

Product Sense Case Study

AI Career Tools Suite

Profile Optimization Tool · Job Application Workflow Agent

Author	Ankit Singh, Operations Leader and AI Product Manager
Date	June 2026
Research	350+ career coaching sessions · Personal job-search experience · Direct user observation
Live at	ankitsingh.net/ai-prototypes

About this document

This document walks through two AI tools I built and shipped: the Profile Optimization Tool and the Job Application Workflow Agent. Rather than just describing what they do, it shows the product thinking behind them: how I scoped the problem, picked the right user, mapped their journey, prioritized what to build, and defined what success looks like. Both tools grew directly from 350+ career coaching sessions and my own active job search.

STEP 1

Scope the Problem

It started with a deliberately vague observation: *job seekers are struggling with their applications*. That's too broad to build anything meaningful, so the first step was to get specific.

The clarifying questions

What kind of struggling? Not finding roles. Plenty of job boards already handle that. The real pain is in *converting* applications: getting past the first filter and into an actual conversation with a hiring team.

Where in the process? Two distinct moments come up: before applying, when someone wants to know if their resume and LinkedIn profile are ready for a specific role, and while applying, when they need to tailor their materials to stand out against the JD.

What channel? Desktop web. Job seekers prep materials sitting at a desk, not on their phone.

What's the core mechanism? Resumes and LinkedIn profiles get evaluated against a job description. The gap between what someone has and what the JD asks for is the problem worth solving.

Refined problem statement

The real problem

Job seekers don't know how well their profile matches a specific role, and they don't have a fast way to close those gaps before they apply.

This costs them in two ways. They waste time applying to roles they were never realistically a fit for. And they miss out on roles they could have gotten, because their materials didn't make the case clearly enough.

What this problem is NOT

Out of scope	Why
Helping people find jobs or discover roles	Job boards like LinkedIn and Indeed already solve discovery
Interview preparation	A separate problem with different tools and a different point in the journey
Salary negotiation or offer comparison	Happens after the application, outside the scope of these tools
Auto-applying to jobs on behalf of users	Takes human judgment out of the equation, and the accuracy and trust risks are too high

STEP 2

Segment the Users

Not everyone in a job search has the same problem. Breaking the audience into real groups reveals where the pain is sharpest and where a tool can actually make a difference.

Segment	Who they are	Application behavior	Core challenge
Early-career (0 to 3 yrs exp)	Recent grads, first or second job	High volume, 20 to 50 apps per month	Generic resumes with little to differentiate; don't know which keywords matter to recruiters or ATS systems
Mid-career (3 to 10 yrs exp)	Experienced professionals, often pivoting roles or industries	Selective, 5 to 15 apps per month	Rich experience but struggle to translate it for a new role; bullets describe tasks instead of outcomes
Senior and career changers (10+ yrs or domain switch)	Directors, VPs, or professionals making a major pivot	Very selective, 2 to 8 apps per month	Hard to condense years of scope into a clean narrative; risk of appearing over- or under-qualified; major reframing required

Segments based on behavioral patterns observed across 350+ career coaching sessions.

STEP 3

Pick the Target Segment

Target: Mid-career professionals with 3 to 10 years of experience

Specifically those who are pivoting into a new role or industry. Think an operations manager moving into product management, or a finance analyst transitioning into strategy. They have real experience but it's framed for where they've been, not where they're going.

Why this segment?

Reason	Detail
Highest pain-to-value ratio	The experience already exists. It just needs to be reframed. AI is very good at that kind of translation, and it doesn't need to invent anything to do it well.

Reason	Detail
Most common case in coaching	Across 350+ sessions, this group consistently came up the most when it came to tailoring help and bullet rewrites.
They care enough to invest time per application	Unlike early-career applicants who send hundreds of applications, this group is selective and wants a real return on the time they spend tailoring.
Outcomes are measurable	We can track fit scores before and after tailoring, and see whether guided steps improve completion rates over time.

Staying focused on this segment

Every design decision in both tools stays anchored to this one user: someone with relevant but untranslated experience, applying to 5 to 15 roles a month, who knows the job they want but doesn't know how to make their materials say it clearly.

STEP 4

Build a Persona and Map the Journey

Meet Priya, the frustrated career pivoter

Attribute	Detail
Name and age	Priya, 31
Background	Five years in operations at a SaaS company. She has managed vendor workflows, cross-functional projects, and tools like JIRA and Notion.
Goal	Move into a Product Manager role at a tech company.
Behavior	Applies to 8 to 12 PM roles per month. Spends 30 to 45 minutes tailoring each application. Sends follow-ups. Rarely hears back.
In her own words	"I know I have the skills. I've done PM work. But I have no idea why I'm not getting calls."
Tools she uses today	Google Docs for her resume. She reads the JD and tries to edit manually. Sometimes she asks ChatGPT for help, but the results are hit or miss.
Research grounding	Directly representative of patterns seen across 350+ coaching sessions, and validated through personal job-search experience.

