

Category

Data & AI

Data & MLOps Foundations Playbook (From POC to Production)

A concrete path for taking a data science proof of concept through the operational, monitoring, and governance gates required to run safely in production.

- Practitioner
- Advanced
- Template Included

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

Overview

A technical playbook for moving machine learning models from proof of concept to production, covering monitoring, governance, templates, and a worked churn-model example with real deployment numbers.

Do we need a full MLOps platform to follow this playbook?

No. The five layers describe capabilities, not specific tools — a small team can implement a minimal version with open-source components (a Git-based feature store, MLflow for the registry, a simple statistical drift check) before investing in an enterprise platform. Start minimal and add tooling as the number of production models grows.

How is this different from Google's or Microsoft's MLOps maturity models?

Those models describe an organization's overall automation maturity (level 0 through 2). This playbook is narrower and more actionable: it's a per-model checklist you can run this week against a specific POC, regardless of where your organization sits on the broader maturity curve.

What's the difference between data drift and prediction drift, and do we need to monitor both?

Data drift is a shift in the input feature distributions (e.g., customer age, usage patterns); prediction drift is a shift in the model's output distribution. Monitor both — data drift often precedes a business-metric problem by weeks, giving you an early warning prediction drift

alone would miss.

How do we decide a model's risk tier?

Ask what happens when the model is wrong: a wrong product recommendation is low risk (minor inconvenience), a wrong fraud flag is medium risk (customer friction, potential revenue loss), and a wrong credit or hiring decision is high risk (legal, regulatory, reputational exposure). High-tier models need explainability and human review built into the Governance layer from day one.

How often should we retrain a production model?

It depends on how fast the underlying data changes, but a good default is a scheduled retrain (monthly or quarterly) as a floor, plus a drift-triggered retrain whenever monitoring crosses your defined threshold — relying on schedule alone will miss faster-moving shifts like the login-pattern example above.

```
{ "@context": "https://schema.org", "@type": "FAQPage", "mainEntity": [ { "@type": "Question", "name": "Do we need a full MLOps platform to follow this playbook?", "acceptedAnswer": { "@type": "Answer", "text": "No. The five layers describe capabilities, not specific tools — a small team can implement a minimal version with open-source components (a Git-based feature store, MLflow for the registry, a simple statistical drift check) before investing in an enterprise platform. Start minimal and add tooling as the number of production models grows." } }, { "@type": "Question", "name": "How is this different from Google's or Microsoft's MLOps maturity models?", "acceptedAnswer": { "@type": "Answer", "text": "Those models describe an organization's overall automation maturity (level 0 through 2). This playbook is narrower and more actionable: it's a per-model checklist you can run this week against a specific POC, regardless of where your organization sits on the broader maturity curve." } }, { "@type": "Question", "name": "What's the difference between data drift and prediction drift, and do we need to monitor both?",
```

```
"acceptedAnswer": { "@type": "Answer", "text": "Data drift is a shift in the input feature distributions (e.g., customer age, usage patterns); prediction drift is a shift in the model's output distribution. Monitor both — data drift often precedes a business-metric problem by weeks, giving you an early warning prediction drift alone would miss." } }, { "@type": "Question", "name": "How do we decide a model's risk tier?", "acceptedAnswer": { "@type": "Answer", "text": "Ask what happens when the model is wrong: a wrong product recommendation is low risk (minor inconvenience), a wrong fraud flag is medium risk (customer friction, potential revenue loss), and a wrong credit or hiring decision is high risk (legal, regulatory, reputational exposure). High-tier models need explainability and human review built into the Governance layer from day one." } }, { "@type": "Question", "name": "How often should we retrain a production model?", "acceptedAnswer": { "@type": "Answer", "text": "It depends on how fast the underlying data changes, but a good default is a scheduled retrain (monthly or quarterly) as a floor, plus a drift-triggered retrain whenever monitoring crosses your defined threshold — relying on schedule alone will miss faster-moving shifts like the login-pattern example above." } } ] }
```

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 data & MLOps 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.