PATRICK ORTELL

VP Engineering | CTO | Technical AI Product Manager | AI/ML Leader

Fort Collins, CO (Remote-First) | 720-229-5528 | patrick.ortell@arus.io LinkedIn | GitHub | arus.io

---

SUMMARY

Engineering leader and AI/ML practitioner with 14+ years shipping production systems — not demos, not decks, not prototypes. 3 exits. Teams scaled from 2 to 80+. Products from $0 to $50M ARR. I've been CTO, VP Engineering, Technical Product Owner, ML Engineer, and Fractional CTO across startups, enterprise AI platforms, and decentralized protocol teams.

What I actually build: Custom PyTorch object detection pipelines for spectral imaging. RAG systems ingesting 2,000+ academic papers as daily copilots. Langchain agents that convert natural language into cross-domain database queries. Predictive models integrated with environmental control systems and CFD-powered digital twins. Open-source systems for portable, vendor-agnostic LLM context stored on-chain. LoRA fine-tuning pipelines that duplicate domain experts from peer-reviewed literature. Multi-agent synthetic user research systems grounded in federal data. I've shipped every layer of the AI stack — from prompt engineering to model training to MLOps to production deployment.

What I actually deliver: 70% support cost reduction via LLM automation. 85% fewer production incidents through MLOps discipline. 30% ML model performance improvement. 100% on-time delivery across 8 consecutive quarters. The first AI chatbot in AgTech — a full year ahead of anyone else. Automated technical due diligence pipelines for Techstars. Privacy-preserving Rust security nodes that give users control over leaked PII.

What I do when nobody's paying me: Build AI/ML infrastructure for nonprofits through arus impact — OSINT pipelines for Bellingcat, press freedom intelligence for Reporters Without Borders, humanitarian early warning systems for climate-vulnerable communities. Mentor at Techstars (#GiveFirst, All Star Mentor 3 consecutive years). Youth hackathon mentor. Code for America. Open-source projects that are equal parts unhinged and useful.

I'm at my best when the problem is hard, the team is small, and the stakes are real.

Currently open to senior AI/ML leadership, Technical Product, or engineering leadership roles (remote-first).

KEY ACHIEVEMENTS

AI/ML & Product:

- **Shipped the first AI chatbot in AgTech** — a full year ahead of any competitor; designed prompts, reduced hallucinations, iterated on adoption data
- **Built custom PyTorch + OpenCV object detection pipeline** for spectral imaging cameras analyzing plant health (mold, disease, transpiration rates)
- **Architected a RAG system** ingesting 2,000+ academic studies on plant health and botany — deployed as a daily copilot for greenhouse growers via Azure GPT
- **Designed Langchain agents** (OSS + closed LLMs + NLTK) converting natural language into cross-domain database queries for greenhouse operations
- **Built predictive models + digital twins** integrated with PRIVA environmental controls and CFD models for real-time greenhouse optimization
- **Led development of on-chain LLM context portability** — vendor-agnostic AI conversations stored cryptographically on the Frequency blockchain (start in GPT, continue in Claude, work in Phi-4)
- **Designed tokenized data sharing framework** — user-approved, consent-based LLM context marketplace on decentralized infrastructure
- **Built Rust-based PII security nodes** — daemon processes giving users real-time control over unauthorized data transmission to brokers

Outcomes & Delivery:

- **70% support cost reduction** via LLM-powered automation (GPT-4, Langchain)
- **85% reduction in production incidents** through MLOps discipline (MLflow, CI/CD, observability)
- **30% ML model performance improvement** via pipeline optimization and training iteration
- **100% on-time delivery** across 8 consecutive quarters directing an 11-person distributed team
- **60% MTTR reduction** via Grafana/Prometheus observability framework
- **25% operational efficiency improvement** through technology introduction
- **20% cloud cost reduction** via infrastructure migration

Leadership & Scale:

