

Mentor Seminar • AI Edition

Seminar Summary — What We Covered and Why It Matters

MHG Charitable Foundation • Make AI Work for You • 2026

This summary distills the full **2026 Mentor Seminar (AI Edition)** — twelve sections spanning how AI works, the tools that matter, the economics behind the boom, and what it all means for jobs and students. Read it as a map: each section pairs the core idea with the practical takeaway.

01 The AI Reality Check: Bloopers & the Gold Rush

Two cautions frame the day — AI's failures, and the hype around it.

- **Bloopers:** glue-on-pizza advice, a "nine McNuggets" miscount, a chatbot agreeing to a \$1 car deal — AI is a **pattern machine, not a truth machine**. Demos are easy; the real world is messy; without guardrails, small errors become big problems.
- **Gold rush:** a sneaker brand rebranding "Allbirds → Smartbird"; a company adding ".AI" and moving its stock ~146% in a day; a crypto miner pivoting to CoreWeave (a genuine pivot — the GPUs really transferred).

Takeaway Treat AI output critically, and separate real value from a name-change narrative.

02 Where AI Stands: Adoption & the Prompting Warm-Up

Usage is widespread but shallow; better prompts are the fastest upgrade.

- Adoption is **broad but uneven** — ChatGPT leads while rivals rise, students already learn with AI, and productivity gains are real but workflow change lags. Everyone uses AI; few use it systematically.
- **Starter pack — six everyday uses:** Ask, Summarize, Translate, Create, Digest, Produce.
- **The prompt recipe (R-G-C-F-E):** Role, Goal, Context, Format, Examples. Don't ask for more — ask more **clearly**, then iterate.

Takeaway The clearer the request, the better the result — a structured prompt beats a vague one every time.

03 How AI Works: The Fundamentals

Under the hood: from how a model is built to how values are trained in.

- AI **doesn't know — it predicts**. A model is a probability engine; the model is the **engine**, while apps, assistants, and agents sit on top.
- **Constitutional AI:** Claude is given a written set of values and learns to check and fix its own answers against them — the why, not just rules. Amanda Askell leads Claude's character work (the "soul doc").
- **The debate:** who decides the values, when one company writes the rules and grades itself?

Takeaway Knowing AI predicts (rather than knows) explains both its power and its mistakes.

04 The AI Landscape: Frontier Models, Assistants & Subscriptions

Who builds the models, how the assistants differ, and what to pay for.

- Four groups now define the LLM landscape; the **model lab is becoming the new platform company**, and foundation models are becoming the operating system for knowledge work. Winning takes five "moats".
- **Assistants, side by side:** three assistants with different strengths; **Perplexity** is the answer engine (cited answers for early research).
- **Subscriptions:** match what you spend to what you need — budget vs. service.

Takeaway Pick tools by task and by moat, not by brand loyalty.

05 The AI Economy: Capex, Circular Financing & the Supply Chain

AI is no longer just software — it is a capital-intensive infrastructure race.

- **Spending arms race:** four giants heading for ~\$640B of capex in 2026 (~4× their 2023 level). Bigger picture: ~\$1.1T hyperscaler capex by 2027; ~\$5.5T for AI, data centers & power; ~\$7T of data centers by 2030.
- Every AI dollar buys **chips, data centers, and power** — a full infrastructure buildout, not software alone.
- **Circular financing:** backers and clouds invest in the AI labs, who spend it right back on compute — strong while AI revenue grows, but a circular risk (boom or bubble?).
- **Supply chain:** AI is built, not just coded — from sand to intelligence. Bottlenecks are physical capacity, not ideas; NVIDIA captures the value today, but value can move.

Takeaway Capital scale is now the moat — and the open question is productive buildout vs. circular bubble.

06 AI in Action: Games & Speech-to-Text

Two hands-on use cases — building software, and digesting media.

- **Games:** from hand-built pixels to AI-built worlds — **describe it** instead of coding it. Same jobs done faster (code, art, dialogue, voice, NPCs); AI doesn't replace the maker, it speeds each one up.
- **Speech-to-text:** turn any podcast or video into readable, searchable, summarizable text — no tech skills. Four steps: **Capture → Transcribe (TurboScribe) → Summarize (Claude/ChatGPT) → Keep.**
- A reusable **summary prompt:** one-sentence point, top 5 ideas, why it matters, the action it leads to, key facts, caveats, and a 60-second version — adding nothing the source doesn't say.
- **Voice ecosystem:** cloning, speech-to-speech (keeps tone & emotion), dubbing, live agents. Players: specialists (ElevenLabs) vs. platforms (OpenAI, Google, AWS).

Takeaway You can now build and digest content by describing what you want — the barrier is clarity, not code.

07 Learning Faster with AI: the Accelerator, RAG & NotebookLM

Use AI to learn faster — not to think less. AI as tutor, examiner, and coach.

- **The 8-step Learning Accelerator:** Goal → Source → Prime → Learn → Retrieve → Correct → Apply → Teach.
- **Tool stack:** Find (Perplexity), Capture (TurboScribe), Understand (NotebookLM), Remember (Quizlet/Anki), Apply (Gamma/Canva).
- **NotebookLM** — your AI study room: grounded with citations, upload PDF/web/YouTube, auto study guides, audio & video overviews, mind maps & flashcards, no hallucination. One source pack generated six study assets (guide, mind map, checklist, audio/video, flashcards, quiz).
- **RAG (Retrieval-Augmented Generation):** the AI reads your sources first, slipping the right passages into the prompt for grounded, cited answers. Same core (memory, retrieval, tools) across NotebookLM, Claude Projects, and ChatGPT Projects — different emphasis.

