

AI EDGE · MILES DEUTSCHER

I FIRED 5 EMPLOYEES WITH AI

The Complete Workflow Playbook

5 AI systems that replace full-time hires — build them yourself.

- 01 OpusClip Clipping & Distribution
- 02 AI Financial Advisor (Personal OS)
- 03 Personal AI Operating System
- 04 YouTube Ideation Engine
- 05 BD Outreach Automation

What's inside

These are the 5 AI workflows Miles uses to run AI Edge operations — each one replacing a function that used to require a dedicated hire. Every workflow is broken down into what it does, how to replicate it yourself, and the tools you need to get started.

BEFORE YOU START

- You don't need to be a developer. Every workflow here is buildable with AI coding tools like Claude Code or Cursor.
- Each workflow is modular — build one at a time. You don't need all five.
- Start with the workflow closest to a pain point you have right now.
- Treat these as starting points. Customise every step to your stack and situation.

01 OpusClip Clipping & Distribution

Turn one long-form video into a week of short clips across all channels automatically.

02 AI Financial Advisor

A personal fiduciary advisor that reads your portfolio, knows your rules, and pings your phone.

03 Personal AI Operating System

A live dashboard replacing tasks, CRM, journal, habits, nutrition, finance — one AI brain.

04 YouTube Ideation Engine

Structure competitor research into four buckets so AI can generate better content ideas.

05 BD Outreach Automation

Turn prospect lists into personalised, scored outreach drafts ready for human approval.

WORKFLOW ONE

01 OpusClip Clipping & Distribution

Turn one long-form video into a week of short-form content across every channel — automatically.

What This Workflow Does

One long-form video (YouTube, podcast, livestream) gets fed into OpusClip. The AI identifies the best moments, scores them by virality potential, auto-reframes to 9:16 vertical, adds captions, and applies your brand kit — all in one pass. You review, pick the top clips, and push to TikTok, YouTube Shorts, and Instagram in a single queue.

The Problem It Solves

Short-form distribution used to require a dedicated video editor to watch every video, manually clip highlights, reformat for vertical, add captions, and schedule across platforms. That's 4-8 hours per long-form video. This workflow cuts it to under 30 minutes — with higher output consistency than most editors.

How to Build It

STEP 1 · IMPORT YOUR VIDEO

Upload the long-form file directly or paste a YouTube/TikTok/Vimeo link. Credits consume roughly 1 per minute of footage.

STEP 2 · SET THE BRIEF

Choose clip length (30s, 60s, or custom). Set the format to 9:16 vertical. Use ClipAnything for non-talking-head content. Add a keyword or topic focus if you're targeting a specific theme.

STEP 3 · GENERATE & SCORE

AI returns a batch of clips, each with a Virality Score from 1-100. Higher scores indicate stronger hooks and payoff moments. Sort descending and start reviewing from the top.

STEP 4 · APPLY BRAND KIT

One click applies your logo, fonts, colours, caption style, and intro/outro. Set this up once and it applies to every clip automatically from that point forward.

STEP 5 · HUMAN REVIEW (NEVER SKIP THIS)

Always human-check before publishing: fix caption errors, trim any dead air, confirm the hook lands in the first 1-2 seconds. Never ship on defaults alone.

STEP 6 · SCHEDULE & DISTRIBUTE

Push clips to the built-in OpusClip scheduler or export to an external tool. Queue across all 6 channels at your set post times.

STEP 7 · TRACK & ITERATE

Log which clips perform best. Feed those topics back into your next clipping session. Over time this creates a performance feedback loop that improves output quality.

QUALITY BAR — EVERY CLIP BEFORE IT SHIPS

- Hook lands in the first 1-2 seconds — if not, trim or skip
- Captions on, accurate, and on-brand style
- Vertical 9:16, subject stays centred throughout
- One clear idea per clip — no dead air at the start
- Brand Kit applied and checked

LOCK THESE SETTINGS ONCE

- Brand Kit: logo, fonts, colours, caption template, intro/outro
- Default clip length (30s or 60s depending on platform)
- Number of clips to keep per long-form video
- Platforms + channels + scheduler + post times
- Who approves before publish (you, an editor, or auto?)