- **$0 → $50M ARR**; grew engineering org from senior IC to VP level
- **3 successful exits** (SpexAI, biemedia, Pegcog/Techstars)
- **Teams scaled from 2 → 80+** across CTO, VP Engineering, and Technical Product Owner roles
- **IC → Director → VP → CTO** progressive leadership across multiple companies
- **95% team satisfaction** in weekly 1:1s; built high-trust remote-first engineering cultures
- **Built automated technical + product due diligence pipelines** for Techstars applicant evaluation

Open Source & Community:

- **4 OSS projects** — synthetic user research (Bureau of Synthetic Affairs), expert duplication via LoRA (Xerox Project), multi-agent PRD validation (PRD Executioner), ML-powered social network (Narwal)
- **11 nonprofit AI/ML infrastructure projects** via arus impact — OSINT for Bellingcat, press freedom intelligence for RSF, humanitarian early warning systems, children's digital rights
- **Techstars All Star Mentor** — 3 consecutive years
- **Code for America** FOCO chapter — 9+ years
- **Youth hackathon mentor** — Poudre School District, 9+ years

EXPERIENCE

Mentor | Techstars

Jun 2023 – Present | Boulder, CO (Part-Time, Hybrid) | #GiveFirst

Pro bono lead mentor for early-stage companies in the Techstars program — providing the full spectrum of fractional CTO/CPO advisory as a give-back to the founder community.

AI/ML & Technical Advisory:

- AI/ML model evaluation, framework selection, vector database architecture, and vendor assessment for early-stage companies building AI-powered products
- Hands-on guidance on LLM integration, prompt engineering, RAG architectures, and ML infrastructure design
- Architecture reviews and technology selection tailored to early-stage constraints

Product & Go-to-Market:

- Product-market fit validation, product discovery, and roadmap prioritization
- Compliance guidance including SOC2, DevOps best practices, and technical due diligence prep
- SDLC process design and agile implementation for small, fast-moving teams

Teams & Hiring:

- Technical hiring support — sourcing, assessment, and onboarding for engineering and ML roles
- Building high-functioning cross-functional teams and engineering culture from scratch

Automated Diligence for Techstars:

- Built automated technical and product due diligence pipelines for Techstars — grading program applicants across a structured evaluation matrix to streamline and standardize the admissions process

---

Founder & Technical Advisor | arus.io

Feb 2010 – Present | Colorado, United States (Remote)

Fractional CTO/CPO and technical advisor for early-stage startups. 16 years of engineering, product, and leadership experience distilled into high-impact advisory — a few hours a month, async-first, no chaos.

AI/ML & LLMs:

- LLM strategy and implementation (RAG, agents, prompt engineering, model selection)
- ML pipeline architecture, MLOps, and model deployment guidance
- AI product discovery — helping founders identify where AI creates real value vs. hype
- Evaluating build vs. buy for AI tooling and infrastructure

Product & Strategy:

- Product-market fit validation, product discovery, and roadmap development
- Requirements definition, success metrics, and adoption tracking
- B2B and B2C SaaS product strategy from zero to scale
- Fundraising support — technical due diligence prep, pitch deck review, investor Q&A coaching

Engineering & Architecture:

- Architecture reviews, code reviews, and technology selection
- Cloud infrastructure design (AWS, GCP, Kubernetes, Terraform)
- CI/CD, DevOps, observability, and SOC2 compliance guidance

- Technical debt assessment and migration planning

Teams & Leadership:

- Technical hiring — sourcing, assessment, and onboarding for engineering and ML roles

- Engineering org design, team scaling, and cross-functional team building

- Sprint planning, agile process implementation, and backlog management

- Founder sanity checks — the honest, low-ego perspective on what to build, what to skip, and what to fix

→ https://arus.io

Technical Product Manager | Project Liberty — Frequency Protocol

Nov 2022 – Nov 2025 (3 years 1 month) | Los Angeles, CA (Remote) | Decentralized Social Protocol

Technical Product Owner for Frequency, a decentralized social protocol. Led both core protocol delivery and AI/ML initiatives focused on data sovereignty, LLM context portability, and privacy-preserving infrastructure.

