

AI 101 for Product Managers

PART 1: FOUNDATIONS

1/ What was the name of the 2017 research paper published by Google that started the AI race?

0

"Teaching Machines to Talk"

0

"Attention Is All You Need"

0

"Bigger Is Better: Scaling Neural
Networks"

0

"Deep Learning Is Enough"

TAKEAWAY:

Attention Is All You Need

Ashish Vaswani, Noam Shazeer, and others at
Google Brain team

2/ What does "GPT" in ChatGPT stands for?

0

General Purpose Technology

0

Generative Pre-trained
Transformer

0

Global Prediction Tool

0

Generated Prompt Text

TAKEAWAY:

Generative: it creates new text instead of just sorting or labelling existing text.

Pre-trained: it first learns from a massive amount of text, then gets fine-tuned to act as a helpful assistant.

Transformer: the design that lets it look at all the words at once and work out which ones matter to each other.

2/ What does "GPT" in ChatGPT stands for?

0 

General Purpose Technology

0 

Generative Pre-trained
Transformer

0 

Global Prediction Tool

0 

Generated Prompt Text

TAKEAWAY:

Generative: it creates new text instead of just sorting or labelling existing text.

Pre-trained: it first learns from a massive amount of text, then gets fine-tuned to act as a helpful assistant.

Transformer: the design that lets it look at all the words at once and work out which ones matter to each other.

3/ Which of these tasks is AI most likely to get WRONG without extra tools?

Which one trips AI up?

A	B	C	D
			
Reply to an upset customer	Summarize a 10-page doc	Translate a help article to Spanish	Add up 5,000 rows of contract values

0 ×

A) Drafting a polite reply to an upset customer

0 ×

B) Summarizing a 10-page document

0 ×

C) Translating a help article into Spanish

0 ✓

D) Adding up 5,000 rows of contract values exactly

TAKEAWAY:

AI's skills are "jagged". It's brilliant with language but unreliable at exact maths, counting and very recent facts unless you give it tools.

Quiz leaderboard

No results yet

Top Quiz participants will be displayed here once there are results!

PART 2: HOW A MODEL IS TRAINED

4/ A base model (pre-trained only) sees this prompt. What is it most likely to write next?

Base model, fresh out of pre-training

> What is the capital of France?

> █

What does it type next?

0 ×

A) "Paris! Anything else I can help with?"

0 ×

B) "I'm sorry, I can't answer that."

0 ✓

C) "What is the capital of Germany?"

0 ×

D) "Could you clarify which France you mean?"

TAKEAWAY:

Pre-training builds raw knowledge by predicting the next word across trillions of words. It's the expensive part, and it produces a "base model" that knows a lot but isn't an assistant yet.

5/ Supervised Fine-Tuning (SFT) mainly teaches the model to...

0

Learn new facts about the world

0

Follow instructions, using example
conversations

0

Run faster and cheaper

0

Browse the internet

TAKEAWAY:

Fine-tuning on example conversations teaches behavior and format. Plays a big role in deciding the tone of the model - professional, casual, direct, etc.

6/ In RLHF (Reinforcement Learning from Human Feedback), what do humans actually do?

0



A) Write the model's code

0



B) Label photos of cats and dogs

0



C) Approve every answer in real time

0



D) Compare and rank the model's responses

TAKEAWAY:

Human preferences shape the model's taste, which is also why models can be people-pleasers. Ask AI for critique, not approval.

7/ "Reasoning" or "thinking" models got much better at maths and coding mainly through...

0 

A) Memorizing more answers

0 

B) Reinforcement learning on
problems with checkable answers

0 

C) Hiring more people to type
answers


0 

D) Faster chips, nothing else

TAKEAWAY:

Models get sharpest where answers can be checked automatically. "Thinking" buys quality at the cost of time and tokens.

8/ These training stages got shuffled. Which order is correct?



