

John Wu

Seattle, WA • johnwu0604@gmail.com • 206-536-9853 • linkedin.com/in/johnwu10

PROFESSIONAL EXPERIENCE

Microsoft — Senior Product Manager, Microsoft Foundry

Oct 2024 – Present

- Grew fine-tuning job creation conversion **+220%** (5%→16%), cut user error rate **71%** (17%→5%), and **doubled weekly active users** (4K→8K) by redesigning the creation flow and fixing reliability defects.
- Enabled **enterprise capabilities** (virtual network support, DNS support, customer-managed keys, managed identities) to unblock Foundry GA and allow enterprise customer migrations.
- Launched serverless fine-tuning for four open-source models (Mistral, Qwen, Llama, gpt-oss), coordinating external model partners (Meta, Mistral), Legal (CELA) on open-source licensing, and Responsible-AI review for compliance.
- Maintained **95%+ uptime** of allocated customer GPU capacity (peaking at 98.8%) through proactive deficit monitoring, automated leaked-node recovery, and capacity coordination with Azure HPC.

Domino Data Lab — Senior Product Manager

Jan 2023 – Oct 2024

- Built an MLflow-based experiment tracking feature, which has been adopted by **30%+ of the total customer base**.
- Delivered **ML pipeline** capabilities, which enabled workflow orchestration and contributed to **\$3M+ in net new ARR**.
- Released an **AI gateway** for enabling secured connections to external LLM providers, including **OpenAI and Anthropic**.
- Enabled users to build custom Generative AI applications by releasing multiple code samples for fine-tuning open-source models, leveraging tools including **Hugging Face, PyTorch, Ray Tune, and NVIDIA NeMo Toolkit**.

Microsoft — Product Manager 2, Azure Machine Learning Platform

Sep 2019 – Jan 2023

- Built and optimized the admin experiences for managing compute, virtual networks, and identities in the Azure ML Studio UI, growing it from initial release to **over 30,000 monthly active users**.
- Improved data scientist productivity by enabling custom dashboards using tracked MLflow metrics, auto-logging of hardware utilization metrics, and serverless training, which contributed to a **30+ point increase in the net promoter score (NPS)** of the Azure ML service.
- Enabled **role-based access control** for data plane assets in Azure ML, which unblocked adoption for **1000+ enterprise customers** on Azure's managed account list.
- Integrated Microsoft's AI supercomputer into Azure ML, used internally by Microsoft Research teams for large-scale, globally-distributed training jobs that execute on Nvidia's largest GPUs.
- **Received promotion offer to Senior Product Manager** (as a retention offer).

Microsoft — Product Manager Intern, Azure Batch AI Service

May 2018 – Aug 2018

- Authored the design specification for adding hardware utilization metrics to the Azure Batch AI service, and created public tutorials for distributed training using **TensorFlow and Horovod**.
- Earned a returning full-time offer **one level above entry**.

EDUCATION

McGill University — B.Eng., Computer Engineering · Montreal, Canada

Sep 2014 – May 2019

- **Co-President, McGill Artificial Intelligence Society** (2017–2019) — created a 12-week intro-to-ML bootcamp (30 AI applications built) and organized McGill's first AI hackathon (150 participants, now hosted annually).
- **Cansbridge Fellowship** — \$6,000 award to 1 of 15 Canadian students (0.5% acceptance) for leadership and entrepreneurial capabilities.
- **Two-time 1st place**, McGill Engineering Design Competition (2016, 2018).