Decentralized AI Context & Sovereignty:

- Led development of an open-source system for storing LLM conversation contexts on the Frequency blockchain — enabling portable, vendor-agnostic AI interactions (e.g., start a conversation in GPT, continue in Claude, work in Phi-4) without vendor lock-in, fully cryptographically secured on-chain

- Solved cold-start and context sovereignty problems for LLMs by allowing users to own, port, and control their AI interaction history across any model or platform

Tokenized Data Sharing & User-Controlled Context:

- Designed B2B user-approved LLM context sharing framework — users could opt in to sell portions of their context data in exchange for tokens, creating a consent-based data marketplace on decentralized infrastructure

Privacy-Preserving Security Infrastructure:

- Directed development of genetic security nodes written in Rust — daemon processes running on infrastructure that gave users real-time control over leaked PII (e.g., social media data being passed to data brokers without consent)
- Enabled users to cryptographically block unauthorized data transmission, providing a decentralized alternative to traditional data broker pipelines

Core Protocol Delivery:

- Directed 11-person distributed team across 4 time zones building a Rust-based decentralized social protocol
- Achieved **100% on-time delivery** across 8 consecutive quarters, shipping **4 major protocol releases**
- Converted business and protocol requirements into clear product specs, workflows, and success metrics for 4 major releases
- Reduced mean time to recovery (MTTR) by **60%** through Grafana/Prometheus observability framework
- Tracked adoption and performance continuously
- Strong stakeholder management across engineering, product, and executive teams in a distributed environment

Stack: Rust, TypeScript, Svelte, AWS, Kubernetes, Terraform, GraphQL, LLMs (GPT, Claude, Phi), cryptographic protocols, on-chain storage

Machine Learning Engineer & Product Owner | SpexAI GmbH

Sep 2022 – Mar 2025 (2 years 7 months) | Germany (Remote) | Enterprise AI Platform, AgTech | **Acquired**

Technical Product Owner & MLOps Lead for enterprise AI platform serving agricultural customers. Led the full AI/ML product portfolio — from computer vision pipelines to LLM-powered copilots — across the entire greenhouse operations stack.

Computer Vision & Object Detection:

- Built and optimized a custom OpenCV + PyTorch object detection pipeline for proprietary spectral imaging cameras analyzing plant health in real time — detecting mold, disease, and measuring transpiration rates

- Improved ML model performance **30%** through pipeline optimization, training iteration, and model tuning

LLM Agents & Natural Language Interfaces:

- Designed and deployed custom Langchain agents using both open-source and closed LLMs combined with NLTK to convert natural language queries into cross-domain database queries — enabling growers, operations, and sales teams to interact with greenhouse data conversationally

- **Shipped the first AI chatbot/assistant in AgTech** — a full year ahead of any competitor — designed prompts, optimized for accuracy, reduced hallucinations, and iterated based on adoption and performance data

- Reduced support costs **70%** by building and launching LLM-powered automation (GPT-4, Langchain) — end-to-end from prompt engineering through deployment and monitoring

RAG & Knowledge Systems:

- Architected and deployed a RAG system ingesting 2,000+ academic studies on plant health, botany, and greenhouse science, served through an Azure GPT interface as a daily copilot for growers —

real-time guidance on pest identification, plant health diagnostics, and remediation techniques based on current best practices

Predictive Modeling & Digital Twins:

- Built custom predictive models that integrated plant health analytics and spectral imaging model outputs with PRIVA environmental control systems within greenhouses
- Enabled both real-time environmental optimization and digital twin scenario modeling using computational fluid dynamics (CFD) models for climate simulation and predictive greenhouse management

MLOps & Platform:

- Led product strategy and discovery for the enterprise AI platform, partnering with business stakeholders to identify high-impact problems and convert needs into product requirements and success metrics
- Cut production incidents **85%** through MLOps discipline (MLflow, monitoring, CI/CD)
- Ensured AI agents followed security, privacy, and data access guidelines across enterprise deployments

Stack: Python, PyTorch, OpenCV, GPT-4, Azure OpenAI, Langchain, NLTK, AWS SageMaker, Kubernetes, Terraform, MLflow, PRIVA, Vector DBs

---

Vice President of Engineering | Choi Holdings

Sep 2021 – Jun 2022 (10 months) | Los Angeles Metropolitan Area (Remote)

Led engineering for an NFT marketplace platform built on the Flow blockchain — from smart contracts to consumer-facing product.