TOOLS USED

OPUSCLIP	CLIPANYTHING	BRAND KIT	TIKTOK	YOUTUBE SHORTS	INSTAGRAM
----------	--------------	-----------	--------	-------------------	-----------

WORKFLOW TWO

02 AI Financial Advisor

A personally-tailored fiduciary advisor that reads your portfolio, knows your rules, and pings your phone with a morning brief.

What This Workflow Does

A locally-running AI advisor that holds three things simultaneously: your written investment rules (a philosophy doc), your live portfolio, and your full conversation history. It generates structured morning briefs, scans overnight market moves for impact on your positions, and sends a Telegram alert to your phone — all grounded in your personal investing philosophy, not generic advice.

The Problem It Solves

A good human financial advisor costs \$10-15K/year and still doesn't know your full context. This system costs under \$20/month in API fees, remembers every conversation, applies your specific tax situation and risk rules to every recommendation, and is available 24/7. The edge isn't the AI model — it's that the model always has your complete context injected.

The Foundation: Your Investor One-Pager

This is the single most important step. Before building anything, write a `philosophy.md` file — roughly one page covering your North Star, core investment rules, time horizons, mindset pre-commitments, personal nuances (tax, income, time budget), and an anti-portfolio (what you will never buy). The AI reads this on every single call. A sharp one-pager turns a generic chatbot into an actual advisor.

WHAT GOES IN YOUR INVESTOR ONE-PAGER

- North Star: what you're optimising for, over what horizon, with what floor
- Core Philosophy: 3-5 rules about HOW you invest, not WHAT you buy
- Time Horizon: how you split capital across tactical / core / generational books
- Mindset Rules: numbered pre-commitments to your future self
- Personal Nuances: tax jurisdiction, income, time budget, public-narrative blind spot, family constraints
- Anti-Portfolio: what you will never buy regardless of upside

How to Build It

STEP 1 · WRITE YOUR PHILOSOPHY DOC

Create `philosophy.md`. Use the structure above. Be specific — 'long-term thinker' is useless. '30-year horizon, won't panic-sell below -40%' is useful. Every rule should be testable. Budget 30 minutes.

STEP 2 · CREATE YOUR PROJECT FOLDER

`mkdir ai-advisor && cd ai-advisor`. Create three files: `philosophy.md` (your one-pager), `portfolio.json` (your holdings with tickers, shares, avg cost), and `.env` (your API keys).

STEP 3 · INSTALL CLAUDE CODE OR CODEX

`npm install -g @anthropic-ai/claude-code`, then run `claude` in your project folder. Or use Codex via `npm install -g @openai/codex`. Both produce the same result.

STEP 4 · TIER 1 — CHAT ADVISOR (5 MINS)

Paste this prompt: 'I'm starting our advisory relationship. Read my philosophy.md and portfolio.json. Confirm you have my context, then ask me the one question that matters most given my current situation.' Done — you have a working advisor.

STEP 5 · TIER 2 — FULL DASHBOARD (30 MINS)

Paste the full build prompt into Claude Code. It builds: a Streamlit dashboard, portfolio charts, live price feeds via yfinance and CoinGecko, a chat interface, SQLite memory, and Telegram integration.

STEP 6 · SET UP TELEGRAM ALERTS

Create a bot via @BotFather, get your bot token and chat ID, wire both into your .env, then click 'Send Morning Update' in the dashboard. Your phone buzzes within 10 seconds.

HOW MEMORY WORKS

Three layers are injected into every AI call — this is the entire trick:

- Layer 1 — Philosophy (philosophy.md): your written rules, read fresh on every call
- Layer 2 — Portfolio (portfolio.json + live prices): your current state, updated daily
- Layer 3 — Conversation history (SQLite): last 30 messages auto-appended every turn

TOOLS USED

CLAUDE CODE	STREAMLIT	SQLITE	YFINANCE	COINGECKO	TELEGRAM	PYTHON 3.11+
-------------	-----------	--------	----------	-----------	----------	--------------

WORKFLOW THREE

03 Personal AI Operating System

A live dashboard replacing tasks, CRM, journal, habits, nutrition, finance — one AI brain that knows everything about how you work.

