

SAHIL SAPRA

Charlotte, NC 28273 | (216) 848-7236 | sahilsapra391@gmail.com | sahilsapra.com | [LinkedIn](#) | [GitHub](#)

PROGRAM MANAGER | SCALING OPERATIONS, WORKFLOW AUTOMATION & AI SYSTEMS

August 20, 2026

Duolingo Hiring Team

I am applying for the Senior Localization Program Manager role in Pittsburgh. The honest part first: I have not run a localization program, and I am not going to pretend the language side is already mine. What I have done, four years running, is take high-volume manual work that did not scale and turn it into a system that did, which is the part of this posting that reads like my actual job.

The clearest example is a 75-customer migration off competing platforms onto our next-generation product at CAMP Systems. Onboarding one enterprise customer took over two weeks of manual data work. Seventy-five of those, run one after another, is roughly three years. That is not a schedule problem, it is an arithmetic one. Rather than ask for headcount I broke the pipeline into its real steps and automated the data injection, taking capacity to three customers per day. The program only existed because the process changed first, and all 75 moved with zero operational downtime and zero churn on a forced platform change.

Getting there meant working through teams that did not report to me. Before anything moved I validated every workflow against the legacy system and found gaps that would have broken live airport operations for specific customers, so I ran a feature parity program, translated those gaps into prioritized requirements and acceptance criteria for engineering, re-sequenced the migrations, and told leadership and customers about the slips before anyone asked. Influence without authority is not a phrase I picked off your posting. It was the only tool I had.

On owning the numbers: I set the reporting for that program and for the products under it. I also built the analytics people actually used. Our second-largest account had a CEO receiving 25 separate weekly revenue emails from 25 airport locations, so I built a dynamic PowerBI digest with cumulative revenue and per-location drill-down, verifying every figure against raw Snowflake data before it reached his inbox. Twenty-five messy emails became one verified digest at zero engineering cost.

On tech-savviness, I build rather than review. I founded SpecHawk, an AI-powered deploy-aware no-code agentic QA platform, raised \$100K pre-seed, and run four design partners on weekly loops against real usage. Inside it I grounded an AI test case generator in retrieval-augmented generation over roughly 2,000 of our own ticket-to-test conversions, so output followed house conventions and improved with each example instead of needing retraining. At CAMP I drove adoption of LaunchDarkly and, when its dashboard did not fit how we worked, wrote a web application against their API for release control, which cut team OPEX 7%.

What I would be learning: CAT tools, translation management systems, quality frameworks for linguistic review, and the internationalization engineering side, Unicode handling and message format. I would rather name that half plainly than imply I have shipped 30 UI languages. The half I bring is the one that is harder to teach, owning a program end to end and building the operating model that keeps it from breaking as volume grows. I am happy to be in Pittsburgh and would welcome the chance to walk through how that migration was really run.

Sincerely,
[Sahil Sapra](#)