Hugo AFFATICATI

haffaticati@gmail.com
haffaticati.github.io
https://www.linkedin.com/in/hugo-affaticati/

Technical Experience

Microsoft Seattle, WA, USA

Senior Cloud Infrastructure Engineer **Sep 2025 – current**

- Delivered record-setting performance of 865,000 tokens/sec for LLAMA 70B Inference on a single NVL72 GB200 rack featured in Satya Nadella's Build keynote and acknowledged by NVIDIA SLT at annual conference GTC.
- Closed a 20% performance gap across 14 AI training models within three weeks—after six months of stalled progress—by leading 50+ engineers across Microsoft and NVIDIA and spearheading full stack root cause analysis and resolution.
- Led largest MLPerf Training v4.1 submission on H200 GPUs (512 GPUs), delivering 28% cost savings to customers (vs competitors with previous hardware generation) leading to a white paper publication with external reviewer Signal65.
- Produced competitive intelligence reports to Microsoft SLT analyzing accelerator performance and cost data used for partner strategy, investor confidence, and seven new $M+ customer contracts
- Mentored five full-time employees and two interns (accepted return offers) as Technical Lead for the workloads team

Cloud Infrastructure Engineer II **Mar 2024 – Aug 2025**

- Developed and launched the AI Benchmarking Guide, standardizing performance of cloud AI infrastructure (GPUs, network, storage, etc.) across hardware vendors by leading 15 engineers across three companies
- Reduced generative AI inference latency by 46% on LLAMA models by optimizing model configuration (memory, batch size, and KV cache, etc.) to fully leverage next-gen GPU architecture and reduce cost for enterprise customers
- Showcased Azure's leadership via two technical talks (NVIDIA GTC, SC24) and three on demand videos

Technical Program Manager II **Sep 2023 – Feb 2024**

- Set a scale record in large language model training on Azure Eagle (world's largest AI supercomputer) showcasing cutting-edge advancements in AI software and hardware performance on NVIDIA H100 GPUs
- Assessed Microsoft's latest H100-based virtual machines with NVIDIA benchmarks, including NeMo Megatron and MLPerf inference & training, to establish performance standards and total cost of ownership for customers in Generative AI
- Optimized performance for key AI clients, accelerating training and reducing costs up by six through latest code and software stack implementation using Python
- Influenced marketing strategies by creating 10+ technical blog posts and tutorials, accumulating over 20k views

Technical Program Manager I **Sep 2022 – Aug 2023**

- Executed the first-ever public proof of concept for scale training LLMs (GPT-3, 530B parameters) in the cloud using NVIDIA NeMo framework - catalyzing the global rush for AI infrastructure including OpenAI's move to Azure
- Reached sub-two-min BERT Model training by implementing flash attention mechanism with Stanford AI Research Group
- Headlined key conferences Microsoft Build (most attended technical talk) and SC22 about AI infrastructure capabilities

Program Manager **Aug 2021 – Aug 2022**

- Led Microsoft's MLPerf Training and Inference submissions by optimizing Linux and Python ML code on A100 GPUs demonstrating unmatched performance, latency, and accuracy on virtual machines
- Benchmarked four GPU generations across ML and DL models, showcasing 2X performance and cost gain per generation
- Decoded AI workloads through writing 15+ technical blog posts and documentation (approx. 35k views) for beginners

KARL Tech (startup) Bordeaux, France

Co-founder, Entrepreneur, CEO **Mar 2019 – Apr 2020**

- Coded color rendering software for online retail within three months of incubation at Station F (Europe leading incubator)
- Earned first prize in Innovation among 48,000 students at Paris-Saclay University for Technical potential
- Exhibited and pitched to decision makers at CES 2020 (world's leading tech tradeshow) with France's official delegation

Leadership Experience

Microsoft, *Chair of the ERGs for LGBTQ+ employees at Microsoft (Southeast and Azure Core)* **Mar 2022 – current**

- Managed a 13-person leadership team and a global six-person board with bi-weekly meetings and dedicated mentorship
- Tripled annual budget to $30,000 by starting strategic partnerships between Microsoft's ERGs (geo and business based)
- Built a community with quarterly morale events and monthly safe-space meetings to improve D&I culture at Microsoft
- Received Microsoft's Leadership Award for cultural impact in 2023 (highest membership growth and budget growth)

Education

Yale University, Graduate School of Arts and Sciences – MS in Applied Physics 2020-2021

Research: Magic State Fidelity for Quantum Computation Optimization with Prof. S. Puri
Improving precision for Optical and Quantum Electronics with entangled photons with Prof. P. Rakich
Adapting Maxwell's equation for a general laser theory with Prof. A. Stone, Deputy Director of Yale Quantum Institute

Paris-Saclay University, Institut d'Optique Graduate School – BS and MS in Engineering & Quantum Physics 2018-2020

BS GPA: 3.97/4.00 and MS GPA 4.00/4.00, Mentored by Nobel Prize-winning physicist Gérard Mourou
Relevant coursework: Atomic Physics, Quantum Mechanics, Fourier and Non-linear Optics, Experimental Research

PSL University, Université Paris-Dauphine – BS in Applied Economics 2019-2020

Bachelor in Applied Economics, condensed into one year, for students of Top Engineering Schools
Relevant coursework: Macroeconomics, Microeconomics, Econometrics, and International Economics