- Managed distributed team of **18 engineers** across 4 time zones

- Architected and shipped NFT minting, marketplace, and wallet integration on the Flow blockchain using the Flow SDK
- Achieved **95% satisfaction rating** in weekly 1:1s; built high-trust engineering culture
- Partnered with business teams to align engineering priorities with product and business outcomes

Stack: Nest.js, TypeScript, React, Flow SDK, Vue.js, Swift, Kotlin, Terraform, AWS, Kubernetes

---

Chief Technology Officer | TiLT

Jun 2020 – Mar 2021 (10 months)

CTO for an FMLA leave management platform navigating one of the most complex compliance landscapes in HR tech — HIPAA, insurance provider integrations, payroll systems, internal HR workflows, and federal government reporting.

- Led team of **7 engineers**; owned technical roadmap and architecture decisions
- Architected integrations across HIPAA-compliant data pipelines connecting insurance companies, payroll providers, HR systems, and federal reporting requirements
- Reduced cloud costs **20%** by migrating from Heroku to AWS while maintaining HIPAA compliance throughout the transition

Stack: Node.js, TypeScript, GraphQL, React, AWS

VP of Engineering | Radar Relay

Apr 2018 – Jun 2020 (2 years 3 months) | Promoted from Senior Engineer → Director → VP

One of the earliest decentralized exchanges (DEX) on Ethereum — an open order book relayer built on the 0x protocol. Backed by Blockchain Capital, Breyer Capital, Digital Currency Group, SV Angel, and others. Wallet-to-wallet, non-custodial trading with no accounts, no deposits, no trading limits.

- Grew from IC to VP, managing **15 direct reports** building decentralized trading infrastructure
- Shipped **4 technical strategies** aligned with SEC compliance requirements — navigating novel regulatory territory for tokenized asset trading
- Built the engineering architecture for 0x protocol integration, smart DEX routing, multi-liquidity source aggregation (Uniswap, Kyber, Curve), and margin trading via dYdX
- Collaborated with SEC on compliance frameworks for decentralized asset trading — one of the first teams to operate a DEX within a regulatory compliance model

Stack: Node.js, TypeScript, GraphQL, Vue.js, React, Python, Golang, Ethereum, 0x Protocol, Hyperledger, AWS, GCP

Chief Technology Officer | TurboTenant

Oct 2016 – Apr 2018 (1 year 7 months) | Fort Collins, CO

CTO for a property management SaaS platform for independent landlords — now serving 900K+ landlords nationwide. Full rental lifecycle: listing syndication (28+ sites), tenant screening (TransUnion), online applications, lease generation, rent collection, and maintenance management.

- Managed **9 direct reports**; owned technical roadmap, architecture, and hiring
- Introduced 3 new technologies driving **25% operational efficiency improvement**
- Built the core platform infrastructure that still powers the product today

Stack: Node.js, React, Ruby, PostgreSQL, AWS

Senior Software Engineer → Director of R&D | CommercialTribe

Apr 2015 – Dec 2016 (1 year 9 months) | Denver Metropolitan Area

Director of R&D for a sales enablement and LMS SaaS platform — helping enterprise sales teams shorten ramp times and boost quota attainment through video-based practice, coaching, and certification workflows.

- Managed **12 engineers** across 3 time zones; achieved **90% sprint completion rate**
- Built the WebRTC-based real-time video platform powering the core product — enabling reps to record, submit, and receive coaching on sales pitches asynchronously
- Shipped React Native mobile application extending the platform to mobile-first sales workflows
- Led R&D on video analysis and coaching feedback features

Stack: WebRTC, React Native, React, Node.js, AWS

Software Engineer | Cadence & Cause

Dec 2014 – 2015 | Greater Denver Area

Ruby on Rails, real-time streaming.