What This Workflow Does

A custom-built personal dashboard at [your-domain.vercel.app](#) — seven AI-powered cards covering your entire work and life operation. A Telegram bot captures voice notes from anywhere and routes them into the right database table automatically. Every piece of data is searchable via natural language. Total cost is around \$30/month at active use.

The Problem It Solves

Most people run 5-7 different apps to manage tasks, CRM, journaling, habits, nutrition, finances, and calendar. None of them talk to each other. This system unifies everything into one AI-native interface where data captured anywhere (voice, text, mobile) flows automatically to the right place, and you can ask questions across all of it at once.

The Seven Cards

Operator + Session	Your name, location, current focus, and today's top 3 AI-ranked tasks. The anchor card — renders first, loads fastest.
Calendar	Pulls from Google Calendar via iCal. 14-day week strip, click any day to view events, auto-scrolls to a NOW marker.
Habit Tracker	Six daily checks, click-to-complete, cross-device sync via Supabase, local midnight reset per your timezone.
CRM (Tasks)	Four urgency tiers (Today / This Week / Month / Someday), Kanban + Smart + Category views, drag-drop reorder, AI-powered natural language search.
Nutrition	Type a meal, AI returns macros. Edit any macro and calories recompute. Edit calories and AI redistributes macros. Daily auto-reset.
Goals	Weekly and monthly goal lists that never auto-clear — stored on a sentinel date so you manage them manually.
Finance Pulse	Reads your messy Google Sheet, AI figures out net worth and categories. Page loads never trigger AI — only a manual refresh or daily cron.

How to Build It

STEP 1 · PICK YOUR STACK

Frontend: Next.js 15 (App Router). Database: Supabase (Postgres + pgvector). AI: Claude as primary, OpenAI as fallback + Whisper for transcription. Hosting: Vercel (free tier). Capture: Telegram bot.

STEP 2 · DESIGN BEFORE YOU BUILD

Spend 30 minutes in Claude Design or v0.dev sketching the layout before writing any code. Describe it in plain English — 'dark mode, glassmorphism cards, 3-column grid' — and get a working HTML mockup. Give this to Claude Code as a target.

STEP 3 · BUILD THE CAPTURE PIPELINE FIRST

Everything downstream reads from what this pipeline writes. Voice note hits Telegram -> Whisper transcribes -> Claude classifies (kind, urgency, entity, tags) -> row written to right table -> embedded to memory_chunks. Target end-to-end: under 5 seconds.

STEP 4 · BUILD CARDS IN ORDER

Easiest first: Operator/Session -> Calendar -> Habit Tracker -> CRM -> Nutrition -> Goals -> Finance Pulse. Each card is a self-contained component with its own API route.

STEP 5 · ADD THE MEMORY LAYER

Every text artifact gets embedded via OpenAI text-embedding-3-small and stored in memory_chunks with pgvector. The Brain tab lets you ask natural language questions across your entire history.

STEP 6 · DEPLOY TO VERCEL

vercel --prod. Push all env vars (see the playbook appendix). Set up Vercel cron for the Finance daily snapshot. Point Telegram webhook at your live URL.

CRITICAL BUILD RULE

Never let page loads trigger AI calls. Pages read from the latest Supabase snapshot. Only manual refresh buttons or scheduled cron jobs call the AI. This single rule prevents burning your API budget on every navigation event.

TOOLS USED

NEXT.JS 15	SUPABASE	CLAUDE API	OPENAI WHISPER	VERCEL	TELEGRAM	PGVECTOR
------------	----------	------------	-------------------	--------	----------	----------

WORKFLOW FOUR

04 YouTube Ideation Engine

Structure competitor research into four buckets so AI generates better content ideas — not random ones.

What This Workflow Does

The Ideation Engine turns YouTube research into structured content ideas. Rather than asking AI to invent random video ideas from nothing, it first maps the competitive landscape into four research buckets — then generates title and thumbnail concepts grounded in what's already working. The result is ideas that are both original and market-validated.