Priya's journey from open to close

This maps what Priya actually does at each stage, from spotting a role to waiting for a response. Each stage surfaces a real pain point, and severity tells us where to focus first.

Stage	What Priya does	Pain point	Severity
1. Discover role	Finds a PM role on LinkedIn. Reads the JD. Thinks it looks like a fit.	No quick way to validate if she is actually a fit before she commits time to tailoring	High
2. Assess fit	Reads her own resume and compares it to the JD manually.	The comparison is slow and subjective. She is not confident about which keywords actually matter to ATS filters.	High

Stage	What Priya does	Pain point	Severity
3. Tailor resume	Opens Google Docs and tries to rewrite bullets to match the JD.	Her bullets describe tasks, not outcomes. She doesn't know how to reframe ops experience as PM experience. Often she gives up and sends the same resume.	High
4. Update LinkedIn	Realizes her LinkedIn still says Operations Lead. She doesn't know what to change.	Her profile is static and generic, not positioned for the roles she is actually targeting.	Moderate
5. Write cover letter	Starts from scratch. Writes a generic opening paragraph.	There is no real connection between her cover letter and the specific JD. After ten applications it starts to feel pointless.	Moderate
6. Submit	Hits submit. Takes a screenshot. Logs it in a spreadsheet.	No confidence she has put her best version forward	Low (action already taken)
7. Wait	Checks email. No response after two weeks.	No feedback loop. She has no way of knowing what went wrong or how to do better next time.	Low (outside our control)

STEP 5

Prioritize the Pain Points

Not all pain points are worth solving. The ones that rise to the top share three traits: they come up often, they cost Priya real time or confidence, and they are something AI can handle reliably without making things up.

Pain point	Priority	Frequency	AI suitability	Why it ranks here
Can't quickly tell if she's a fit before investing time in tailoring	HIGH	Every application	High	Pattern-matching between resume and JD keywords is exactly what LLMs do well. Answering this question saves the most time per application.
Can't rewrite bullets to show outcomes in JD language	HIGH	Every application	High	The experience already exists, it just needs reframing. AI can do this cleanly because it's grounded in what the user already wrote, so there's no risk of hallucination.
LinkedIn profile isn't positioned for the target role	MODE RATE	Once per campaign	High	Important, but it only needs to happen once per job search campaign. A single thorough audit goes a long way.
Cover letter starts from scratch every time	MODE RATE	Every application	High	AI can draft this well, but it works best after the fit analysis and bullet rewrites are done. It's downstream of the core pain.
No feedback loop after rejection	LOW	Post-application	Low	Completely outside our control. We have no visibility into recruiter or ATS decisions.

Top pains to solve

- 1. Instant fit assessment.** Give Priya a clear, grounded answer to "Am I a realistic fit for this role?" in under 60 seconds, before she invests any time in tailoring.
- 2. Bullet rewriting.** Take her existing experience bullets and reframe them in the language of the JD, with outcomes front and center, so she is not starting from a blank page.

3. LinkedIn audit. Scan her public profile against her target role and give her specific, section-level fixes. This sits in scope as a complementary tool, not the primary unlock.

STEP 6

List the Solutions (Features)

Once the pain points are prioritized, the features follow naturally. Every solution on this list ties back to a specific pain. Nothing is on here just because it seemed cool to build.

Feature	Pain it solves	How it works
Fit score engine	Can't tell if she's a fit	Priya pastes her resume and the JD. The system returns a 0 to 100 score with a tier (strong, reach, or weak), a keyword overlap breakdown, and JD line citations that show exactly why she scored where she did.
Skill gap analysis	Doesn't know what's missing	Pulls required skills from the JD, compares them against the resume, and surfaces the gap with JD evidence attached. Not generic advice, but specific proof of what's missing.
Bullet rewriter	Can't reframe experience in JD language	Takes each of Priya's existing experience bullets and rewrites them to reflect outcomes and the terminology from the JD. She can accept or reject each suggestion one by one.
Cover letter generator	Starts from scratch every time	Streams a personalized cover letter built from the fit analysis and her accepted bullet rewrites. It's specific to the role, not a template with a name swapped in.
LinkedIn section audit	Profile isn't positioned for target role	Scans her headline, About section, Experience, Skills, and photo and banner. Returns section-level gaps and prioritized fixes, each tied back to JD evidence.
Match report and print	No record of what she submitted	An exportable report for each application. Both the Resume and LinkedIn tabs share the same visual format: score, keywords, actions, and a print to PDF option.

