Principles in Action Playbook: Ideation

Should you use AI?

- [ ] Are you building something good for the world?
- [ ] Do you need to use AI?
- [ ] How realistic is your solution?
- [ ] What are your success metrics?
- [ ] How can you mitigate risks?

Are you building something good for the world?

Instead of approaching your brainstorming with "How can we use AI to ___?"

→ Move to "How might we solve [human need]?"

→ And then "Can AI solve [this need] in a unique and helpful way?"

'How might we' statements force you to start with a human need and steer you away from suggesting a solution, so that you can be open to generating further possibilities.

How might we:

What evidence do you have that this is a problem? (Listen to people, look at data, watch behaviours)

Do you need to use AI?

Let's consider whether you actually need AI in your product to solve your problem.

What is the need, task, or user experience you want to improve?

Have your users expressed concerns about using AI? What did they say?

Using the chart below as a starting point, consider whether adding AI will improve your product experience, do nothing for your product, or maybe even degrade it.

AI is good at ...	Avoid AI when ...
☐ Automating tasks that people don't know how to do, or find boring, repetitive, or dangerous	☐ People want to do a task without help
☐ Augmenting tasks that people enjoy doing	☐ People want creative control to see their vision through
☐ Recommending different content to different users	☐ Being predictable is essential, always
☐ Predicting and forecasting trends and events	☐ The cost of errors is greater than the benefits of a small increase in success rate
☐ Personalising a user experience	☐ You and your customers need to understand exactly why something happened
☐ Understanding human language	☐ Shipping fast is priority
☐ Recognizing patterns in images, text, or numbers	☐ Solving novel situations where there is limited data available for training
☐ Detecting anomalies	☐ Making ethical decisions
☐ Generating custom text, images, or music	☐ Giving users manual control is a better user experience
	☐ A rule-based solution will do the job

If leveraging AI makes sense for your product, let's think about:

Who is your user? Who is not your user? How will they both be impacted?

Who might be indirectly impacted? (e.g., individuals, communities, organisations, society, planet)

How have you consulted marginalised populations and public safety groups?

Have you chosen the appropriate level of automation for the task?

- Will the user have complete control over how to proceed with the prediction?
- Will the user be presented with suggestions on how to proceed?
- Will the system choose how to proceed on behalf of the user?

How realistic is your solution?

Evaluate capacity before you start building to decide whether your team can implement your idea.

Do you have diverse team members who have the technical expertise to build, evaluate, and deploy AI products? List them.

Do you have subject matter experts who can commit to being a part of the entire product development process and share insights? List them.

How did you obtain your data?

How do you feel about dedicating significant time to experimental work, even if there's a chance it may not yield satisfactory results?

What are the applicable local and international laws and regulations, including those related to data protection and privacy, security, copyright, and intellectual property?

What are your success metrics?

Let's consider when your AI system will be good enough for people to use.

What is the **action** or **behaviour** you are trying to optimise? What are the possible outcomes?

What are the consequences of false positive and false negative predictions? Weigh the cost of these errors.

Next, consider how your model metrics translate to your product metrics. When choosing high-level product metrics, such as engagement, speed, or cost savings, consider the following:

List which metrics you will track. Include proxy metrics and counter metrics.

Do you have baseline measurements to compare against? What are they?

Are you able to slice your metrics across user subgroups? You need to know if your feature is benefiting all user types or negatively impacting some people.

How will you collect meaningful feedback? (e.g., through user surveys)

Who will be responsible for owning and reporting the metrics?

How can you mitigate risks?

AI is not perfect; it's probabilistic. You should expect your product to give users incorrect or unforeseen output at some point, and those consequences can have their own consequences, also known as *second-order effects*.

Plan to design your user experience around these error possibilities:

List the errors your users might encounter. For each, what will help the users move forward?

Will you explain the AI system's output to your users? How?

How will you give users the ability to intervene? (e.g., preview, review, edit, undo, dismiss, ignore, manually take over control)

How will you mitigate bias, avoid discrimination, and ensure fairness?

What ways can you allow users to provide product feedback and report issues?

How often will you monitor your performance metrics? What are your benchmarks?

Remember, the default reaction to poor AI system output doesn't always have to be to fix your AI model to get better results; you can make design changes to the user experience, too.