

AI Roadmap Revolution

Six scenarios — expanded edition. Each scenario shows the data the AI saw, what it recommended, the techniques it used, and questions to discuss in lab.

How to use this walkthrough

Each scenario takes one full page. Read top to bottom: the data the AI was given, the recommendation it produced, the reasoning path it followed. The bottom of each page lists **the specific NLP techniques the AI used** — match these against the toolkit reference at the end to build fluency. The discussion prompts are designed to push past the first answer and surface the trade-offs every Product Owner has to weigh.

Six domains, six approaches:

1. FITNESS Digital Fitness App	2. EDUCATION Online Learning Platform	3. LANGUAGE Language Learning App
4. DELIVERY Food Delivery Platform	5. WELLNESS Mental Wellness App	6. TRAVEL Travel Planning Website

A note on technique-spotting. One of the most useful skills a Product Owner can develop is recognizing which AI technique fits a given problem. As you read each scenario, try to predict which techniques the AI used *before* reaching the callout — then check your guess. Over six scenarios, the patterns become visible.

SCENARIO 1 · FITNESS

Digital Fitness App

When user-facing content fails to match user readiness, the friction shows up first in abandonment data — and second in the language users use to describe the failure.

DATA PRESENTED TO THE AI

- User feedback surveys flagged that workout videos were too generic or mismatched with user fitness levels.
- Session abandonment rates spiked for beginner users after 3–5 sessions.
- Reviews repeatedly mentioned phrases like “too hard,” “not beginner-friendly,” and “confusing instructions.”

AI RECOMMENDATION

Prioritize a personalized workout path generator that adapts video difficulty and pacing based on user fitness assessments and past activity.

HOW THE AI ARRIVED AT THIS

The AI performed semantic clustering on feedback to identify a pattern around perceived difficulty and engagement drop-off. Sentiment analysis confirmed frustration concentrated among newer users. Abandonment rates were correlated with account metadata — specifically fitness-level tags — which surfaced this as a user-type-specific need rather than a universal problem.

AI TECHNIQUES USED

Semantic Clustering	grouped similar complaints by meaning, not just keywords
Sentiment Analysis	isolated emotional tone — frustration vs. neutral feedback
Correlation Analysis	linked abandonment patterns to fitness-level metadata

FOR DISCUSSION

1. What other user metadata could sharpen the personalization (age, prior injury, time of day)?
2. If the AI's recommendation requires fitness assessments, what's the cost to user onboarding — and is that trade worth it?
3. How would you measure whether the personalized path is actually working after launch?

SCENARIO 2 · EDUCATION

Online Learning Platform

Course completion data tells you where users get stuck. Support tickets tell you why. Pairing them is what turns observation into a backlog item.

DATA PRESENTED TO THE AI

- High dropout rates in technical courses compared to soft-skills courses.
- Users submitted support tickets asking “Where do I start?” or “This is too hard to follow.”
- Completion data showed users jumping between modules or skipping foundational content.

AI RECOMMENDATION

Develop a pre-assessment-driven course onboarding assistant that recommends personalized learning paths based on prior knowledge and career goals.

HOW THE AI ARRIVED AT THIS

The AI used sequence pattern analysis on course progression data and paired it with language-based ticket analysis to detect mismatches between user preparedness and course design. Topic modeling on the support tickets uncovered a recurring theme of confusion around course sequencing — students didn't know which module was a prerequisite for which.

AI TECHNIQUES USED

Sequence Pattern Analysis	found common navigation paths and where users diverge
Topic Modeling	extracted themes from unstructured support tickets
Text Classification	categorized tickets by problem type for quantification

FOR DISCUSSION

1. Could a simpler intervention — like a recommended learning path on each course page — solve this without an AI assistant?
2. What would a pre-assessment look like that takes less than 60 seconds to complete?
3. How might this approach unintentionally exclude self-directed learners who don't want assessments?

SCENARIO 3 · LANGUAGE

Language Learning App

User complaints about “boring” content rarely mean the content is boring in absolute terms — they usually mean the content has stopped progressing.

DATA PRESENTED TO THE AI

- App store reviews and in-app feedback frequently mentioned “repetitive” or “boring” exercises.
- Engagement data showed sharp drop-offs after 7 consecutive days for mid-level learners.
- Popular user forum posts suggested users wanted more real-world, conversational practice.

AI RECOMMENDATION

Introduce an AI-powered conversation simulator that lets users practice dynamic dialogues with context-aware feedback.

HOW THE AI ARRIVED AT THIS

Frequency analysis of review terms highlighted repetition-related complaints. Engagement analytics pinpointed exactly when user interest waned — around the seven-day mark for mid-level learners. Clustering of forum comments revealed a strong unmet desire for natural conversation practice, which the AI matched against broader market trends in language learning to validate the recommendation.

AI TECHNIQUES USED

Frequency Analysis	identified the most-repeated complaint terms in reviews
Clustering / Similarity	grouped forum posts by underlying user need
Market Trend Matching	validated that the recommendation aligns with industry direction

FOR DISCUSSION

1. Conversational practice is expensive to build well. What's a smaller first step that tests the demand?
2. How do you handle the privacy and content-safety concerns of an AI conversation partner?
3. Mid-level learners drop off — but what about beginners? Should they see this feature too?

SCENARIO 4 · DELIVERY

Food Delivery Platform

When low ratings cluster around two specific complaints (“cold food,” “long wait”), the data is telling you exactly what to fix — but the fix lives upstream of the rating itself.

