

AI Laboratory Podcast

AI-Assisted Podcast Production & Live Experimentation

2025

7 Episodes Produced	3 Streaming Platforms	~80% Post-Production Time Saved	2 hrs Post-Work Per Episode
-------------------------------	---------------------------------	--	---------------------------------------

Section 1: What Was Built

The AI Laboratory Podcast was a live-streamed video podcast produced and broadcast simultaneously on YouTube, X (formerly Twitter), and Twitch. The show ran for seven episodes over approximately two months, with each episode running approximately one hour in length.

The podcast served a dual purpose: it was both a platform for solving real problems with AI — originally alongside guests, and later independently — and an AI experiment in its own right. The production workflow was deliberately built around automation tools to reduce the post-production labor that had killed previous podcast attempts.

Section 2: Why It Was Built

The podcast was built to solve two distinct problems simultaneously.

The first was a personal productivity problem. Prior podcast attempts had required two to twelve hours of post-production work per episode — video editing, publishing, clipping — a labor burden that reliably drained motivation and killed continuity. The goal was to determine whether AI-assisted automation could reduce that burden enough to make a regular cadence sustainable.

The second was a forcing function problem. Without a structured external commitment, AI experimentation tends to stay shallow. The podcast created accountability to conduct a new experiment every episode — exploring new tools, solving real problems, and producing something shareable regardless of the outcome.

Core Question

"Can an AI-automated production workflow reduce post-production labor enough to make consistent podcast production sustainable — and can a live format serve as a reliable forcing function for real AI experimentation?"

The purpose was simultaneously personal (skill development and AI exploration), professional (a demonstrable portfolio of AI experiments), and experimental (testing the limits of workflow automation with first-generation AI tools).

Section 3: Starting Competence

Prior Podcast Experience	Yes — but previous attempts required 2–12 hours of post-production per episode using traditional tools
Prior Streaming Experience	Limited — multi-platform live streaming was new
Prior AI Tool Use	Early stage — this was one of the first major AI projects undertaken
Video Editing Background	Practical but manual — familiar with the process, not with automation
Skills Brought In	Organizational instincts, clear design intent, and structured approach to workflow design

This project was executed at an early stage of AI fluency — making it a genuine first-generation experiment rather than a refined application of established methods.

Section 4: How AI Was Used

Workflow Design & Platform Selection

Before recording began, an extended structured conversation with Claude was used to design the show: target audience, format, content strategy, episode structure, and production workflow. This front-loaded design phase was a deliberate application of the author's planning instincts — define the architecture before executing.

Claude was also used to identify and evaluate automation tools that would reduce post-production friction. Two key platforms emerged from that evaluation:

- Yardstream — enabled simultaneous live streaming to YouTube, X, and Twitch; recorded full episodes; and generated short-form clips for YouTube Shorts

- Opus — provided an additional layer of automated clip generation for short-form social content

Together, these tools reduced post-production time from a potential 2–12 hours per episode down to approximately 1–2 hours — roughly 20% of prior effort.

Episode Design & Scripting

Claude was used to generate episode ideas, outline content, and build hooks and talking points for each session. This reduced the cognitive load of show preparation and ensured each episode had a clear structure before going live.

Live Experimentation

Each episode featured a live AI experiment — a problem attempted in real time with AI tools, sometimes succeeding, sometimes failing. Sample experiments included:

- Holiday gift idea generation — using AI to produce creative, personalized recommendations
- AI-generated artwork and concept imagery
- A fantasy sports advice chatbot — designed to ingest current player stats and produce lineup recommendations (a partial failure due to scope limitations in the first implementation)
- Live website builds
- AI agent design sessions for book publishing, advertising, and content workflows — several of which became standalone portfolio projects

Key Workflow Insight

The most valuable contribution of AI to this project was not just automating production — it was enabling a complete workflow redesign. By replacing manual video editing with automated clipping and multi-platform streaming, the show became operationally viable in a way prior podcast attempts never were.

