Ziyad Mir

ziyadmir.com
linkedin.com/in/ziyadmir

Small-town Canada to Silicon Valley

- Port Elgin, Ontario, Canada (pop. 8000)
- In 2009, **Walmart Supercentre opened**, bringing 100s of jobs*

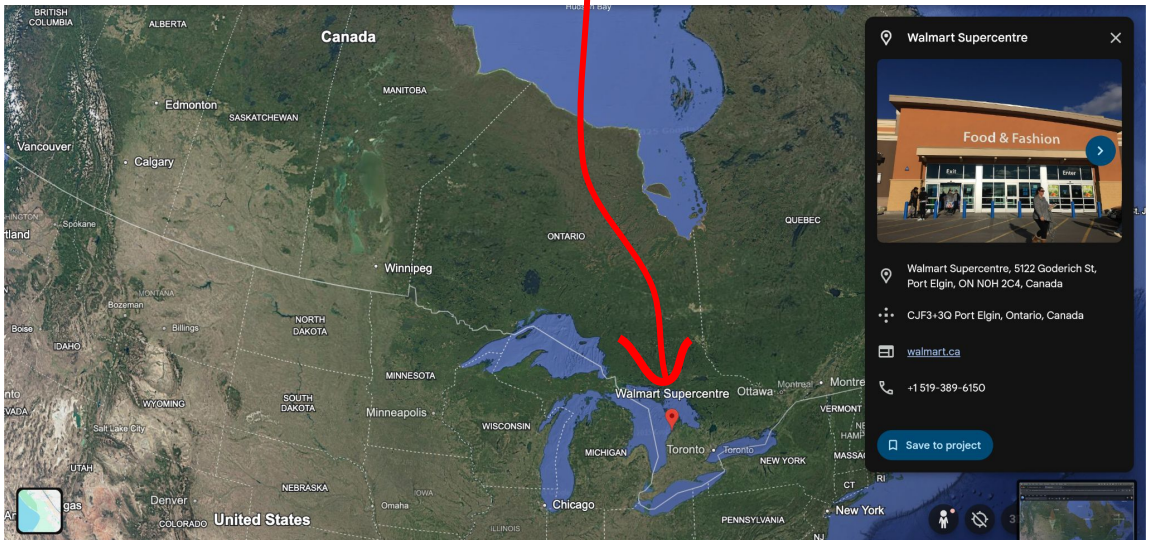

* One in particular holds special significance.