DATA PRESENTED TO THE AI

- Users gave 1–2 star ratings citing “cold food” or “long wait times.”
- Delivery logs showed wide variance in driver delivery routes and restaurant prep times.
- Chat logs and service tickets indicated confusion about delivery time estimates.

AI RECOMMENDATION

Implement a predictive delivery ETA system that uses live kitchen prep data and real-time traffic patterns to give more accurate time estimates and reroute drivers intelligently.

HOW THE AI ARRIVED AT THIS

Correlation analysis linked low satisfaction scores directly to delivery data inconsistencies. Named entity recognition on chat logs extracted the specific complaint terms — “lateness,” “cold meals” — and let the AI quantify their frequency. Anomaly detection flagged a small set of restaurants with consistently inconsistent prep-to-dispatch times, suggesting the problem wasn't system-wide.

AI TECHNIQUES USED

Correlation Analysis	linked rating drops to specific delivery data anomalies
Named Entity Recognition (NER)	extracted complaint nouns from unstructured chat text
Anomaly Detection	flagged outlier restaurants for targeted intervention

FOR DISCUSSION

1. If only a few restaurants drive most of the problem, does the AI ETA system help — or is the right move to fix those restaurants?
2. What's the failure mode if your ETA prediction is consistently 5 minutes too optimistic?
3. Could showing users a prep-time progress bar reduce complaints without changing the actual delivery time?

SCENARIO 5 · WELLNESS

Mental Wellness App

Engagement is a proxy for value — but only when the user's goal aligns with using the app. In wellness, low engagement after a heavy session may signal that the app delivered something the user couldn't process alone.

DATA PRESENTED TO THE AI

- Journaling feature was used heavily at night but had low retention after the first week.
- Sentiment scores of journal entries showed high emotional distress but no follow-up engagement.
- Exit surveys mentioned, “I didn't know what to do after writing,” or “It felt like venting into a void.”

AI RECOMMENDATION

Build an AI-guided reflection coach that reads journal entries and recommends appropriate mental health activities (breathing exercises, CBT prompts, therapist referrals).

HOW THE AI ARRIVED AT THIS

Emotion classification on journal text revealed high distress, while behavioral tracking showed those same users disengaged within hours. The pattern flagged a gap in the support flow — the app surfaced emotional content but offered nothing in response. Topic clustering of exit surveys emphasized a missing “next step,” which directly informed the recommendation.

AI TECHNIQUES USED

Emotion Classification	detected fine-grained emotional states beyond positive/negative
Behavioral Tracking	correlated emotional content with subsequent app activity
Topic Clustering	grouped exit-survey complaints by underlying theme

FOR DISCUSSION

1. What's the ethical line between an AI that suggests CBT exercises and one that should be referring users to a human?
2. How do you avoid the AI making users feel surveilled when it responds to their journal content?
3. If a journal entry suggests crisis, what should the AI do — and who decides the threshold?

SCENARIO 6 · TRAVEL

Travel Planning Website

When users abandon a complex flow more than a simple one, the answer isn't always to simplify the flow — sometimes it's to change the input modality entirely.

DATA PRESENTED TO THE AI

- Users abandoned multi-city trip booking flows more than single-destination trips.
- Clickstream data showed frequent toggling between tabs like “map,” “dates,” and “compare fares.”
- Feedback forms contained comments like “too complicated,” “wish I could just tell it what I want,” and “I got lost in the options.”

AI RECOMMENDATION

Introduce an AI travel assistant chat feature that allows users to input natural language queries like “Plan me a 2-week Europe trip with museums and beaches” and returns optimized trip plans.

HOW THE AI ARRIVED AT THIS

Journey mapping based on session data and interaction heat-maps revealed high decision friction in complex trip planning. Natural language processing on feedback comments revealed a strong, repeated desire for simplified, conversational planning — leading directly to the chatbot recommendation. The data suggested users didn't want fewer options; they wanted a different way to navigate them.

AI TECHNIQUES USED

Journey Mapping	modeled the actual paths users took through the booking flow
Heat-Map Analysis	visualized where attention concentrated on each page
Intent Detection	extracted what users were trying to accomplish from feedback

FOR DISCUSSION

1. A chatbot is a major UX shift. What's a low-risk way to test demand before committing?
2. What happens when the AI mis-interprets a query and books the wrong trip? Who's responsible?
3. Heat-maps show users toggling between map/dates/fares. Could a unified single-screen view solve this without any AI?

NLP Toolkit Reference

The AI techniques referenced throughout the scenarios belong to a small family of NLP tools. Use this as a quick reference when matching scenarios to techniques.

Task	What it does	Common tools
Sentiment Analysis	Detects whether text is positive, negative, or neutral.	<i>MonkeyLearn · Amazon Comprehend</i>
Topic Extraction	Identifies main themes or key phrases in a body of text.	<i>spaCy · GPT APIs</i>
Named Entity Recognition	Finds names of people, places, brands, products in unstructured text.	<i>Hugging Face Transformers</i>
Text Classification	Sorts text into user-defined groups (spam, complaint, praise, etc.).	<i>Google Cloud NLP · Azure Text Analytics</i>
Summarization	Generates a short summary of longer content.	<i>OpenAI (GPT-4) · Cohere</i>
Intent Detection	Identifies the user's goal (to buy, complain, ask for help).	<i>Dialogflow · Rasa</i>
Text Generation	Writes new content or completes what the user started.	<i>OpenAI ChatGPT · Jasper AI</i>
Clustering / Similarity	Groups similar pieces of text together based on meaning.	<i>Haystack · SentenceTransformers</i>

Prepared for Rocket Nine Solutions course materials — *AI for Product Owners*.