meant → each point links to the original comment
3. Catch conflicting and vague requests early (urgent)	Longer demo vs shorter video took ≈ 1 h to settle (Q3); feedback like “make this more engaging” was vague and needed follow-up (Q2)	Conflict Flag – shows two opposite requests on the same part of the video side by side and marks vague comments that need a follow-up question; a person decides	Conflict (contradiction) detection; vague-comment detection	Hardest AI task here. Start narrow: same timestamp plus opposite instructions. Needs real comment data to test (?).	False alarms annoy; missed conflicts give false confidence → AI only flags, never picks a side
4. No comments on outdated versions (next)	Comments about scenes that “no longer existed” (Q2); she wants “version tracking that flags outdated comments” (Q2)	Version Diff – compares two versions' transcripts and captions, flags comments made on an outdated version, lists what is addressed and what is open	Version diff; semantic matching of comments to changes	Text comparison on caption data VEED already produces; low compute cost.	Marks a comment “done” when it isn't → shown as “likely addressed”, user confirms
5. Stop repeating caption and translation fixes (secondary)	“The main frustration was repeating the same corrections across projects.” (Q4); 20–30 min captions, 1.5–2.25 h with translation (Q4); outside the problem statement	Brand Glossary – a reusable approved glossary applied to captions and translations; regional reviewers edit translated captions without access to the whole project (Q4)	Entity recognition; custom vocabulary for captions and translation	Builds on VEED's flagship auto-captioning in 100+ languages (case). Needs a reviewer-only access level (?).	One wrong fix spreads everywhere; uneven quality across languages → regional reviewer stays in the loop

Why this order

- In version 1 the Feedback Inbox came first. After reading Q5 we moved clear sign-off to the top: Maya says approvals affect “timelines, capacity and output volume” more than anything else (Q5), and her own fixes (one contact per group, a feedback deadline) are about sign-off.
- Needs 2 and 3 are the other two barriers in our problem statement. Need 3 depends on need 2: you can't flag conflicts in comments you haven't collected.
- Need 4 matters but depends on 1–3 working first.
- Need 5 is real (Q4, and it's on the persona) but it isn't part of the problem statement, and Maya ranks it lower (Q5). We keep it as a secondary idea because it uses VEED's strongest existing feature.
- For a first prototype in Week 7 we'd suggest needs 1–3 together (Approval Tracker + Feedback Inbox + Conflict Flag) – to agree.

Options we considered and dropped (keep these – they're our decision trail)

- Faster AI auto-editing. Decided in Week 5: this was assumption A1. Maya: “Faster editing delivers diminishing returns if videos still get stuck in approvals.” (Q5)
- AI that approves the video or picks the “winning” comment. Proposed for dropping: it takes the decision and the accountability away from people, and it breaks our first red line.
- AI that merges conflicting comments into one instruction for the editor. Proposed for dropping: it hides the disagreement instead of showing it, so the people who disagree never sort it out.
- Scoring reviewers (e.g. a “slow reviewer” or “difficult reviewer” rating). Proposed for dropping: it turns a workflow tool into monitoring of colleagues, with a real risk of bias. Maya never asked for it.
- Auto-sending reminders or replies to stakeholders. Proposed for dropping in its automatic form: messages to colleagues should come from Maya. AI could draft a reminder for her to send (?).
- Avatar or voice-clone features from SpiritMe. Proposed for dropping for this problem: they're VEED's strongest in-house AI, but they don't touch the review problem and they add consent risk for real people on screen.
- A separate project-management app. Proposed for dropping: Maya wants review to happen “within the editing environment rather than across VEED, Slack, email and spreadsheets” (Q2), and a new app pulls VEED towards the broader creative-suite route the case study warns can dilute focus.
- Fully automatic translation with no regional reviewer (for need 5). Proposed for dropping: Maya wants a native speaker to review terminology and tone (Q1), and errors would spread to every version.

We also want to keep one non-AI option in view: one feedback channel, a named approver per project and a feedback deadline. Maya already does a version of this by hand (Q2), and she still loses hours. When we test, we should check that our features beat her current workaround, not just that they work.