Scaling startups to big, bold Enterprise bets

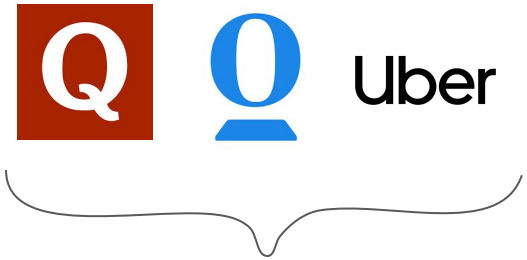

Scaling at startup pace and
pre deep-learning products

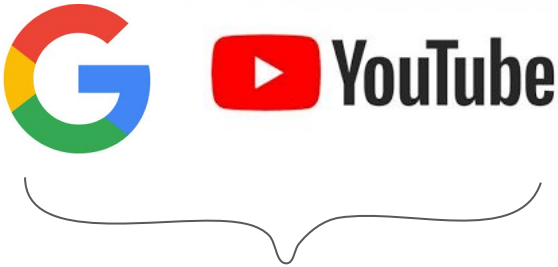

ML-enabled adtech
driving 500M+

Executive steering and
Company-level impact

“T-Shaped” Archetype

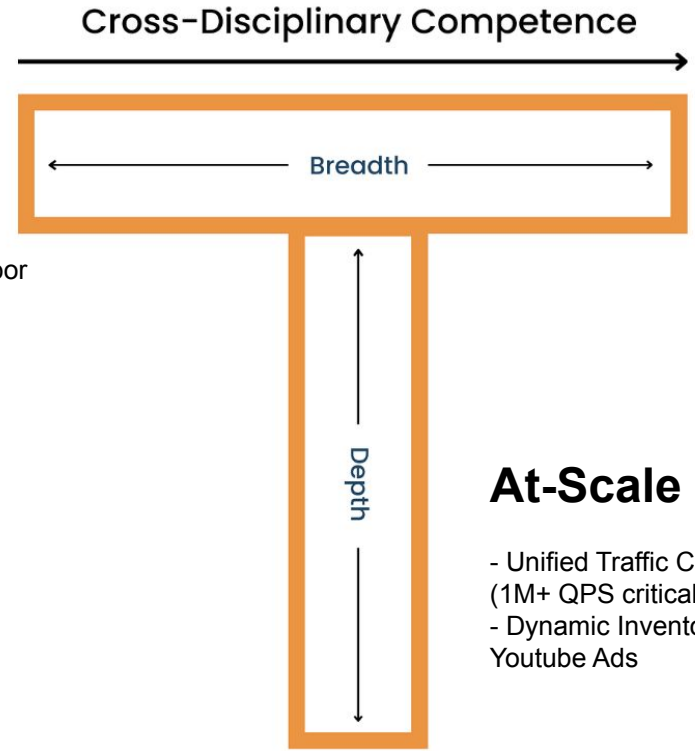

Product

- Growth at Quora
- Full-stack at Opendoor
- i18n at Uber
- Brand at Youtube

Sales

- Brand Management (Apple, Lays, etc.) at Youtube
- Vendor Procurement at Coupang

At-Scale Online Systems

- Unified Traffic Control/Global Rate Limiter at Coupang (1M+ QPS critical path)
- Dynamic Inventory Management and Fallbacks for Youtube Ads

Bet 1: What if we could model emotions?

- Win-win-win for Youtube **viewers, advertisers and the platform**.
- Thrive in a **cookie-less world.**
- Deliver **justified, high ROAS spend**.

Deepdive link: Buying Brand Identity at Youtube

Bet 2: What if you never had to work alone?

- I’m no longer a bottleneck; **making me feel enabled**.
- I’d work faster and at higher quality; **making my work feel more meaningful**.
- I’d take on “impossible” or “not worth it” projects; **inspiring and motivating me.**

Deepdive link: Deploying a Centralized and Secure Enterprise AI Platform

What if _____?

Walmart

What's the question that nobody is asking today?

Deepdive link: Global Governance @ Walmart, AI/ML Northstar Vision

Global Governance @ Walmart, AI/ML Northstar Vision

TL;DR -- 30% cut in risk-adjusted loss (≈$XXX M/yr) by **unifying Global Governance** (Physical, Digital, Financial, Legal, Supply-Chain & ESG, Cyber/Workforce) onto one **continuously improving AI platform**. **Crawl -> Walk -> Run** roadmap delivers value inside **24 months**, **upskills talent** and exits **high-cost vendors**.

Where we sit (H2 2025 snapshot)

Ranked No. 1 on the Fortune 500, Walmart is a global-scale giant nearing 1T in market capitalization -- and at Walmart scale, all risks are amplified. Global Governance sits centrally and acts as Walmart's immune system; derisking the existing business and unlocking future upside via compliant, international growth.

Over time, with growing complexity and aging systems, Global Governance itself faces internal execution risks; slowdowns in employee productivity, ballooning operating costs, and business demands outpacing delivery. Meanwhile, socioeconomic trends have led to a rise in vigilantism (e.g., 2024 United Healthcare shooting), increasing global security requirements across the workforce. Alongside, tough macroeconomic headwinds (high interest rates, abrupt tariff decisions, debt spiral and limited debt monetization options) limit the effectiveness and feasibility of traditional scale-through-headcount growth options (i.e., reduction in effective-ROIC on headcount spend).

Machine Learning sits perfectly positioned to help address these issues.