Software Developer / Product Owner | Levels Beyond (Reach Engine)

Jun 2013 – Oct 2014 (1 year 5 months) | Denver Metropolitan Area

Media asset management platform. Elasticsearch, video pipelines.

Software Developer | biemedia (acquired)

Jul 2011 – Sep 2012 (1 year 3 months)

Ruby on Rails, video streaming. *Exit*.

Software Engineer | Techstars — Pegcog (acquired)

May 2011 – Dec 2011 (8 months) | Boulder, CO

Scala, D3.js, MongoDB. *Exit*.

Intern | Mozilla

Jan 2011 – Jun 2011 (6 months) | Remote

Completed 1 project integrated into the main product, improving user interface responsiveness by **5%**. Basecamp, GitHub, Jenkins CI/CD.

TECHNICAL SKILLS

AI/ML & Prompt Engineering

LLMs, PyTorch, TensorFlow, OpenCV, Langchain/LangGraph, A2A, NLTK, prompt optimization, hallucination reduction, AI assistant/agent design, RAG architectures, NL2SQL, computer vision, object detection, digital twins, CFD modeling, predictive modeling, MLflow, SageMaker/Bedrock, Azure OpenAI, Vertex, MLOps, Deep Learning, Neural Networks, NLP, Model Deployment, Data Engineering, Generative AI, ML pipelines, model optimization, AI context portability, data sovereignty

Product Management

Product discovery, requirements definition, success metrics, stakeholder management, workflow design, adoption tracking, PMF validation, technical due diligence, Agile/Scrum, Sprint Planning, Roadmap Development, Backlog Management, Product Ownership, User Stories, Jira

Languages

Python, TypeScript/JavaScript, Rust, Golang, Ruby, Scala, Swift, Kotlin

Frontend

React, Next.js, Svelte, React Native/Expo, Vue.js, AngularJS, D3.js, CSS

Backend

Node.js, Nest.js, GraphQL, REST APIs, PostgreSQL, MongoDB, Elasticsearch, Redis, Vector DBs, data pipelines, video pipelines, real-time streaming, WebRTC, logic flows, JSON

Infrastructure & DevOps

AWS (SageMaker, Bedrock, Lambda, ECS/EKS, S3, RDS, DynamoDB, CloudFront, Route53, IAM, CloudWatch), GCP (Vertex AI, BigQuery, Cloud Run, GKE), Microsoft Azure (Azure OpenAI), Kubernetes, Helm, Terraform, Docker, CI/CD, SOC2 compliance, Grafana, Prometheus

Blockchain

Ethereum, Hyperledger, Substrate/Rust protocols, SEC-compliant trading systems, cryptographic protocols, on-chain data storage, decentralized identity

Leadership

Org Scaling & Team Building: IC → Director → VP → CTO progression across multiple companies; scaled teams from 2 to 80+; up to 18 direct reports; 95% team satisfaction rating; built high-trust remote-first and async engineering cultures from scratch

Distributed Teams: Led teams across 4+ time zones; 100% on-time delivery across 8 consecutive quarters; 90% sprint completion rates

Technical Hiring: Full-cycle technical recruiting for engineering and ML roles — sourcing, assessment, onboarding; supported first AI/ML hires for Techstars portfolio companies

Exits & Scaling: 3 successful exits (SpexAI, biemedia, Pegcog/Techstars); scaled product from $0 to $50M ARR

Architecture & Strategy: Architecture reviews, technology selection, cloud migration strategy (20% cost reduction), operational efficiency improvements (25%), infrastructure modernization

Stakeholder Management: Cross-functional alignment across engineering, product, and executive teams; SEC and SOC2 compliance leadership; technical due diligence for investors and accelerators

Fundraising & Advisory: Technical due diligence prep, pitch deck review, investor Q&A coaching; built automated diligence pipelines for Techstars applicant evaluation

Mentorship & Community: Techstars All Star Mentor (3 consecutive years); Code for America FOCO chapter (9+ years); Poudre School District youth hackathon mentor (9+ years); Sox Place homeless youth mentor

> **Don't take my word for it** — these skills are endorsed by the people I've actually worked with. See it on <u>LinkedIn</u>.