0 ✓

0 ✗

0 ✗

0 ✗

A) B → A → C → D

B) C → B → A → D

C) B → C → A → D

D) D → A → B → C

TAKEAWAY:

Data → pre-train → fine-tune → preference tuning & reasoning RL → safety testing → deploy & monitor. It's product development: platform, experience, test, ship, iterate.

9/ Your AI feature needs to know our latest policy docs. What is the smartest first move?

0 

A) Pre-train our own model from scratch

0 

B) Fine-tune a model right away

0 

C) Fetch the relevant docs and add them to the prompt (RAG)

0 

D) Wait for the next model release

TAKEAWAY:

Climb the ladder cheapest-first:
prompt → RAG → fine-tune → train. To give AI
fresh facts, retrieve them; don't retrain.

Quiz leaderboard

No results yet

Top Quiz participants will be displayed here once there are results!

PART 3: EVALS, MCPS, AND AGENTS

10/ In AI product development, what is an "eval"?

0 

A) A set of test cases that measures output quality on your task

0 

B) A premium pricing tier

0 

C) A legal review of the model


0 

D) A set of responses generated by LLM

TAKEAWAY:

Evals are the PM's new acceptance criteria. If you can't define "good" with test cases, you can't ship AI with confidence.

11/ What is MCP?



0 ✓

A) An open standard for plugging AI apps into tools and data

0 ✗

B) A new AI model from a big tech company

0 ✗


C) A security certification for AI

0 ✗

D) A replacement for all APIs

TAKEAWAY:

MCP doesn't replace APIs. It wraps them so AI can find and use them itself, turning $N \times M$ custom integrations into $N + M$.



12/ What makes something an "AI agent" rather than a chatbot?

0 

A) It has a name and a personality

0 

B) It plans and takes actions with tools, in a loop

0 

C) It can talk using voice

0 

D) It runs on a bigger model

TAKEAWAY:

Agent = model + tools + a loop + a goal. More autonomy needs more guardrails, especially before irreversible actions.

Quiz leaderboard

No results yet

Top Quiz participants will be displayed here once there are results!

PART 4: AI IN YOUR PM DAY-TO-DAY

13/ Four PMs asked AI to review their PRD. Whose prompt gets the most useful feedback?

Prompt-off: whose PRD review prompt wins?

Priya

P

"Review my PRD."

Marco

M

"Make this PRD better."

Jen


J

"Is this PRD good?
Yes or no."

Sam

S

"You're a skeptical eng lead.
List the 5 biggest gaps in this PRD for SMB admins, citing the section for each."



Participant	Votes	Status
Priya	0	Incorrect (x)
Marco	0	Incorrect (x)
Jen	0	Incorrect (x)
Sam	0	Correct (✓)

TAKEAWAY:

Specific beats vague: role + task + context + format. And ask AI to critique, not compliment.

14/ You have 40 customer interview transcripts and a readout on Friday. Best way to use AI?

0 

A) Ask for "key insights" and paste them straight into the deck

0 

B) Don't. AI can't read transcripts

0 

C) Have AI tag themes with direct quotes, then spot-check the sources

0 

D) Only summarize the first five; the rest are probably the same

TAKEAWAY:

Let AI do the first pass of synthesis, **but** demand quotes with sources and **verify** a sample yourself.

15/ You want stakeholder feedback on a new feature idea by Friday. Fastest route?

0



A) Wait for design capacity next sprint

0



B) Write a 20-page spec first

0



C) Build a clickable prototype with an AI tool

0



D) Ask engineering for a two-week spike

TAKEAWAY:

Show, don't tell: an AI-built prototype in an afternoon gets sharper feedback than a 20-page spec.

16/ AI drafted your weekly stakeholder update in 30 seconds. What's your job now?

0

A) Hit send. Efficiency!

0

B) Give it real data, then edit for accuracy and judgment

0

C) Nothing. The AI owns it now

0

D) Fix the typos and go to lunch

TAKEAWAY:

AI writes the first draft; **you** own the facts and the judgment.

Quiz leaderboard

No results yet

Top Quiz participants will be displayed here once there are results!

Honest impression of the Session in one word?