AI PM ROADMAP

120 concepts to become an AI product manager

8 areas, 15 concepts each. Every concept explained in one line, the order to study each area, and two practice questions to test yourself.

How to use it: one area at a time, tick each concept only when you can explain it out loud in under a minute, then answer the two practice questions without notes.

@chalkesahil · chalkesahil.ai

AREA 01 / 08

Product basics

Study order: Learn the core PM loop on any product before adding AI.

01	**What a PM does**	Decide what to build and why, then get it shipped.
02	**User problems**	Start from a pain, not a feature.
03	**Jobs to be done**	What the user is really hiring the product for.
04	**Personas**	Who exactly you are building for.
05	**User interviews**	Ask about past behaviour, not future wishes.
06	**Problem statements**	One clear sentence the team can rally around.
07	**Prioritisation**	RICE or impact vs effort to pick what's next.
08	**Roadmaps**	Themes and outcomes, not a list of dates.
09	**PRDs**	The doc that explains the problem, scope and success.
10	**User stories**	As a user, I want, so that.
11	**MVP**	The smallest thing that tests your riskiest assumption.
12	**Metrics**	North star, inputs and guardrails.
13	**Working with engineers**	Clear problems, not dictated solutions.
14	**Stakeholder management**	Keep leadership, sales and support aligned.
15	**Launch planning**	Rollout, comms and what you'll watch after.

Practice

1. Prioritise five feature requests for a food delivery app.
2. Write a one-page PRD for a new feature in an app you use daily.

AREA 02 / 08

AI fundamentals

Study order: You don't need to code models, but you must know what they can and can't do.

01	**AI vs ML vs deep learning**	Nested ideas, from broad to specific.
02	**Supervised learning**	Learn from labelled examples.
03	**Unsupervised learning**	Find patterns without labels.
04	**Training vs inference**	Building the model vs using it.
05	**Training data**	The model is only as good as what it learned from.
06	**Overfitting**	Great on old data, poor on new data.
07	**Precision and recall**	Wrong alarms vs missed cases.
08	**Confidence scores**	How sure the model is, and when to trust it.
09	**Bias in data**	Skewed data gives skewed decisions.
10	**Recommendation systems**	How feeds and suggestions are ranked.
11	**Computer vision**	Models that read images and video.
12	**Speech and audio**	Speech to text and text to speech.
13	**Generative AI**	Models that create text, images and code.
14	**Foundation models**	Big general models you adapt instead of training.
15	**Build vs buy**	Call an API, fine-tune, or train your own.

Practice

1. Explain precision vs recall using a spam filter.
2. Should a small startup train its own model? Argue both sides.

AREA 03 / 08

LLM capabilities

Study order: Use the top models daily so you know their strengths and failures first-hand.

01	**How LLMs work**	Predict the next token from huge amounts of text.
02	**Context window**	How much the model can read at once.
03	**Tokens and cost**	Every word in and out has a price.
04	**Latency**	Users notice every extra second.
05	**Hallucinations**	Confident wrong answers; design around them.
06	**Prompting**	Clear instructions and examples change quality a lot.
07	**RAG**	Let the model answer from your company's documents.
08	**Fine-tuning**	Teach a model a style or task with examples.
09	**Tool use**	Models that call APIs and take actions.
10	**Agents**	Models that plan and act over many steps.
11	**Multimodal**	Text, images, audio and files in one model.
12	**Model comparison**	Claude, GPT, Gemini and open models each have trade-offs.
13	**Open vs closed models**	Control and cost vs convenience and quality.
14	**Model updates**	A new version can change your product's behaviour.
15	**Limits of current AI**	Know what still needs a human.

Practice

1. When would you use RAG instead of fine-tuning?
2. A model update made your bot's answers worse. What do you do?

AREA 04 / 08

AI product discovery

Study order: Find problems where AI is clearly better, not just newer.

01	**AI-first vs AI-added**	Rebuilding the workflow vs adding a chat box.
02	**Good AI use cases**	Repetitive, language-heavy, tolerant of small errors.
03	**Bad AI use cases**	Zero error tolerance with no human check.
04	**Automation vs assistance**	Do it for the user, or help the user do it.
05	**Human in the loop**	Where a person reviews before it counts.
06	**Data availability**	No data, no AI feature.
07	**Feasibility checks**	Test with a prompt before writing a PRD.
08	**Cost per task**	AI features have a running cost per use.
09	**Value per task**	Time or money saved for the user.
10	**Trust design**	Show sources, confidence and an undo.
11	**Failure modes**	List how it can go wrong before launch.
12	**Competitive moats**	Your data and workflow, not the model.
13	**Pricing AI features**	Seats, usage or outcomes.
14	**Wizard of Oz testing**	Fake the AI with a human to test demand.
15	**Writing an AI PRD**	Add sections for data, evals, risks and fallbacks.