The Problem It Solves

Most creators either copy competitors directly (derivative content) or brainstorm in a vacuum (ideas with no market signal). The Ideation Engine solves both problems. Competitor research provides the signal; the four-bucket system provides the structure that stops AI from generating generic output. The human strategist makes the final creative call — the AI handles the research load.

The Four Research Buckets

Direct Competitors

Same niche, same audience, similar format.

- 1 Used mainly to understand what NOT to copy. Tells you what's already saturated.

Indirect Competitors

Same niche or audience, but different style or angle.

- 2 Shows other approaches that are working in your space. Good for differentiation ideas.

Adjacent Competitors

Nearby audience or problem, different niche.

- 3 The most valuable bucket for fresh ideas. Concepts that transfer across niches tend to feel original to your audience.

Distant Competitors

Outside the niche entirely.

- 4 Use for formats, stakes, storytelling mechanics, and packaging ideas. How does a finance channel structure a hook that you could apply to AI content?

How to Build It

STEP 1 · PICK YOUR RESEARCH PLATFORM

YouTube is the default, but the same architecture works for TikTok, X, Reddit, or newsletters. The bucket logic is platform-agnostic.

STEP 2 · DEFINE AND FILL YOUR BUCKETS

Start with your own channel. Pull recent videos and channel metadata. Use AI to create a basic profile of your channel and audience. Then search for channels that fit each of the four buckets. Aim for 5-10 channels per bucket.

STEP 3 · CLASSIFY WITH AI

Feed candidate channels to Claude with your channel profile and the four bucket definitions. Ask it to classify each one and explain why. Review the classifications — remove anything that doesn't fit, add channels you know should be there.

STEP 4 · ANALYSE PATTERNS IN EACH BUCKET

For each bucket: what formats are performing (list, story, tutorial, reaction)? What hook styles are working? What thumbnail patterns repeat? What topics are getting outsized views relative to subscriber count?

STEP 5 · GENERATE IDEAS ONLY AFTER RESEARCH

Once buckets are populated and analysed, prompt AI to generate title and thumbnail ideas. Give it the patterns you found plus your channel context. The research is what separates useful AI output from generic suggestions.

STEP 6 · EDIT, FILTER, AND DECIDE

AI output is a starting point. The strategist edits, combines, discards, and improves. Keep outputs in an editable format (spreadsheet or simple dashboard). The final creative judgment stays human.

STEP 7 · SAVE AND REVISIT

Store your research so you can return to it without starting over. A good database of classified competitor channels becomes more valuable over time as you add more data points.

SIMPLE VERSION VS FULL APP

- Simple version: Google Sheet for research storage + Claude for classification and ideation. Can be running in an afternoon.
- Intermediate: Zapier or Make automation to pull channel data + Claude API for classification + Airtable for storage.
- Full app: Visual frontend + backend YouTube Data API integration + AI layer + saved search history. 1-2 week build with Claude Code.

TOOLS USED

YOUTUBE DATA API

CLAUDE API

AIRTABLE / SHEETS

ZAPIER / MAKE

CLAUDE CODE

WORKFLOW FIVE

05 BD Outreach Automation

Turn prospect lists into personalised, scored outreach drafts ready for human approval — without sending a single generic message.

What This Workflow Does

The BD Automation workflow converts raw prospect lists into enriched, AI-researched, scored leads with personalised outreach drafts — ready for a human to approve and send. The goal is not to automate sending. It's to automate the research, context-gathering, and first-draft work that makes outreach worth sending in the first place.

The Problem It Solves

The hard part of business development isn't writing the email. It's knowing who is worth contacting, why now, and what angle makes the message actually relevant to that specific person. That research process — which used to take 20-30 minutes per prospect — can be compressed to under 2 minutes per lead with AI. A human still makes the final call. They just spend their time on decisions, not research.

