

Category

Data & AI

AI for Product Management & Roadmapping Playbook

A structured way for product managers to evaluate, prioritize, and roadmap AI-powered features without defaulting to "add a chatbot" or chasing every model release.

- Practitioner
- Intermediate
- Template Included
- Workshop Ready

- Mithun A. Sridharan
- 17 July 2024
- 8 minute(s)

Overview

A hands-on playbook for product managers deciding which AI features actually belong on the roadmap, using a four-lens evaluation framework, prioritization templates, and a worked SaaS example.

How is this different from a standard RICE or ICE prioritization framework?

RICE and ICE assume the main uncertainty is impact and effort. AI features carry two additional, often-larger uncertainties — whether the right data exists and what happens when the feature is wrong — that standard frameworks don't surface until you're already mid-build. The 4-Lens framework adds those as first-class scoring criteria.

What if we don't have an ML engineer to help score Feasibility and Data Readiness?

You can run a lightweight version with any technical stakeholder who can answer "does this kind of data exist, and can a reasonably capable model do this task today." The goal at this stage is a directional score, not a technical spec — bring in deeper ML expertise once an idea clears the initial gate.

Should every AI feature start in the assistive tier?

Not always, but it should be the default unless the task is low-stakes and well-understood (e.g., auto-formatting) where automated tier is safe from day one. The assistive-first approach mainly protects you against features with weak Responsibility scores, where a wrong automated action has real cost.

How often should we re-score the backlog?

Quarterly is usually the right cadence — frequent enough to catch improving data readiness or new evidence of desirability, infrequent enough that the scoring stays credible instead of becoming a constant re-litigation of the roadmap.

What do we do with ideas that score well on Desirability but poorly on Data Readiness?

Don't kill them — convert them into a data-prep task. A short, scoped effort to label or collect the missing data can turn a 2 into a 4 within a quarter, and those converted ideas often become the strongest releases because the underlying user problem was already validated.

```
{ "@context": "https://schema.org", "@type": "FAQPage", "mainEntity": [ { "@type": "Question", "name": "How is this different from a standard RICE or ICE prioritization framework?", "acceptedAnswer": { "@type": "Answer", "text": "RICE and ICE assume the main uncertainty is impact and effort. AI features carry two additional, often-larger uncertainties — whether the right data exists and what happens when the feature is wrong — that standard frameworks don't surface until you're already mid-build. The 4-Lens framework adds those as first-class scoring criteria." } }, { "@type": "Question", "name": "What if we don't have an ML engineer to help score Feasibility and Data Readiness?", "acceptedAnswer": { "@type": "Answer", "text": "You can run a lightweight version with any technical stakeholder who can answer \"does this kind of data exist, and can a reasonably capable model do this task today.\" The goal at this stage is a directional score, not a technical spec — bring in deeper ML expertise once an idea clears the initial gate." } }, { "@type": "Question", "name": "Should every AI feature start in the assistive tier?", "acceptedAnswer": { "@type": "Answer", "text": "Not always, but it should be the default unless the task is low-stakes and well-understood (e.g., auto-formatting) where automated tier is safe from day one. The assistive-first approach mainly protects you against features with weak Responsibility scores, where a wrong automated action has real cost." } }, {
```

```
"@type": "Question", "name": "How often should we re-score the backlog?",  
"acceptedAnswer": { "@type": "Answer", "text": "Quarterly is usually the right cadence —  
frequent enough to catch improving data readiness or new evidence of desirability,  
infrequent enough that the scoring stays credible instead of becoming a constant re-  
litigation of the roadmap." } }, { "@type": "Question", "name": "What do we do with ideas  
that score well on Desirability but poorly on Data Readiness?", "acceptedAnswer": {  
"@type": "Answer", "text": "Don't kill them — convert them into a data-prep task. A short,  
scoped effort to label or collect the missing data can turn a 2 into a 4 within a quarter, and  
those converted ideas often become the strongest releases because the underlying user  
problem was already validated." } } ] }
```

Subscriber access

Unlock this playbook

This playbook — including every framework, template, and step-by-step section — is available free to Think Insights subscribers. Enter your email to unlock it instantly and get our weekly insights newsletter. No account needed, and access is remembered on this device.

First Name

Last Name

Email

Country

I agree to data processing in accordance with GDPR

Leave this field blank

References

Related playbooks

Data & AI

Multimodal AI for Document, Image, Audio & Video Workflows Playbook

A framework for using multimodal AI across document, image, audio, and video workflows — building genuinely integrated processing that handles multiple content

- Practitioner

- Intermediate
- Template Included

[Read playbook](#)

Data & AI

AI for Sustainability, Climate Risk & Transition Planning Playbook

A framework for using AI in sustainability, climate risk, and transition planning — modeling transition pathway scenarios and climate risk exposure with genuine

- Executive
- Advanced
- Template Included

[Read playbook](#)

Data & AI

AI Storytelling, Internal Communication & Adoption Campaigns Playbook

A framework for AI storytelling, internal communication, and adoption campaigns — using specific, concrete narratives of genuine AI-driven outcomes rather than

- Practitioner
- Beginner
- Template Included

[Read playbook](#)

Talk to our AI product strategy team

Author

Mithun A. Sridharan

I'm Mithun A. Sridharan, Founder of this website - Think Insights - on Strategy, Management Consulting, Leadership, Digital Transformation, and Data Literacy. Follow me on social media or connect with me on LinkedIn for updates.