5. Risks and ethics

Linked to our Week 2–3 ethics table (10 concerns) where it fits:

- Hallucination: a summary can leave out or change what a reviewer said, and Maya might act on it. → Every summary point links to the original comment; the accuracy metric checks it.
- Over-reliance (ethics table row 9): if the summaries look good, people may stop reading the original comments. → Originals are always one click away; the Week 9 test should look at whether testers still check them.

- Bias and accessibility (row 3): sorting comments into must-fix or nice-to-have could treat short, indirect or non-native English comments as less important. Comments may arrive in several languages. Summaries need to work with screen readers. → Maya confirms the tags; test with mixed-language comments (?).
- Privacy (row 1): review comments and unreleased footage are sensitive client material, and linking Slack or email gives the tool access to more than video comments. Regional reviewers need access to translated captions only, not the whole project (Q4). → Client footage stays private (red line 3); only messages the user chooses get pulled in.
- Accountability (row 7): if AI tags or flags something wrongly and a video goes out with an error, who is responsible? → Approval stays with named people; the tool records who approved what.
- Workplace monitoring: an approval tracker can easily become a way to watch who is slow. → No reviewer scores (dropped above); status only, no rankings.
- Growth pressure and cost (row 10 and Sabba's cost problem): pressure to ship fast can cut testing; extra AI calls raise the cost per video. → Start with the cheap parts (tracker, version diff), and check the summarisation cost before building it (?).

6. What we still need to validate

- Do real marketing professionals who use VEED have the same problem? Everything so far comes from AI personas (?).
- Are Maya's numbers anywhere near typical – 3–4 of 5–6 projects, 3–4 hours, one extra day (?).
- VEED already has a review link (Maya Q2). Why did reviewers still use email and Slack – habit, access, or something the link can't do (?).
- How many reviewers and channels does a typical business video go through? Maya had three groups and talks of “five people” (Q2) (?).
- Does VEED have its own text/language model, or would summarisation run on a third-party API (?) – and what would that cost per video (?).
- Would Maya trust an AI summary enough to act on it, or would she re-read everything anyway (?).
- Would clients let a tool read their Slack and email (?).
- How often do conflicts actually happen, and can they be spotted from the text of the comments alone (?).
- Does Maya's “at least half” prediction (Q3) hold when we test a prototype (?).

Questions to take back to the personas in Week 7

- Maya: “You shared a VEED review link last time. Who used it, who didn't, and what did they do instead?”
- Maya: “Last time you waited for a sign-off, who were you waiting for, and how did you find out?”
- Maya: “When you sorted comments into must-fix, should-fix and nice-to-have, how long did that take, and did anyone disagree with your sorting?”
- Sabba: “When you last decided whether a feature should run on your own models or an outside API, what decided it?”

7. Midpoint Reflection (Week 5 seminar)

- What have we learnt from user discovery so far? Our starting assumption (editing takes too long) didn't hold up. Maya's first cut took about an hour (Q1); the time went on review: scattered feedback, conflicting requests and unclear sign-off.
- What initial problem(s) have we identified? Fragmented feedback and approval for marketing teams using VEED; second, repeated fixes to AI terminology and translations.
- Why is this a high priority problem? 3–4 of Maya's last 5–6 projects were hit (Q3), and she ranks it first (Q5). It fits Sabha's focus on customer experience (Q1) and his suggested next bet of lean teams making marketing video (Q3 – a recommendation, not a VEED decision).
- What else do we need to discover about our users? Whether real marketers say the same; why reviewers bypass VEED's review link for email and Slack; how many reviewers are typical; whether users would trust AI summaries (all ?).
- What's the next step? Week 7: agree the goal and targets, sketch a low-fidelity Approval Tracker + Feedback Inbox + Conflict Flag, and test it back with both personas.

8. How we split the work (Week 6–7, proposed – confirm in the group chat)

Roles follow our Team Charter. A name means that person leads that piece; we still do most of it together.