This paper focuses on a strategic vision for ML (and AI) for Global Governance at Walmart:

- Inefficiencies and Opportunities
- ML Framework: Towards a Single Global GovernanceAI Future
- Success Metrics and Notes

Deepdive link: Global Governance @ Walmart, AI/ML Northstar Vision

What if **Global Governance was a single AI?**

Walmart

- **Modeling**: Governance ultimately optimizes **expected risk-adjusted loss** under resource constraints.
  - The probability-weighted financial or reputational impact of every event across Global Governance domains + a tunable cost term. Expressing risk this way condenses all heterogeneous signals (e.g., dangerous object detection video models, escalate-or-ignore AML-models) into one optimisation target, enabling a single MoE-style “WalmartAI” to continuously learn company-wide policies while still honouring domain-specific constraints.

Deepdive link: Global Governance @ Walmart, AI/ML Northstar Vision

What if **Global Governance was a single AI?**

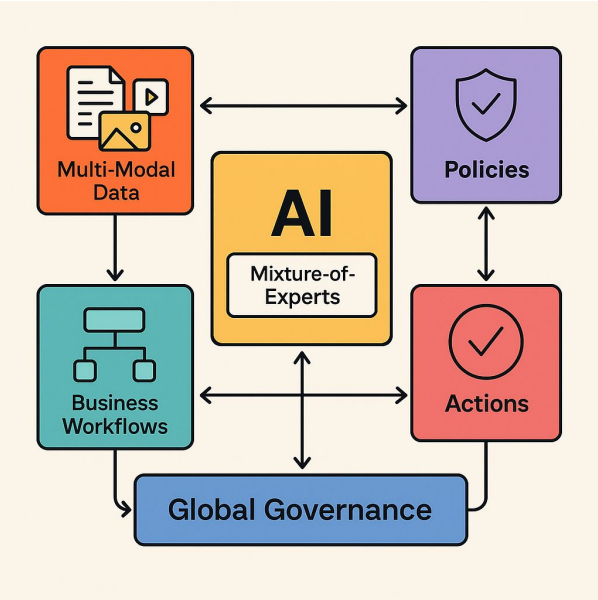

MoE/Mixture-of-Experts
Core architecture has **~30 expert cells:** six governance domains (Physical-Security, Digital-Fraud/AML, Cyber & Workforce, Supply-Chain & ESG, Legal, Product-Safety) and five data modalities (Vision, Text, Tabular, Graph, Time-series). Each cell shares a frozen, cross-modal backbone (T5-XXL + ViT-G) but **unfreezes the final two Transformer blocks and adds a LoRA head** so we can specialise rapidly without forgetting. Training runs in three passes:

RL Environment

- **States:** 15k-dimension vector from the MoE's fused embeddings (Vision + Text + Tabular + Graph + Time-series) X curated risk features (e.g., entity_age, sku_margin, CCTV crowd-density) guarantees partial observability is minimised; no raw PII leaves the sandbox.
- **Actions:** Discrete-set of hierarchically grouped: *Hard fences* -- block TX, lock account, recall SKU and *Soft controls* -- require MFA, increase audit sample rate, *Comms* -- human escalation, Slack push, Jira ticket mirrors real governance levers; keeps effective latency <= 500 ms
- **Transitions:** Real-time event stream *digital-twin* simulator allows billions of off-policy updates/day without touching prod
- **Rewards:** fn = - (A x Residual Risk) - (B x Opportunity Cost) - (C x Regulation Penalty) aligns with risk-adjusted EBITDA single objective

- **Modeling**: Governance ultimately optimizes **expected risk-adjusted loss** under resource constraints.
  - The probability-weighted financial or reputational impact of every event across Global Governance domains + a tunable cost term. Expressing risk this way condenses all heterogeneous signals (e.g., dangerous object detection video models, escalate-or-ignore AML-models) into one optimisation target, enabling a single MoE-style "WalmartAI" to continuously learn company-wide policies while still honouring domain-specific constraints.

Deepdive link: [Global Governance @ Walmart, AI/ML Northstar Vision]

A personal note, on why Walmart

Port Elgin Walmart

associate #7 was (my mother)
Arjamand Mir.

Currently active, 15+ years.

A personal note, on why Walmart

- Deployment tradeoff?
  - Edge, store-local, regional, etc.
- P/R tradeoff?
  - Acceptable risk and false-negatives?
- Escalation tradeoff?
  - Thresholds and wait-safe times?

Thoughtfully making these tradeoffs is hard.