OPEN SOURCE & NONPROFIT IMPACT

<u>arus impact</u> — AI/ML infrastructure for nonprofits doing real impact work. Data pipelines, ML tools, and integration middleware — scoped for mission, not margin. 501(c)(3) focused, pro bono.

Featured Projects (Top 5)

<u>Bellingcat — OSINT Investigation Infrastructure</u> Entity resolution, computer vision geolocation, and multilingual monitoring pipelines for open-source investigations supporting human rights accountability and war crimes documentation. *Tech: Entity Resolution, Computer Vision, NLP*

<u>Reporters Without Borders — Press Freedom Intelligence</u> Automated violation monitoring, Press Freedom Index pipeline acceleration, and Journalism Trust Initiative (JTI) certification tooling for the world's leading press freedom organization. *Tech: NLP, Classification, Data Pipelines*

<u>CyberPeace Institute — Cybersecurity Intelligence for the Humanitarian Sector</u> Scaling threat intelligence, volunteer matching, and cyber-resilience tooling from 240 NGOs to 10,000 — protecting the organizations that protect people. *Tech: Threat Intelligence, NLP, Matching Algorithms*

<u>Data-Pop Alliance — Humanitarian Early Warning Infrastructure</u> Flood prediction pipelines, geospatial analytics, and humanitarian data infrastructure for climate-vulnerable communities. ML that

saves lives. *Tech: Climate ML, Geospatial, Data Pipelines*

Fairplay — Children's Tech Accountability Tools Automated dark pattern detection, platform policy monitoring, and COPPA compliance scanning for children's digital rights advocacy. *Tech: Computer Vision, NLP, Platform Monitoring*

→ View all 11 projects

Open Source Projects (GitHub)

Bureau of Synthetic Affairs AI-powered synthetic user research system that intercepts anonymous Reddit confessions, cross-references them with federal government data (Census, BLS, ATUS, Consumer Expenditure), and manufactures 19-dimension synthetic humans to convene as a tribunal to validate product ideas before you build them. Replaces $15k focus groups with statistically-grounded synthetic personas that simulate real user behavior — complete with cognitive biases, trust issues, and app abandonment histories — enforced at the prompt level. *Tech: Python, Claude/Anthropic, BAML, Firebase, Reddit API, Census/BLS/ATUS APIs*

The Xerox Project Expert duplication system that ingests academic papers, clinical research, and domain expertise, then synthesizes LoRA-weighted personas that *learned* the field rather than pretending to know it. First subject: a burnout recovery specialist trained on 127 peer-reviewed papers and 43 clinical case studies — with 94.7% voice match. Not prompt engineering — actual LoRA fine-tuning that puts knowledge in the weights, not the instructions. Runs on your hardware for $0/query forever. *Tech: Python, Mistral 7B, LoRA/QLoRA, PyTorch, PEFT, BAML, Claude (synthesis), Web Search*

PRD Executioner Multi-agent system that generates 50–300 synthetic user personas, has them brutally review your PRD, and surfaces hidden conflicts through multi-agent debates — all in under 5 minutes. Features type-safe LLM outputs via BAML (99.9% reliability), LangGraph multi-agent orchestration, LangSmith observability, real-time streaming UI, and conditional agent debates triggered when sentiment exceeds conflict thresholds. *Tech: Python, TypeScript, Next.js, BAML, LangGraph, LangSmith, Claude, FastAPI, SSE*

Narwal Fully open-source social network built on user sovereignty and transparency. Features ML-powered video content labeling (replacing hashtags), serverless content moderation using Google Cloud NLP, Perspective API, and Video Intelligence API for automated toxicity detection and explicit content

filtering. Includes real-time video feeds, in-app purchases, ad management, and a web admin for manual content moderation. *Tech: React Native, Firebase, Google Cloud ML APIs, Perspective API, Video Intelligence API, Cloud Functions*

→ <u>View all repos</u>

EDUCATION

Metropolitan State College of Denver (Metropolitan State University of Denver)

Dual BSA, Technical Communications & Computer Science Grant writing for non-profits, Digital design | 2009 – 2013

