

Piyush Vyas

AI product, engineering & delivery

work@piyushvyas.com · piyushvyas.com · github.com/pyyush

EXPERIENCE

Agentic Intelligence LLC

Founder · Jun 2026-Present

- Founded Agentic Intelligence to design and deploy AI agents for customer-facing and operational workflows.
- Build browser-use, computer-use, and conversational agents across voice, messaging, email, and web.

Amazon Web Services (AWS)

AI Delivery & Engagement Lead · Sep 2024 - May 2026

- Led 40+ AI/ML and generative AI engagements across AMER and EMEA, serving startups through Fortune 500 companies in financial services, software & internet, healthcare & life sciences, and automotive.
- Led multidisciplinary teams of consultants, partners, account teams, and customer engineers through discovery, architecture, prototyping, evaluation, and production readiness, managing scope, milestones, delivery risk, and executive communication.
- Led 25+ executive workshops with C-suite, business, and engineering leaders to prioritize use cases, define technical and business KPIs, and turn broad AI goals into practical delivery roadmaps.
- Designed and prototyped agentic systems, including browser-use, computer-use, and conversational agents, alongside RAG and computer vision; turned recurring patterns into reusable architectures and accelerators for AWS teams and partners.
- When customers moved to production, partnered with AWS Professional Services and customer production teams to support deployment, then led adoption, enablement, and stakeholder alignment through rollout.

Numerade Labs, Inc.

AI Product & Engineering Lead · Apr 2023-May 2024

- Partnered directly with the CTO and Head of Product to shape the roadmap and technical direction for Ace AI Tutor and Numerade's AI video platform, turning student behavior, educator feedback, and business goals into product priorities.
- Built and led a six-person ML team - 4 engineers and 2 interns - setting technical direction, managing delivery, and promoting one engineer to senior within 12 months.
- Architected Ace as an agentic tutoring system using GPT-4, tool and function calling, image generation, and sandboxed Python code execution; built an automated grader and feedback loop to improve prompts, tools, and answers.
- Added visual problem-solving to Ace, driving 45% higher student engagement, more follow-up questions, and stronger retention.
- Built a multimodal video-generation pipeline using image generation, ElevenLabs text-to-speech, and HeyGen avatars, producing videos near-instantly at roughly one-fifth the cost of educator-created content; rollout across 50+ academic partners reached 90%+ adoption and contributed to growth in paid users.

EXPERIENCE / CONTINUED

Spectrum Labs, Inc.

Senior Machine Learning Scientist · Oct 2021-Apr 2023

- Built a voice toxicity detection product from the ground up - created a proprietary labeled audio dataset and trained voice-native models to identify abuse, bullying, hate, and other harmful behavior in live conversations.
- Worked with the CTO and engineering team to turn the research into a product sold to multiple customers, then expanded it from English to 12+ languages using multilingual speech models, including XLS-R 53.
- Took ownership of text moderation, moving the stack from classical classifiers to transformer models - training custom models from scratch, continuing pretraining on proprietary data, and deploying them through ONNX.
- Grew into the technical lead across voice and text moderation, working with the CTO and VP of Engineering on model strategy, evaluation, production integration, and the handoff from research to engineering.
- Joined customer and sales calls as the AI lead, translating trust and safety requirements into product capabilities and helping the team win and expand customer deployments.

Comcast Applied AI

Research Intern (NLP) · Aug 2021-Oct 2021

- Built a wav2vec 2.0-based speech recognition system for Comcast voice products, achieving a 30% reduction in word error rate for experiences serving millions of users.
- Co-authored research on temporal early exiting for low-latency streaming speech-command recognition, published at IEEE ICASSP 2022.

RESEARCH AWARDS

Best Student Paper Award · INTERSPEECH 2021

Optimally Encoding Inductive Biases into the Transformer Improves End-to-End Speech Translation

Coauthored with Anastasia Kuznetsova and Donald S. Williamson.

[Award record](#)

Outstanding Student Paper Award · IEEE ICASSP 2021

An End-to-End Non-Intrusive Model for Subjective and Objective Real-World Speech Assessment Using a Multi-Task Framework

Coauthored with Zhuohuang Zhang, Xuan Dong, and Donald S. Williamson.

[Award record](#)