STEP 7

Prioritize the Features: Top Two

With six features on the table, the question is which two to build first. The criteria are simple: it has to solve the highest-severity pain, it has to work reliably with AI, and it has to be shippable without depending on features that don't exist yet.

Feature	Pain severity	AI reliability	Dependency-free?	Ship priority
Fit score and gap analysis	HIGH	High	Yes, needs only resume and JD	SHIP FIRST
Bullet rewriter	HIGH	High	Yes, works on the resume the user already has	SHIP SECOND
Cover letter generator	Moderate	High	No, works best after bullets are rewritten	v1 (after top 2)
LinkedIn section audit	Moderate	High	Yes, standalone pipeline	v1 (parallel)
Match report and print	Low	N/A	No, needs the score first	v1 (bundled)
Skill gap analysis	High	High	No, part of the fit score flow	Bundled with feature 1

The top two features

Feature 1: Job Fit Score and Gap Analysis

This lives in Steps 1 and 2 of the **Job Application Workflow Agent**. Priya pastes her resume and the JD, and within 60 seconds she has a 0 to 100 score, a tier classification, a keyword breakdown, and a list of skill gaps with JD citations. This is the entry point. She knows whether it's worth tailoring before she spends a minute doing it.

Feature 2: Bullet Rewriter

This is Step 3 of the **Job Application Workflow Agent**. Each bullet from Priya's existing resume gets rewritten to reflect outcomes and the language of the JD. She reviews every suggestion and accepts or rejects it. AI makes the suggestions, she makes the call. There is no hallucination risk here because the source is always her own text.

Why the other features are still in scope

The cover letter, which is Step 4 of the Workflow Agent, ships in the same product because it is downstream of everything else. By the time Priya reaches Step 4, the fit analysis and bullet rewrites are already done, and the cover letter can be built from that data. It is a natural completion of the workflow, not a separate thing to build.

The LinkedIn audit lives in a separate product, the Profile Optimization Tool, because it serves a different moment. It is not for when Priya is actively applying. It is for when she wants to know whether her baseline profile is ready for the kind of role she is targeting. Same user, but a different trigger and a different session.

STEP 8

Wireframes and Success Metrics

Tool 1: Profile Optimization Tool

Solves: LinkedIn profile not positioned for target role. Moment: Before the active application, as a baseline readiness check.

Wireframe

The tool lives at one URL with two tabs: Resume and LinkedIn. The layout is scan-first: the score is the first thing Priya sees, then the gaps, then the actions she can take.

Input	Analysis	Output
Resume tab: Paste resume text + Job description	LLM keyword extraction + Deterministic scoring (45% keywords / 35% checklist / 20% section audit)	Match % hero score Keyword gap matrix Section checklist Prioritized action plan Print → PDF export
LinkedIn tab: Paste LinkedIn URL + Target JD(s)	Profile fetch (Proxycurl) + Same LLM + scoring pipeline	Alignment % + per-JD chips Section audit: headline, About, Experience, Skills Role conflict flags Print → PDF export

Key design decisions

Score first: The match percentage is the first thing Priya sees. Before she reads a single recommendation, she already knows whether the profile is in good shape or not. That number sets the context for everything that follows.

Every recommendation cites the JD: Each suggested fix points back to the exact line in the job description that triggered it. No generic advice. Every action Priya takes is connected to a specific requirement.

Nothing gets stored: No resume or LinkedIn content is saved between sessions. Priya pastes in, gets her report, and nothing persists afterward. This is a deliberate trust decision, not just a technical one.

Success metrics for the Profile Optimization Tool

Metric	Type	Target	Why it matters
Analysis completion rate	Primary	> 85%	Most users who start should finish. High drop-off is a signal of UX friction or results that feel off.
Time to report (p50)	Primary	< 60 seconds	This is the core promise. If it takes three minutes, Priya will just do it manually.

Metric	Type	Target	Why it matters
Action plan engagement	Secondary	> 60% scroll past score	If users only look at the score and leave, the recommendations aren't landing.
Print and export rate	Secondary	> 20% of sessions	Exporting means Priya treated the output as a real work product. That's the strongest quality signal we have.
AI hallucination rate	Guardrail	0 per 100 analyses	The LLM must never invent employers, skills, or JD requirements. This is enforced through prompt guardrails and monitored continuously.

Tool 2: Job Application Workflow Agent

Solves: Can't tell if she's a fit, can't rewrite bullets, and faces a blank page every time she writes a cover letter. Moment: At the point of application, when she is actively tailoring for a specific role.

Wireframe: 4-step guided workflow

The agent breaks the tailoring process into four sequential steps, each one unlocking the next. This structure mirrors exactly how a career coach works through an application with someone: you do not hand a person everything at once. You go step by step.