Activities & Societies: Mhdivine Club, Collegiate Entrepreneurship Organization

- Tutor GED and College prep to inner city and at-risk youth for Mhdivine club — <u>http://mhdivine.com/</u>
- **Entrepreneur of the Year Award 2010**
- Volunteer work for Soxplace with homeless teens (2009) — <u>http://www.soxplace.com/</u>

CERTIFICATIONS & LICENSES

Certification	Issuer	Date
Amazon Web Services: Data Analytics	AWS	Feb 2020
Data Science on Google Cloud Platform: Architecting Solutions	Google	Feb 2020
AWS End User Computing Competency Partner	AWS	Dec 2019
Planning and Releasing Software with JIRA	Lynda.com	Sep 2018

Certification	Issuer	Date
Ethereum: Building Blockchain Decentralized Apps (DApps)	Lynda.com	Sep 2018
GraphQL Essential Training	Lynda.com	Feb 2018
Elasticsearch Essential Training	Lynda.com	Jan 2018
Blockchain: Beyond the Basics	Lynda.com	Dec 2017
Expert Scrum Master Certified (ESMC)	Mountain Goat Software	Jan 2014
M101JS: MongoDB for Node.js	MongoDB University	Aug 2013

HONORS & AWARDS

- **Techstars All Star Mentor** — 3 consecutive years (2023, 2024, 2025)
- **Entrepreneur of the Year Award** — Metropolitan State College of Denver (2010)

VOLUNTEERING & COMMUNITY

- **Code for America** — FOCO Chapter (Jun 2016 – Present, 9+ years) — Using technology to solve everyday challenges for nonprofits, schools, government, and grassroots organizations in Fort Collins
- **Youth Hackathon Mentor** — Poudre School District (Sep 2016 – Present, 9+ years) — Mentoring students building apps for mental health, nutrition, and physical activity
- **Mentor** — Sox Place (Jan 2008 – Sep 2009) — Drop-in center for homeless youth in Denver; meal prep, GED teaching, one-on-one counseling

RECOMMENDATIONS

Ben Niehaus — CTO, SpexAI

November 18, 2025 | Managed Patrick directly

> Patrick and I worked together at SpexAI. He got came recommended through the TechStars Network.
> And he came at the perfect timing. We needed a Technical Product Owner for our ML Pipeline giving
> me space to breath and focus more on the product roadmap and team development. Brainstorming
> sessions turned into get shit done sessions. He will not say what you want to hear, but need to hear.
> Forging a solution that is fit for the task with as little ego as possible. His first Product to market was
> the first ChatBot in a Ag Tech application, at least a year ahead of everyone else. Why, because he
> can cut through the noise. And probably equally important deliver high quality products efficiently
> and fast. I can only recommend working with Patrick in any AI, LLM, MLOps capacity.

Kenne Ives — Head of Product, Project Liberty

November 18, 2025 | Managed Patrick directly

> I had the pleasure of working with Patrick while he served as the Technical Product Manager at
> Project Liberty, and he was an exceptional partner across every dimension. Patrick brought a rare
> combination of deep technical expertise — especially in AI — and a broad command of web3,
> crypto, wallets, and blockchain systems. His knowledge of LLMs and emerging AI tooling was
> consistently ahead of the curve, and he was often the person I turned to for clarity on what was
> possible and what was coming next. Patrick played a critical role in driving the development of our
> Frequency L1 blockchain and several key partner integration tools. He approached every challenge
> with ownership, thoughtful problem-solving, and a strong grasp of agile best practices that kept the
> team focused and moving quickly. His ability to translate complex systems into clear product
> direction was invaluable. He's a deep thinker, a deeply technical leader, and one of the most
> knowledgeable people in AI I've worked with. Any team would be lucky to have him.

Larissa Licha — Product Director, AI Context & Memory @ Miro

November 17, 2025 | Worked with Patrick at different companies (Techstars)

> Patrick was one of our mentors at Techstars and easily the most impactful. He has this rare ability to immediately understand a problem and zero in on smart, scrappy ways to solve it. His deep expertise in ML/AI was invaluable as we were scaling our platform; he helped us stay focused, move fast, and maintain a high bar for quality. He even jumped in to support our first AI/ML hire, from assessment to onboarding.

