

# The 90-Day Plan to Ship Your First AI Mobile Feature

By **Bhavik Sadiwala** · Staff Engineer Android | AI Agents & Automation · Procore Technologies

---

Most engineering leaders who want to add AI to their mobile product make the same mistake: they start with the technology and work backwards to the user. They evaluate LLM APIs, debate on-device vs cloud, and prototype features that look impressive in a demo but never ship. The plan I am about to share works in the opposite direction.

**Before Day 1: Answer this question — what does your user struggle to find or do in the app right now? Not what AI feature would be cool. What does your user actually struggle with that intelligence could solve?**

## Days 1-30: Foundation

### Baseline your metrics

Before any AI work, measure how long it takes users to complete the task your AI feature will address. What is the current completion rate? What is the drop-off point? You need a before number to prove a result.

### Adopt AI dev tools

Cursor, GitHub Copilot, and Augment Code. This has nothing to do with your AI feature — it makes your team faster so the feature development itself moves quicker. Do not skip this step.

### Find the right technical lead

One engineer who has shipped an LLM feature in a production mobile app is worth more than five engineers learning from scratch. The rest of the 90 days depends on this person.

## Days 31-60: Build

### Ship to internal users first

Your first AI feature should not launch to all customers. Launch it to your own team. Real usage in a low-stakes environment surfaces problems that staging never catches.

## Iterate on response quality

The technical integration is usually the fast part. Collect the queries your internal users are sending. Identify the 20% of query types that produce wrong answers. Fix those before external launch.

## Days 61-90: Scale

### Staged rollout

5% of users first. Monitor: do users complete the task faster? Do they use the feature more than once? What queries are they sending that you did not anticipate?

### Measure and communicate

Compare your Day 1 baseline metrics to current performance. These numbers are your story to stakeholders and the foundation for the next AI feature investment.

## The One Mistake That Derails Most Teams

Trying to build the perfect AI feature instead of shipping a good one. The first version of our AI Mobile Assistant handled about 60% of queries well. We shipped it. Within two months, accuracy exceeded 90% because real usage data is more valuable than any amount of pre-launch testing. A shipped feature at 60% accuracy that learns from real usage will beat a pre-launch feature at 90% accuracy in a sandbox within three months.



*Bhavik Sadiwala* is a Staff Engineer Android at Procore Technologies, specialising in AI-powered mobile features and automation. [bhaviksadiwala.com](https://bhaviksadiwala.com) | [LinkedIn](#)