The Five-Part Architecture

Lead Source	Where prospects enter the workflow — manual import, LinkedIn export, event attendee list, inbound form, or scraped list. Clean the data into a consistent format before anything else.
Enrichment	Add context about each prospect: company size, recent news, funding status, social activity, mutual connections, relevant content they've published. Tools like Apollo, Clay, or Hunter.io automate much of this.
AI Research	Claude summarises each account in 2-3 sentences and identifies a personalisation angle: why this prospect is relevant right now, what their likely pain point is, and what the credible opener is.
Scoring	Rank prospects by fit, urgency, and likelihood of conversion. Use a simple rubric: ICP match, recency signal (recent funding, hiring, product launch), and relationship warmth. High score = prioritise immediately.
Action	Generate the outreach draft, internal BD brief, or follow-up task. Push approved next steps to the CRM. Log all replies, booked calls, and disqualifications to improve the scoring model over time.

How to Build It

STEP 1 · DEFINE YOUR ICP

Before building anything: write down exactly who you're targeting. Company size, industry, role, signals that indicate buying intent. This definition feeds every AI prompt in the workflow.

STEP 2 · SET UP YOUR LEAD SOURCE

Import an existing list or connect a source (LinkedIn Sales Navigator CSV, Apollo export, event list). Standardise to: name, company, role, email, LinkedIn URL, and any context notes.

STEP 3 · ENRICH EACH RECORD

Run each prospect through an enrichment layer. Clay is the most flexible (pulls from 50+ sources). Apollo covers contact data. Manually add context for high-priority accounts. You want: recent news, company funding, relevant publications, and hiring signals.

STEP 4 · AI RESEARCH & PERSONALISATION

For each enriched record, run a Claude prompt that ingests all available context and outputs: a 2-sentence account summary, a personalisation angle (why them, why now), and a suggested opener line. Store these as fields in your CRM or spreadsheet.

STEP 5 · SCORE AND PRIORITISE

Apply your scoring rubric across the full list. Sort by score. The top 20% get immediate personal outreach from you. The middle 60% go into a nurture sequence. The bottom 20% get deprioritised or disqualified.

STEP 6 · GENERATE OUTREACH DRAFTS

For top-priority prospects, use AI to generate a first draft using the personalisation angle from Step 4. Keep drafts short — 3-4 sentences maximum. The human reviews, edits the voice, and sends.

STEP 7 · LOG, ITERATE, IMPROVE

Track every reply, booking, and disqualification. Feed this data back into the scoring model. Over 3-4 months you'll have a data-backed sense of which signals actually predict conversion for your specific use case.

WHAT AI IS ACTUALLY GOOD AT HERE

- Summarising large amounts of public context quickly (company pages, press releases, LinkedIn)
- Identifying a personalisation angle you might have missed
- Generating multiple first-draft openers to choose from
- Scoring consistency — applying the same rubric at scale without fatigue

WHAT STAYS HUMAN

- The final decision to reach out (AI draft + human send, not AI send)
- Editing the voice — AI drafts sound like AI; your edits make them sound like you
- Relationship context that isn't in any database
- The strategic judgment on timing and positioning

TOOLS USED

CLAY	APOLLO.IO	CLAUDE API	MAKE / ZAPIER	HUBSPOT / NOTION CRM	HUNTER.IO
------	-----------	------------	---------------	-------------------------	-----------

FINAL NOTE

Start With One

You don't need all five of these running simultaneously. The highest-leverage move is picking the one that maps to your biggest time sink right now and building just that one.

Once it's running and saving you real time, build the next one. The compounding effect of five AI systems working in parallel is significant — but it starts with shipping one.

WHICH WORKFLOW FIRST?

- Creating lots of video content? Start with Workflow 01 — OpusClip Distribution.
- Actively investing and checking markets regularly? Start with Workflow 02 — AI Financial Advisor.
- Running too many productivity apps that don't talk to each other? Start with Workflow 03 — Personal OS.
- Publishing content and struggling with ideation? Start with Workflow 04 — Ideation Engine.
- Running BD or sales and spending too much time on research? Start with Workflow 05 — BD Automation.

For more AI workflows, tools, and deep dives — follow @aiedge_ on X.

Miles Deutscher · AI Edge · Giving You The Edge On AI