

Mathematician building production AI systems across video, language, and agentic workflows. Chasing AI since 2011; building language systems since an NLP internship at an Intel-acquired startup. The last two years: multimodal and agentic systems.

EXPERIENCE

Principal Research Data Scientist · iSpot.tv · Los Angeles

Aug 2024 — present

Data Science Manager, Aug 2022 — Aug 2024 · Data Scientist, Feb 2018 — Aug 2022

- Designed and built a multimodal video-intelligence system on **40k+ ads in production**, around one conviction: image captions without the broader context are worthless. A Gemini audio first pass — speech plus the sound effects, music, tone, and intonation that catch sarcasm, irony, laughs, and crashes — informs the storyline, and every keyframe caption is then written with that storyline as context.
- Curate and maintain the **regression set that governs LLM model and vendor migrations** — GPT-4, then Claude on Bedrock — output quality compared across highly varied ads before any model change ships.
- Proposed and prototyped **Sage, the company's agentic analytics platform**: atomic operators (aggregate, sort, filter, transform) mapped to deterministic tool sequences, so the agent is auditable instead of a black box. The prototype shaped the production design.
- Turned a decade of hand-written data-science libraries into **MCP tools** for pulling data from multiple sources and analyzing it auditably.
- Built contextual theme classification that uses video storylines as context rather than ad text alone; lifted accuracy versus text-only baselines on held-out creative.
- Built semantic search over ad-survey responses on PostgreSQL + pgvector — each ad scored by how many response sentences cleared the query's similarity cutoff, deduplicated. Also stood up a PyTorch inference path for BERT classifiers.
- Work the full loop with executives, sales, and product: scope the job, constrain the model, ship the integration, and explain what the system can and cannot do.

AI Technical Advisor · Perception Grid

Jun 2024 — present

- Advise and help build the AI architecture for a platform that correlates documents, images, video, and 3D/4D assets into searchable digital twins; video analytics tying storylines to script elements and moodboard character ID.

Independent AI Consultant · The House of Us

Jul 2025 — Feb 2026

- Deployed an LLM document pipeline (Claude) extracting structured fields from thousands of scanned and text PDFs under a JSON schema, with an analytics dashboard for the operator team. Published with synthetic data: github.com/ctmitts/order-intelligence.

SELECTED PROJECTS · DEMOS AT COLINMITTS.COM

Video Intelligence — give it a video, get back a report that explains it: the storyline in beats, every scene with what was said and what the picture shows. Reads the video first to decide what kind of report to write. ~\$0.21 and under a minute per video; live reports on the site.

iMessage Baby Tracker — built for my infant son: four caregivers text the family group chat in two languages; an LLM pipeline (own eval suite, 48 cases) turns chatter into typed events with goals, trends, a dashboard, and twice-daily age-aware briefings texted back.

SKILLS

Systems Production multimodal pipelines · LLM eval / regression harnesses · tool-using agents · MCP tool development · RAG and vector retrieval · structured extraction · FastAPI / Docker / AWS (Bedrock)

Languages & data Python · SQL · JavaScript/TypeScript · PostgreSQL / pgvector · PyTorch inference · ffmpeg · Whisper · Git

Field Scoping ambiguous workflows · eval design with domain partners · turning model behavior into a product contract · advising non-technical leadership

EDUCATION

Boston University — B.A., Applied Mathematics and Economics · directed study on virtual-assistant architecture, 2011

Santa Clara University — MBA, Business Analytics & Data Science

General Assembly — Data Science Immersive · recursive Wikipedia semantic database with TF-IDF/LSA vector search

Off hours: piano, skiing, cooking, proud new father.