CASE STUDY

How CodeContent produced 100 engineer-written tutorials generating 250,000+ views for Weights & Biases

100 Articles
Engineer-written tutorials and blog posts delivered over a 3+ year engagement.

250K+ Views
Total organic views across all published articles on W&B's platform.

6 Topic Clusters
LLMs, RAG, MLOps, computer vision, NLP, and ML fundamentals.

About W&B

Weights & Biases is the MLOps platform that ML teams at OpenAI, NVIDIA, and Microsoft use to track experiments, evaluate models, and ship production ML. Their Fully Connected content platform publishes tutorials for the ML practitioner community. Learn more at wandb.ai

The Challenge

W&B's content team could define topics and provide outlines, but translating those briefs into tutorials with working code, real W&B SDK integration, and validated model outputs required a writer who could actually train models.

Before the engagement, tutorials shipped sporadically and required heavy engineering review to verify that code examples produced the claimed results. The constraint was not ideas; it was finding writers with genuine ML implementation skills.

Why CodeContent

W&B evaluated CodeContent's existing ML content portfolio and engineering depth. The quality of implementation in published tutorials (working code, real experiment logs, validated model outputs) demonstrated the technical capability their content program required. What started as an initial engagement scaled into a multi-year partnership.

W&B needed technical content that was:

- ✓ Implementation-first with code that actually runs
- ✓ Integrated with W&B SDK (Tables, Sweeps, Weave)
- ✓ Publishable with minimal review from engineering
- ✓ Consistent: 4+ articles/month on a fixed cadence

> "Mostafa has worked with me as an author for a couple years now, and writes top-notch content related to the AI space, including fully-functioning tutorials to help educate our readers and give them immediate hands-on experience. He's a pleasure to work with, professional, hits his deadlines, and clearly knows the AI space across a variety of use cases and frameworks.
>
> **Dave Davies**
> Content Team Lead at Weights & Biases

The Work

The assigned engineer-writer had hands-on experience with PyTorch, Hugging Face Transformers, PEFT, LangChain, and the W&B SDK. Each tutorial followed a fixed process: set up the environment, implement the full pipeline, log experiments to W&B, validate outputs, then write.

PyTorch · Transformers · PEFT / LoRA · LangChain · ChromaDB · W&B SDK · W&B Weave · W&B Tables

Technical Depth

The GPT-5.4 computer use agents tutorial (March 2026) builds a fully observable browser and desktop automation agent using the OpenAI Agents SDK, tracing every screenshot, click, and decision through W&B Weave. It represents the current frontier of agentic AI content.

The production-ready fraud triage copilot builds, evaluates, guards, and monitors an LLM fraud analyst end-to-end with W&B Weave, covering the full lifecycle from prompt engineering through production guardrails.

The document QA system with LangChain and Weave builds a RAG pipeline for financial documents, then uses Weave to systematically detect and eliminate hallucinations at every retrieval stage.

Across the 100 articles produced, each averaged over 2,500 words and 30 code blocks, with consistent depth across the entire portfolio. The top-performing article alone generated over 40,000 views.

Navigating Change

The engagement spanned a period where W&B's product surface expanded rapidly into Sweeps, Artifacts, Tables, Inference, and Weave. Covering new features sometimes meant working with limited documentation, requiring direct SDK experimentation. That ramp-up time was real, but it meant the tutorials often became the first comprehensive implementation guides available for those features.

Before & After

DIMENSION	BEFORE	AFTER
Content velocity	Sporadic, no cadence	4+ articles/month, consistent
Technical depth	Docs-level overviews	2,600+ words, 30 code blocks avg
Topic coverage	Core tracking features	6 clusters: LLMs, RAG, MLOps, CV, NLP
Code validation	Varied; some untested	Every example implemented + logged
Product coverage	Experiment tracking only	Tables, Sweeps, Artifacts, Inference, Weave

Top Performing Content

The portfolio includes several high-traffic articles that continue to drive organic discovery:

ARTICLE	VIEWS
Learning Curves in ML	40,000+
Guide to LLMs	27,000+
The Basics of ResNet50	15,000+
Intro to Advanced RAG	7,000+
Fine-Tuning LLaMA 2	9,000+

Verify it yourself: Every article is live on wandb.ai. The code runs. The experiment logs are real. See all published work →

The Lesson

For AI/ML platforms, the gap between "we have docs" and "developers learn from our content" is bridged by implementation depth, not publishing volume. Tutorials where the writer configures 4-bit quantization with LoRA and builds custom evaluation pipelines compound in organic value because they answer what ML practitioners actually search for.

Engineer-Written
Every tutorial built by a writer who runs the code first.

SEO-Optimized
Structured for organic search and developer discovery.

Product-Integrated
Real SDK integration, not generic documentation snippets.

Scalable Pipeline
4+ articles/month sustained over a multi-year partnership.