5 Things I Learned Leading AI Integration on a Global Mobile Team

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

Procore builds construction management software used on job sites across the world. When we started integrating AI into the mobile product, we made assumptions that turned out to be wrong. Here are the five most useful things I learned.

1. The Environment Is the Feature

We built the AI voice agent in our Bengaluru office. We tested it in meeting rooms. It worked well. We shipped it to beta users on construction sites. It did not work well. Construction sites are loud — wind, heavy equipment, power tools, other workers nearby. The feature is not just the code. It is the code working in the real environment. Test in conditions that match real usage before you call the feature done.

2. Accuracy at 60% Ships. Accuracy at 100% Does Not.

We had an internal debate before launching the AI Mobile Assistant. Should we wait until accuracy is high enough that no user is confused? Or ship and learn? We shipped at roughly 60% accuracy. Within eight weeks, accuracy exceeded 90% — because we had thousands of real queries to learn from. The risk of shipping early is some users have a bad experience. The risk of shipping late is your competitors ship first.

3. The Prompt Is the Product

Engineering teams think of the LLM as infrastructure and the prompt as a configuration detail. The reality is the opposite. The LLM is infrastructure. The prompt is the product. A poorly designed prompt produces responses that are technically correct but not useful. A well-designed prompt makes the same model feel built specifically for your product.

We have a prompt engineer on our AI feature work. Not an ML engineer, not a data scientist — someone who thinks deeply about language, user intent, and how to specify desired output in a way the model reliably follows. This role should exist in more mobile teams.

4. The Skeptical Engineers Become the Strongest Advocates

When we first introduced AI coding tools, the senior engineers were the most skeptical. Six months later, the same engineers have gone deepest with the tools. Their skepticism was healthy — they tested seriously before trusting, and the trust they have now is earned. Forcing adoption does not work. Creating conditions for genuine exploration does. Give skeptical engineers real problems to solve with AI tools, not toy examples. Let them discover the value themselves.

5. The Moat Is Not the Technology

A competitor can build a voice agent with Gemini Live API in months. What they cannot reproduce quickly is the domain knowledge encoded in our prompts — the specific patterns of how construction project managers query data, the edge cases we have tuned for, the user trust built over months of usage. Technology is the starting point, not the finish line. The moat is iteration cycles, domain knowledge, and user trust.

Bhavik Sadiwala is a Staff Engineer Android at Procore Technologies, specialising in AI-powered mobile features and automation. bhaviksadiwala.com | LinkedIn