Step	What the user sees	What the system does	Output
Step 1 Analyze Fit	Single input screen: paste resume + JD	Extracts keywords, scores overlap, classifies tier. Cites JD lines for each strength/gap.	0–100 score Strong / Reach / Weak tier JD citations
Step 2 Surface Gaps	List of missing skills, each with a JD quote	Identifies skills required by JD not evidenced in resume. Returns structured gap list with proof.	Named skill gaps + JD evidence for each
Step 3 Rewrite Bullets	Side-by-side: original bullet vs. AI rewrite. Accept or reject each.	Takes each bullet, rewrites it using JD language and outcome framing. Grounded entirely in user's original text.	Accepted rewrites ready to paste into resume
Step 4 Draft Cover Letter	Streaming text output. Full personalized letter.	Uses the fit score, gaps, and accepted bullets as input. Streams a letter specific to this role, not a generic template.	Role-specific cover letter (streaming)

Key design decisions

Steps, not a single output dump: Getting a fit score, a gap list, rewritten bullets, and a cover letter all at once would be overwhelming. Breaking it into four sequential steps makes each piece digestible and actionable. This is the same reason a good coaching session works the way it does: you go through it together, one thing at a time.

Accept or reject on every bullet: Priya is always in control. The AI makes suggestions and she decides what stays. This is important for trust, but also for accuracy. She knows context about her experience that the model simply does not have access to.

Streaming for the cover letter: Streaming output makes the experience feel fast and responsive. It also signals that the letter is being written live for this specific role, not pulled from somewhere generic.

Regression-tested before every ship: Every prompt change runs through a five-case golden test set before it goes out. The cases cover strong match, weak match, a partial AI skill gap, an overqualified seniority signal, and a minimal-resume edge case. Improvements are measured, not assumed.

Success metrics for the Job Application Workflow Agent

Metric	Type	Target	Why it matters
Workflow completion rate (all 4 steps)	Primary	> 60%	Users who complete all steps get the full value. Looking at drop-off by step tells us exactly where the experience breaks down.
Fit score accuracy (golden set pass rate)	Primary	5/5 per deploy	This is the core quality gate. A score must land in the expected band before any prompt change ships.
Bullet acceptance rate	Secondary	> 50% of rewrites accepted	If Priya rejects most of the AI suggestions, the rewrites are not useful enough. That's a prompt tuning problem.
Cover letter generation rate	Secondary	> 40% of users who reach Step 3	Tells us whether users trust the workflow enough to see it through to the end.
Step 1 to Step 2 drop-off rate	Secondary	< 20%	If users see their fit score and leave immediately, either the score is discouraging or Step 2 isn't communicating its value clearly.
Hallucination and fabrication rate	Guardrail	0 per 100 analyses	The LLM must never invent skills, employers, or JD requirements. This is enforced through prompt guardrails and retry logic.

SUMMARY

What This Shows About How I Think

These two tools are not demos. They are the output of a real product process: starting with actual research, narrowing to a specific user, and making deliberate trade-offs at every step along the way.

Framework step	Decision made	Tool it shaped
Scope the problem	Narrowed from 'help with job search' to 'convert applications by closing the gap between your profile and a specific JD'	Both
Segment users	Three distinct segments identified by experience level and application behavior, not just demographics	Both
Pick target segment	Mid-career pivoters: highest pain, best fit for what AI can do well, and the largest group in my coaching data	Both
Persona and journey	Priya's seven-stage journey maps the real pain at every touchpoint, not just the obvious ones	Both
Prioritize pains	Fit assessment and bullet rewriting came out on top: high frequency, high stakes, and reliably solvable with AI	Workflow Agent
List solutions	Six features, each derived from a specific pain. Nothing on the list exists without a pain to justify it.	Both
Prioritize features	Fit score and gap analysis ships first as the entry point. Bullet rewriter ships second as the conversion step. LinkedIn audit runs in parallel because it serves a different moment.	Both
Wireframes and metrics	Each tool has primary metrics tied to its core promise, secondary engagement signals, and a hallucination guardrail	Both

Where this came from

The persona, pain points, and feature priorities here are not hypothetical. They come from 350+ career coaching sessions, direct observation of job seekers working through tailoring and bullet rewrites in real time, and my own active job search. The specifics, like spending 30 to 45 minutes per application, being unsure about when tailoring is actually worth it, and writing bullets that describe tasks instead of outcomes, are patterns I have seen across hundreds of conversations. Not assumptions.

Try the tools: ankitsingh.net/ai-prototypes · Technical docs: ankitsingh.net/ai-prototypes/tech-docs

Ankit Singh · ankit.singh101@gmail.com · linkedin.com/in/ankitsingh