
SUMMARY

Co-founder and operator of AI systems in production. At Dicta I designed and run a multi-agent pipeline that drafts, reviews, and ships voice-matched executive content for paying B2B clients around the clock: 300 posts through the delivery ledger in its first two months, across 9 executive voices. Alongside it, an independent research program: persistent-memory infrastructure run across 3,300+ logged agent sessions, two preprints, and the first runnable implementation of the Narrative Continuity Test. A decade of shipping applied ML before that, including an NLP system delivered on ATTRACT funding via CERN, a startup grown to €1M revenue and acquired, and \$1.74M in AI consulting billings.

EXPERIENCE

Co-Founder

Aug 2025 – Present

Dicta Technologies Inc. · Canada (remote)

AI-native executive-content company, co-founded with one partner and run end to end on agent infrastructure.

- Designed and operate a multi-agent production pipeline: dispatcher, voice-bound drafters, a blind cold-read agent, a five-dimension review gate, and bounded auto-revision with named repair strategies. 300 posts through the ledger to date, voice-matched across 9 executive voices for 3 B2B client accounts.
- Built the editorial QC system: a three-stage gate chain plus an append-only gate-signal ledger; every flag is dispositioned before the next batch dispatches. The core production prompt is at 22 documented major versions, each one a caught failure, fixed.
- Automated the upstream: agent pipelines have triaged 3,500+ news candidates and extracted 950+ insights from client source material to feed the drafters.
- Built an AI SDR for inbound leads: prospect research, sample content drafted in the prospect's own voice, and personalized outreach.
- Run the company with a two-person human team; agents carry production, QC, scheduling, dashboards, and reporting.

Senior Consultant

2023 – 2025

Humanistic · Canada

- Ran AI/ML product engagements alongside the founder and CEO; 12 client projects totaling \$1.74M in billings, with 9 of 12 advancing to production pilots. Clients included a Fortune 500 automaker and a consumer electronics company.
- Directed development of ML interaction models that cut manual review time from 4 hours to 25 minutes per batch.
- Led technical evaluation for GenAI, NLP, and computer vision initiatives, translating recommendations into board-ready proposals.

Director of Operations

2021 – 2023

SPS Inc. (acquired by Phoenix GmbH) · Canada & Germany

- Dropped into two programs weeks from failure; shipped both on schedule by automating bottlenecks and adding real-time analytics.
- The turnaround caught Phoenix GmbH's attention; the acquisition closed 14 months later.

Co-Founder

2019 – 2021

Miltton Branch · Finland

- Co-founded an ML insight and foresight startup; recruited a 9-person team including 5 ML engineers; reached €1M revenue in year one.
- Signed 7 enterprise clients and 2 EU research partnerships; negotiated and closed acquisition by a multinational buyer (NDA) in summer 2021.

Program Manager

2015 – 2019

emmy.Network · Canada & Finland

- Managed an 8-person research team; delivered a production NLP system on a 6-month timeline.
- Won €410K in grant funding, including ATTRACT funding via CERN; spun out the ML project that became Milton Branch.

RESEARCH & OPEN-SOURCE PROJECTS

INDEPENDENT PROGRAM SINCE FEB 2026 · BUILT IN DISCIPLINED HUMAN-AI PARTNERSHIP; METHOD DISCLOSED

- **Solitaire:** persistent memory and personality infrastructure for LLM agents. 34,782 entries in production; retrieval precision@3 raised from 80% to 100% across four repair rounds; persona-trait alignment improved from 2.25 to 4.43 out of 5. Approved for listing on mcpservers.org; public release in progress.
- **Two preprints:** *From Memory to Partnership* (systems paper: four evaluation tracks over 400+ sessions, confounds stated in the abstract) and *Artificial Relational Intelligence* (position paper with falsification criteria).
- **Narrative Continuity Test, operationalized:** first runnable implementation of the NCT (Natangelo 2025). 25 behavioral cases across five identity axes, with a context-verification harness and a bare-model baseline arm so every score point is attributed: retrieval, behavior, or base-model credit.
- **Deep Applied Research:** a hypothesis-and-verdict protocol for research agents (1.9k LOC, approx. 250 passing tests). Measured its own 77% adopt-rate bias across 96 reports, shipped corrections, and re-measured.
- **33-category interaction-quality taxonomy** for AI text and behavior, with detection heuristics per category; runs daily as a live production gate at Dicta.

SKILLS

Agent systems: multi-agent orchestration · pipeline design · MCP · prompt systems & QC gates · bounded auto-revision

Evaluation: benchmark design & implementation · LLM-judge calibration · baseline attribution · failure taxonomies

Memory & retrieval: knowledge graphs · retrieval ranking · personality persistence · SQLite

Leadership: technical program management · P&L ownership · client delivery · M&A (two exits)

EDUCATION & CERTIFICATIONS

Master of Arts, International Studies · Simon Fraser University

Bachelor of Arts, International Studies · Simon Fraser University

Certified Scrum Master · Scrum Alliance

AFFILIATIONS

Former Associate Member, CERN · Association for Information Science and Technology (ASIS&T)