

Principles in Action Playbook: Post-Deployment

What do you do once people start using your AI?

- Are you building user trust?
- Are you tracking user engagement?
- Are you monitoring the performance of your model?

Are you building user trust?

Your model has been trained; your product is being used by real people. What now?

Are you maintaining user trust by continuing to educate users about your AI system?

Do you inform users when there are significant changes to the AI model or UX that might impact predictions or recommendations?

Have you considered in-app announcements to notify users of major updates, especially active users who may benefit most from new features?

Are you aware of potential confirmation bias, noting that implicit feedback may limit users to commenting only on what they see rather than other potential needs?

Do you have a process for addressing user concerns that can't be resolved immediately, and using feedback to make product adjustments that prevent recurring issues for future users? Outline the process.

Are you tracking user engagement?

Have you conducted a funnel analysis to identify areas where users may drop off in defined paths within your product?

Are you tracking conditions where users are more likely to accept or reject a recommendation?

Are you observing trends and patterns over time?

Are you looking outside your product for feedback, including platforms like Reddit, X, Facebook groups, or app store reviews?

Are you actively monitoring for bad actors and anticipating potential misuse of your product?

Are you performing regular QA testing to ensure product integrity?

Are you regularly evaluating whether your success metrics should be updated as user interaction with the product evolves?

Are you monitoring the performance of your model?

Are you planning regular maintenance for your dataset, beyond the initial data quality analysis?

Are you comparing live data to training data to detect significant differences that could impact performance?

Have you assessed the trade-offs between the cost of retraining and the performance benefits?

Are you regularly monitoring product metrics and fairness metrics?



PLAYBOOK FOR RESPONSIBLE AI PRODUCT DEVELOPMENT: Post-Deployment

VECTOR INSTITUTE

Is someone assigned responsibility for monitoring the model post-deployment and reporting on its performance?

Are you sharing regular model performance reports, including details on improvements made, with stakeholders?

Are you allowing time for machine learning engineers and data scientists to run experiments and gather more data before deploying updates?

Are you staying updated on new regulations and changes in the AI landscape to maintain legal and ethical compliance?