Corbin Long — Product @ Amplica Labs

November 17, 2025 | Worked with Patrick on the same team

> I worked with Patrick for three years on a mission-driven, fast-moving product team. What consistently set Patrick apart was his ability to step into the unknown, especially as we explored the limits of large language models and privacy in AI. Most of our team came from blockchain backgrounds. Patrick brought in needed ML expertise at exactly the right moment. He broke down the ML landscape in simple terms, highlighted which tools and approaches were worth our time, and called out potential pitfalls before we wasted effort. When engineers or stakeholders were about to overbuild or chase a dead end, Patrick cut through the noise, offering grounded advice and pragmatic alternatives. What really impressed me: Patrick didn't just recite technical facts. He translated ML intricacies into choices that fit our product and resource constraints, helping us stay focused and move fast, even during pivots. Whether he was picking apart documentation or demoing Python prototypes to test assumptions, he kept the team anchored and self-sufficient. If you need someone who can make real impact in ambiguous, high-stakes environments, especially when your team's skills are uneven and the path forward keeps changing, Patrick is the person you want in your corner.

Krystallin Baker, CCMP — Creative Strategist & Brand Storyteller

Worked with Patrick on the same team

> Patrick is an incredibly talented professional, who not only delivers industry expertise to projects, but also contributes to an inclusive work culture. With a creative mindset and an ability to think outside the box, working with Patrick is a true privilege.

Robert Schachte — Co-founder @ HYVV

August 5, 2021 | Worked with Patrick on the same team

> Patrick is one of the most talented CTO's I've ever met. He can be anything you need him to be. Architect, developer, Investor experienced, product shepherd, teacher, organizer, leader, phenomenal father, manager, friend, heck he can even frame a house. A bottomless well of skills. I believe he's the best CTO in Colorado and the first stop for me on any technology endeavor.

Paul Ironside — CEO | Revenue-Led Operator

January 26, 2020 | Senior to Patrick

> Among the myriad of professional accolades that I'll highlight next, Patrick has a unique character mix between setting / maintaining high standards and an overall roll with it attitude that positively imprints on a team. Integrity, attention to detail, commitment to deadlines, quality are a handful of traits that represent Patrick. "Hey...come take a look at this." was always something I enjoyed hearing from Patrick Monday morning to show me something he was working on over the weekend (for fun) that furthered the business.

Charisse Bowen, PMP — Vice President, Federal Services

February 6, 2018 | Patrick's client

> Patrick is a total pro. There is no project too big or too small for this guy. He is an expert when it comes to Node.js React.js and AWS. Even more than a rockstar engineer. He always goes above and beyond and delivers on time (usually more than what I've asked for). He also has a huge heart and is incredibly giving of his time to better his community at large. I would jump at the opportunity to hire Patrick again in a heartbeat.

Paul Ronto — Worked with Patrick on the same team

February 5, 2018

> Patrick is a visionary when it comes to solving difficult tech-related problems. He understands the needs of the users and the limitations/opportunities surrounding how to build an effective solution. I've continually been impressed with his forethought and intention that goes into product planning and how he manages his team of developers. It's not often you find a DEV that puts this much thought into the how and why a product is needed, and in the SaaS world that's key, without it products get built that either don't work, or don't solve the specific problem they were designed for. I'd be excited to have Patrick running my DEV team no matter what the product or the problem, I've learned a ton from him and can confidently expect top notch work every time.

Keith Kacsh — Principal DevOps Engineer

February 28, 2015 | Worked with Patrick on different teams

> Patrick is a true champion of ideas and creative innovation. He uses multiple technologies to create new and innovative applications and designs. I have personally worked with him on multiple projects to help design and implement the infrastructure needed for his visions. He has a true desire to learn new technologies and then use them for the best possible products. With a high level of knowledge and understanding of web based technologies, and a side of creative genius, he makes an excellent addition to any organization.