Practice

1. Pick one feature in a banking app that AI should improve, and one it shouldn't.
2. How would you test demand for an AI feature before building it?

AREA 05 / 08

Prototyping

Study order: Build working prototypes yourself; AI PMs who can ship demos stand out.

01	**Prompt prototypes**	Test the core idea in a chat window first.
02	**System prompts**	Write the instructions your product will use.
03	**No-code builders**	Lovable, Bolt or v0 for clickable demos.
04	**Vibe coding**	Describe the app, let AI write the code, then test it.
05	**Claude Code and Cursor**	AI coding tools to build real prototypes.
06	**Workflow tools**	n8n or Zapier to chain AI steps without code.
07	**Custom GPTs and Projects**	Packaged assistants for quick internal tools.
08	**API playgrounds**	Compare models and settings side by side.
09	**Mock data**	Realistic sample data for demos.
10	**Figma with AI**	Faster wireframes and flows.
11	**Demo scripts**	Show the happy path and one recovery from failure.
12	**User testing prototypes**	Put it in front of five users in a week.
13	**Prototype to spec**	Turn what you learned into a clear build plan.
14	**Prompt libraries**	Save prompts that worked, with notes.
15	**Portfolio**	Two or three shipped AI prototypes with write-ups.

Practice

1. Prototype a meeting-notes summariser in one evening.
2. Your demo works in testing but fails with real users. What went wrong?

AREA 06 / 08

Evals and metrics

Study order: Define 'good' before launch; AI quality is a product decision.

01	**What evals are**	Tests that score AI output quality.
02	**Golden sets**	Real examples with agreed good answers.
03	**Rubrics**	Clear criteria for what makes an answer good.
04	**Human evaluation**	Experts grade a sample of outputs.
05	**LLM as a judge**	A model grades outputs at scale.
06	**Offline vs online evals**	Test before launch vs measure in production.
07	**Task success rate**	Did the user get what they came for?
08	**Accuracy vs helpfulness**	Correct but useless still fails.
09	**Latency and cost metrics**	Quality has to be affordable and fast.
10	**Adoption metrics**	Who uses it, how often, and do they return?
11	**Thumbs up and down**	Cheap feedback; read the comments too.
12	**Edit rate**	How much users change the AI's output.
13	**Escalation rate**	How often it hands off to a human.
14	**A/B tests for AI**	Compare models or prompts on real traffic.
15	**Regression checks**	Re-run evals before every model change.

Practice

1. Define success metrics for an AI customer support bot.
2. Thumbs-up rate is high but retention is falling. What's going on?

AREA 07 / 08

AI risk and ethics

Study order: Every AI PRD needs a risk section; interviewers check for it.

01	**Hallucination risk**	Where a wrong answer can hurt a user.
02	**Bias and fairness**	Check outcomes across user groups.
03	**Privacy**	What user data goes to which model provider.
04	**Data retention**	How long prompts and outputs are kept.
05	**Consent**	Users should know when AI is involved.
06	**Transparency**	Explain what the AI did and why.
07	**Prompt injection**	Content that hijacks your AI's instructions.
08	**Abuse and misuse**	How bad actors could use your feature.
09	**Content safety**	Filters for harmful output.
10	**Copyright**	Training data and generated content questions.
11	**India's DPDP Act**	India's data protection law and what it means for your product.
12	**Regulated domains**	Health, finance and legal need extra care.
13	**Human oversight**	Who can override the AI and how fast.
14	**Incident response**	What happens when the AI does harm.
15	**Red teaming**	Try to break your own feature before launch.

Practice

1. List the risks of an AI feature that answers medical questions.
2. A user shows your bot giving biased loan advice. What do you do in the first 24 hours?

AREA 08 / 08

PM interviews

Study order: Practise out loud with a timer; structure beats clever answers.

01	**Product sense**	Design a product for a user and a goal.
02	**Product improvement**	Pick a product, a user, a problem, a fix.
03	**Metrics questions**	Define success and diagnose drops.
04	**Root cause analysis**	Data issues first, then internal, then external.
05	**Estimation**	Clear assumptions beat exact numbers.
06	**Prioritisation questions**	Show your trade-offs out loud.
07	**Strategy questions**	Market, competition and where AI changes the game.
08	**AI product design**	Design an AI feature with data, evals and fallbacks.
09	**Technical depth**	Explain RAG, agents and costs simply.
10	**Behavioural questions**	STAR stories about conflict, failure and impact.
11	**Why PM**	A real moment you owned a product decision.
12	**Why AI PM**	Show an AI prototype you built, not just interest.
13	**Case study presentation**	Problem, insight, solution, metrics, risks.
14	**Resume for PM**	Outcomes and decisions, not tasks.
15	**Questions to ask them**	Ask how they measure AI quality today.

Practice

1. Design an AI feature for Swiggy that increases repeat orders.
2. Daily active users dropped 10% after an AI launch. Investigate.