Section 5: Timeline & Conditions

Episode Count	7 episodes
Total Duration	Approximately 2 months
Episode Length	~1 hour per episode

Streaming Platforms	YouTube, X (Twitter), Twitch — simultaneous via Yardstream
Work Conditions	Side project; built alongside other professional responsibilities
External Deadline	None — self-directed
Competing Projects	The Red Omen publication, luminariexpanse.com, advertising campaign, other portfolio work
Reason for Conclusion	Insufficient guest pipeline; professional work absences; shifted focus to higher-value AI projects

Section 6: Outcome & Evidence

The AI Laboratory Podcast successfully demonstrated that a fully automated, multi-platform podcast production workflow is achievable with first-generation AI tools — and that the resulting time savings are substantial enough to make consistent production operationally viable.

The format shift — from guest-driven problem-solving to solo experimentation — reduced the show's appeal over time. The guest format, which involved solving real problems live alongside other people, was significantly more engaging than independent content generation. The absence of a steady guest pipeline was the primary driver of the show's conclusion, not the workflow or the automation.

Several experiments conducted on the podcast became the foundation for standalone portfolio projects, including AI agent systems for book publishing and content production workflows.

Demo

YouTube channel with episode archive available for review.

Section 7: Key Metrics

Episodes Produced	7
Episode Length	~1 hour each
Streaming Platforms	YouTube, X, Twitch (simultaneous)
Short-Form Content	Clips generated via Yardstream and Opus for YouTube Shorts
Post-Production Time (Pre-AI)	2–12 hours per episode

Post-Production Time (With AI)	~1–2 hours per episode
Time Savings	~80% reduction in post-production effort
AI Tools Used	Claude, Grok, Gemini, ChatGPT + various task-specific platforms
Production Tools	Yardstream (streaming/recording/clipping), Opus (short-form generation)
Project Duration	~2 months
Audience Reach	Small — primary value was process and output, not audience scale

Section 8: Lessons Learned

Lesson 1: Design the Architecture Before You Build

The single most valuable step in this project was the upfront structured conversation with Claude to define what the show should be, who it was for, and how it should be built. That design clarity drove every subsequent tool selection and workflow decision. Vague inputs produce vague outputs — in podcast production as in all AI-assisted work. Starting with a fully articulated concept reduced downstream friction significantly.

Lesson 2: Automation Solves the Production Problem, Not the Motivation Problem

The workflow automation worked. Post-production time dropped dramatically, and the technical barrier to consistent production was effectively eliminated. What automation cannot solve is a missing collaborator. The most enjoyable and engaging format was guest-driven problem solving — and without a reliable pool of guests prebooked before launch, the show's core value proposition eroded. The lesson: for a format that depends on collaboration, guest acquisition is a prerequisite, not an afterthought.

Lesson 3: Failure Is Valid Data

Not every experiment worked. The fantasy sports chatbot underperformed due to scope limitations in the first implementation. Live experiments in front of an audience created real stakes and real feedback — including the kind of failure that refines future approaches. A format designed to embrace failure as a learning mechanism produced more genuine insight than one optimized for polished outcomes.

Lesson 4: Side Projects Compound

Several experiments conducted on-air became standalone portfolio projects. The AI agent systems explored on the podcast directly fed into book publishing workflows documented elsewhere in this portfolio. A podcast designed as a forcing function delivered on that promise: it generated experimentation that produced durable, reusable work beyond the episodes themselves.

What This Project Proves

Workflow automation capability: A multi-platform live podcast was designed, built, and operated with approximately 80% less post-production effort than traditional methods. First-principles AI experimentation: New tools were evaluated, selected, and deployed in a live production environment with no prior experience. Structured design thinking: Upfront architectural design — not just execution — drove the quality and efficiency of the final workflow. Curiosity and exploration: Seven episodes of live AI experimentation across diverse problem domains, with results ranging from strong successes to instructive failures. Adaptability: When the guest-driven format proved unsustainable, the show adapted rather than stopped — demonstrating the same scope management instincts applied across all portfolio work.

--- END OF CASE STUDY ---
AI Laboratory Podcast | YouTube, X, Twitch