- Qiu Qilin (Team Coordinator): run the Week 7 kick-off; get the team to agree the goal wording, the targets and the “proposed” dropped options; update the whiteboard Work Log for Weeks 6–7.
- Yueyuan Fan (Research / Evidence): check every quote and number on the slides against the transcript; find out how VEED's review link works today and why reviewers might bypass it (?).
- Lan Zhang (Visual / Slide Lead): put the clusters, personas, goal and needs → solutions map onto the whiteboard and into the deck style; set the look for the prototype screens.
- Shenglin Yan (Communication Lead): take the goal and the map to the tutor in Week 7 for feedback; ask the follow-up questions to the Sabha persona about cost and fit.
- Rongjun Gao (AI Prompt Tracker): log this week's AI use (both rounds) in the 028 Prompt Log, including what we accepted, changed and rejected; keep the AI section of these notes up to date.
- Shuhan Lei (Prototype Designer): sketch the low-fidelity prototype for needs 1–3; draft the Week 9 test plan using the metrics in section 3.

9. AI use and what we checked ourselves

Tool: Claude (Anthropic), used by Rongjun Gao on 22 October 2026, in two rounds.

- Round 1 – what we gave it: our whiteboard text (D/F/V evidence, interview questions, personas, problem statement), the Week 5 lecture and seminar slides, the good/bad problem statement examples, the VEED case study and our Week 2–3 ethics table. What we asked (gist): turn our agreed problem statement into an Innovation Goal using the lecture formula, with metrics and guardrails; map Maya's needs to features, AI task types, feasibility and risks; answer the Midpoint Reflection; use only evidence in our folder and don't invent quotes or numbers. Output: 9 slides and the first draft of these notes, with highlighted slots where quotes were missing.
- Round 2 – what we added: our interview transcript (Sabba Q1–3, Maya Q1–5). What we asked (gist): fill the quote slots with exact quotes and question numbers; use the transcript numbers on the

metrics slide; add the sub-needs Maya names to the needs map; use Sabba's answers on cost and focus; add the Activity 1 affinity clusters and the two persona templates; check everything against the Week 5 seminar.

- Output of round 2: 12 slides (new: affinity clusters, two persona templates) and this version of the notes.

Where the AI draft needs care

- Every quote was checked automatically against the transcript text: each one appears there character for character. Labels like “video editing takes too long” (our own assumption A1) are our words, not quotes from the interviews.
- The targets ($\leq 35\%$, ≤ 2 hours, 0 extra days, 9 in 10) were proposed in the draft. Two of them are now backed by Maya's “at least half” prediction (Q3), but that is still a prediction, and the team hasn't agreed the targets yet.
- Round 2 moved Approval Tracker to the top of the needs map and changed the hands-on target from ≤ 1 hour to ≤ 2 hours. Both changes follow the transcript; the team should confirm them.
- Most of the dropped options (all except A1) were suggested in the draft. They are marked “proposed” until we agree them.
- Feasibility notes use only the case study and the transcript. Anything about VEED's text models or costs is marked “?”.

What we checked ourselves

Checked by _____ on ____ October 2026:

- ☐ The problem statement on slides 5–6 and in section 1 matches the whiteboard word for word
- ☐ Every quote matches the transcript word for word, with the right question number
- ☐ Every number (3–4 of 5–6 projects, 60–70%, 3–4 h and its breakdown, one extra day, 20–30 min, 1.5–2.25 h) matches Maya Q3–Q4
- ☐ Sabba's \$500k / \$1M figures are only used as his hypothetical examples
- ☐ Persona slides match the whiteboard personas; nothing added without a source
- ☐ Every case-study fact (\$1/minute, SpiritMe, auto-captions in 100+ languages, 70% growth, lean team) matches the case study text
- ☐ Goal wording, targets and red lines agreed by the team
- ☐ Dropped options agreed, or moved back onto the list
- ☐ Each person has confirmed their line in section 8
- ☐ This week's AI use (both rounds) added to the 028 Prompt Log

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.

INOVM0028 Week 5 lecture and seminar slides (2026): Customer Discovery – persona, problem statement and Innovation Goal templates; Consolidation Week homework; Midpoint Reflection. University of Bristol.

INOVM0028 (2026) Examples of good and bad problem statements. University of Bristol. Course material.

Course Custom GPT personas (Sabba Keynejad, VEED CEO; Maya Chen, Video Strategist), Week 4 seminar, INOVM0028. Transcript: VEED Simulated Stakeholder Interview (team file).

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.