Takeaway Ground the AI in your own materials, and it becomes a tutor that can't make things up.

08 Meet Claude: Powers, MCP & Cowork

From chatbot to workflow assistant — and how the pieces connect.

- **Anthropic:** the safety-focused lab behind Claude — the most valuable AI startup (~\$965B), ~\$47B revenue run-rate, S-1 filed June 2026; founded 2021 by ex-OpenAI leaders Dario & Daniela Amodei.
- **Claude's powers:** Chat, Design, Skill, Connector, Plug-in, Cowork, Code. The formula: **Connector + Skill = Plug-in** (packaged, shareable).
- **MCP (Model Context Protocol):** an open standard (Anthropic, 2024) — the "USB-C" for AI tools. Any assistant can discover tools, read data, and take actions across services (Canva, Drive, GitHub, Slack) without custom integrations; OAuth authorizes each connection.
- **Cowork:** a "digital coworker" on your PC — you set the goal; Claude thinks in the cloud, works on your PC, runs code safely in a Linux sandbox, and calls connectors, skills & plugins.

Takeaway Claude has shifted from answering to doing — MCP is the plug that lets it reach your tools.

09 Finance Demo: From Business Data to Insight

Turning raw operational data into decisions — in minutes.

- From **119,000+ hotel booking rows** (30+ columns) to clear decisions, driven by one plain-English prompt ("analyze as a revenue-management analyst...").
- **Four business answers in minutes:** demand & revenue trend, cancellation risk, guest-segment profitability, and a plain-English management memo.
- **Your turn:** do it with your own bank/card statements or a company GL — understand your data, build a few charts, ask good questions, collaborate with AI.

Takeaway AI compresses hours of spreadsheet analysis into minutes — with you driving the questions.

10 AI-Assisted Investing: Finance Agents & Reading an S-1

An analyst's workflow, packaged — and applied to a real IPO filing.

- **Four free Claude finance plugins:** Financial Analysis (DCF/comps), Equity Research (ideas/screens/coverage), Earnings Reviewer (results → rating & target), Market Researcher (sector maps). Workflow: scan sector → shortlist → value → check & decide.
- **Case study — SpaceX's 200-page S-1:** largest U.S. IPO on record; \$18.7B 2025 revenue; \$1.77T IPO valuation (Jun 2026); 10.3M Starlink customers (Starlink = 61% of revenue). Six moves: summarize → map market → build model → value → frame thesis → draft memo.
- **Two ways to drive it:** your own prompt (Path A) vs. a prebuilt skill **/initiate** (Path B). Cowork ships the analyst playbook built in; ChatGPT gives a flexible workspace you assemble yourself.
- **Grounded research:** Consensus searches 200M+ peer-reviewed papers (with a Consensus Meter and study snapshots); Deep Research × Consensus turns evidence into a cited report — from information to understanding.

Takeaway A single command can now run a full analyst playbook — grounded in real evidence.

11 Multimodal, Multi-Model & AI Agents

The architecture shift from single models to coordinated, acting systems.

- **Multimodal** (many inputs → one model) vs. **multi-model** (one task → many models). Patterns: native multimodal, retrieval/tool hub, orchestrator-workers. Multi-model is now the norm — Datadog 2026: **70% run 3+ models** (up from 56%); 6+ models rose 23% → 41%.
- **AI agents:** from chatbots to doers — a system that uses a model to pursue a goal over many steps (reason, act, remember, coordinate). Unlike rule-based programs, an agent chooses its own steps toward your goal.
- **Status (2026):** real but uneven — fast in coding/research, still hard on long office work (~24% best-agent office success; >40% of agent projects dropped by 2027, per Gartner).
- **An ecosystem, not just a model:** orchestration, protocols (MCP = what an agent can reach; A2A = which agents it works with), and oversight. The hard part is now **trust and control**.

Takeaway The frontier is moving from models that answer to agents that act — and standards matter as much as models.

12 AI & the Future of Work

Tasks change first, jobs second — AI reprices human value.

- Slice any job into tasks and score each against AI. **Exposure** = the share AI can already do; the rest stays human. Value shifts toward human skills + AI fluency.
- **Broad exposure, limited job loss so far:** ~40% of global jobs exposed (IMF 2024); 24% of workers have some exposure (ILO 2025); ~17-20% of U.S. firms now use AI (from 3.8%); 37.4% use AI on the job (St. Louis Fed); +78M net new jobs by 2030 (WEF); -13% jobs for young exposed workers (Dallas Fed). Adoption races ahead of the jobs signal — restructuring, not mass layoffs.
- **Capability ≠ impact — redesign decides.** AI is closest to routine data, admin & rules work; furthest from judgment, empathy & complex decisions. Exposed & declining: data-entry, tellers, cashiers, bookkeeping, customer service, claims. Durable & growing: nurse practitioners, data scientists, security analysts, home-health aides, trades, social workers (exposure ≠ decline — developers +15%).

- **What to do:** move up the value chain — own exceptions, go "T-shaped", build real AI fluency, own judgment & context, signal skills early. **For students:** build what AI can't copy — critical thinking, communication, leadership, creativity, domain expertise. The winner = expert + AI.

Takeaway Ask not "will AI take my job?" but "which tasks will it change?" — then move up the value chain.