Northvalley Intelligence

Software for everyone.

CAPABILITY STATEMENT

Capabilities

Northvalley Intelligence builds software for organizations that need it to work with the systems they already run: web and mobile applications with custom integrations, agentic applications, and data science, machine learning and data cleansing.

Software development. Web and mobile applications plus custom integrations with the systems a buyer already runs.

Agentic applications. Software agents that take actions inside a business's own systems, with voice and chat front ends and agent-callable backends (MCP).

Data science, machine learning and data cleansing. Data cleansing, entity normalization and matching, ontology, ranking and search relevance, and machine learning models, on messy real-world data.

Differentiators

- Entity recognition and natural-language processing on messy real-world data: normalization and matching systems at CareerBuilder, a product ontology and machine learning ranking at The Home Depot, and catalog search and discovery with generative AI at Genuine Parts Company.
- Data cleansing as a repeatable method: profile, normalize, match and measure, with an evaluation set agreed before the cleanup starts.
- Agent-callable delivery: systems built so a customer can reach them from the assistants they already use. See the live demo under Capabilities.
- Enterprise release discipline at small-business scale: a staging environment, scenario tests and CI-gated deploys on every engagement, with the founder (Ph.D., 22 years in software) doing the hands-on work.
- Models can be trained and run on local hardware, so data stays in-house where a client needs it.

Past performance

Canon Insurance Advisers

Client: Canon Insurance Advisers, LLC (Georgia insurance agency).
Role: direct engagement, founder as sole developer.
Period: July 2026 to present (currently active).
Situation: a Georgia insurance agency quoted each home and auto prospect by hand, re-keying the same details for several hours per quote.
Task: cut quoting time without replacing the agency's existing tools.
Action: a one-page intake form feeds an automated quoting run that completes the rating steps unattended, with guard rails that never bind, purchase or order a report; an earlier phase added a retention board that flags policies at risk of lapsing before they end.
Result: several hours of manual quoting became a few seconds of form filling and a few minutes of automated quote generation.
Reference: available on request.

Smart Cash Sheet

Client: Smart Cash Sheet.
Role: ongoing engineering.
Period: September 2026 to present (currently active).
Situation: a personal cash-flow app on the web with iOS and Android apps, run by a solo founder, shipped about once a month, each release carrying weeks of accumulated change and risk.
Task: take over the engineering and ship faster without breaking production.
Action: a staging environment, scenario tests and a three-phase release process; bank-data import, subscriptions and billing, and the safety rules around the in-app money coach.
Result: the monthly cycle became two production releases in one week, turning the founder's roadmap into shipped features faster.
Reference: available on request.

COMPANY DATA

Legal name
Northvalley Intelligence LLC

Entity
Georgia limited liability company, founded June 2026

Location
Marietta, Georgia

NAICS (primary)
541511 — Custom Computer Programming Services

NAICS (secondary)
541512 — Computer Systems Design Services

PSC
DA01 — IT and Telecom – Business Application/Application Development Support Services (Labor)
DH01 — IT and Telecom – Platform Support Services: Database, Mainframe, Middleware (Labor)
DD01 — IT and Telecom – Service Delivery Support Services: ITSM, Operations Center, Project/PM (Labor)

Contracting role
Prime contractor or subcontractor

Unique Entity ID (UEI)
JX76HA57NCJ8

CAGE code · SAM.gov
26B16 · registration active

HOW AN ENGAGEMENT RUNS

1. Conversation with the founder.
2. Written scope and acceptance criteria.
3. Build in short increments on a staging environment.
4. Release through CI, with documentation and handover.

COMPANY CONTACT

Ferosh Jacob, Founder
ferosh@northvalleyintel.com
(470) 781-4143
northvalleyintel.com/government
Marietta, Georgia

Key personnel

Ferosh Jacob, Ph.D. — Founder and technical lead

Northvalley Intelligence's work is performed by its founder. Before founding the company in June 2026 he spent 22 years in software, moving from software engineer to senior engineering manager leading cross-functional teams of more than 40 engineers, managers, data scientists and machine learning engineers, across the retail, healthcare, recruiting and automotive-parts domains.

Senior Engineering Manager July 2024 – June 2026
Genuine Parts Company, Atlanta, GA
Led a cross-functional domain of four squads on catalog, search and discovery, using generative AI and classical machine learning. Built a first-of-its-kind agent-based auto-parts assistant on an Agent Development Kit, working with UX and platform teams through A/B testing and production rollout.

Senior Engineering Manager January 2021 – June 2024
Kaiser Permanente, Atlanta, GA
Led a cross-functional team of managers, data scientists, machine learning engineers and software engineers responsible for search relevance and platform performance for internal and external clients. Directed the migration of KP.org search from a vendor product to a homegrown open-source solution, transitioning the endpoints with minimal disruption to other services. Two refereed papers followed on automated search-relevance evaluation (AMIA 2022; CIKM 2022).

Staff Software Engineer September 2015 – January 2021
The Home Depot, Smyrna, GA
Led data-driven initiatives in search relevance, product ranking and ontology development for the online store: machine learning ranking models, a search-evaluation framework that measured feature impact before deployment, and a product ontology that improved search precision and recall from customer behaviour and external datasets.

Senior Software Engineer September 2013 – September 2015
CareerBuilder, Norcross, GA
Led a machine learning school and company normalization system that identified and standardized educational institutions, companies and other entities from unstructured text, and developed profile merging and matching across multiple data sources for the Recruitment Edge project. Co-author of the 2016 Springer chapter "Data Science and Big Data Analytics at CareerBuilder".

Software Engineer June 2004 – July 2007
Stabilix Private Solutions Ltd., Trivandrum, India
First professional role, supporting a United States insurance client from India and building the security module of Health Companion, a generic application and claim-processing system for insurance companies.

EDUCATION AND RESEARCH

Ph.D. in computer science, The University of Alabama (2009–2013) — thesis "Modeling Computation-Intensive Problems for Heterogeneous Computing and Improved Source Code Maintenance", advisor Dr. Jeff Gray; the research was supported by an Air Force Research Laboratory STTR grant. M.S. in computer science, Clarkson University (2007–2009) — thesis "CSeR: A Code Editor for Tracking and Visualizing Detailed Clone Differences". B.Tech. in electronics and communication, National Institute of Technology Calicut, India (2000–2004). Ph.D. intern at Pacific Northwest National Laboratory (2012): two domain-specific languages that let scientists describe and deploy signature-discovery workflows, integrated into the laboratory's pipeline and published in IEEE Computing in Science & Engineering (2014).

TEACHING

Part-time Assistant Professor of computer science, Kennesaw State University, since May 2017 — information retrieval, big-data analytics, text mining, data warehousing and mining, requirements engineering, software testing, software project management and algorithm analysis.

PUBLICATION RECORD

Twenty-seven published items since 2009: five refereed journal articles, a Springer book chapter, sixteen refereed conference papers, four refereed workshop papers and a conference tutorial — in IEEE Computing in Science & Engineering, The Journal of Supercomputing, Computer Physics Communications, the Journal of Parallel and Distributed Computing, IEEE Big Data, the AAAI industry track, AMIA and CIKM. Full list at feroshjacob.github.io.