

AI Playbook for Entrepreneurs & Ministries

Dr. Ron Pedemonte

BlessBack Business Academy v.5 January 2026

Introduction: Why this Playbook exists.....	2
What Does God's Word Say About Artificial Intelligence?.....	8
Chapter 1: Which AI Platform should I use?.....	11
Chapter 2: Settings Quickstart (ChatGPT + Gemini).....	26
Chapter 3: Deep Research with AI.....	28
Chapter 4: Custom GPTs & GEMs.....	44
Chapter 5: Projects in ChatGPT.....	54
Chapter 6: Study & Learn — Guided Learning.....	60
Chapter 7: Agents.....	64
Chapter 8: Text to Image AI Tools.....	69
Chapter 9: Text to Video AI Tools.....	81
Chapter 10: Turn Text Into Visuals (Napkin.ai).....	91
Chapter 11: Google AI Studio.....	94
Chapter 12: Google NotebookLM.....	105
Chapter 13: App Builders.....	114
Chapter 14: Gemini/Google Tools Integration.....	127
Chapter 15: Coze (Agent & App Builder).....	135
Chapter 16: Canva AI.....	148
Chapter 17: AI Slide Generation Platforms (Best Options + Workflows).....	156

Introduction: Why this Playbook exists

AI can amplify our impact—speeding up research, standardizing teaching quality, simplifying bilingual production, and turning repeatable tasks into reliable workflows—while staying rooted in people-first, Scripture-shaped values. This manual translates what works in the field into **repeatable, low-friction workflows** that anyone on your team can run —whether they are instructors, ministry leaders, volunteers, or entrepreneurs.

Why AI, why now

Adoption has moved from hype to habit. In 2025, the U.S. Chamber reports **58% of U.S. small businesses** use generative AI (up from 40% in 2024), and most say restricting AI would hurt growth and operations. (uschamber.com) Microsoft finds **75% of knowledge workers** already use gen-AI at work. ([Microsoft](https://www.microsoft.com)) McKinsey's 2025 global survey shows 71% of organizations now use gen-AI in at least one function (especially marketing/sales, service operations, product development, and software). ([McKinsey & Company](https://www.mckinsey.com)) The OECD estimates that responsible AI at scale can increase labor productivity growth over time. ([OECD](https://www.oecd-ilibrary.org)) QuickBooks' 2025 SMB pulse: **68%+** use AI regularly, and **74%** say it boosts productivity. ([QuickBooks](https://www.quickbooks.com))

Where AI shows up in small businesses

- **Office “copilots”:** Drafts, summaries, meeting notes, slide edits, email replies—embedded in Microsoft 365 and Google tools; usage is now widespread across knowledge workers. ([Microsoft](https://www.microsoft.com))
- **Marketing & content:** SEO briefs, ad variants, social calendars, image/video generation; teams report expanding automation and early agent use. ([McKinsey & Company](https://www.mckinsey.com))
- **Sales & support (CRM-embedded):** Lead scoring, email sequencing, account research, chat/help-center summaries—baked into Salesforce/HubSpot/Zoho-style platforms. ([McKinsey & Company](https://www.mckinsey.com))
- **Finance/back-office:** Bookkeeping assist, invoice/bill classification, **cash-flow forecasts**, payroll/tax prep—QuickBooks reports strong small business uptake and benefits. ([QuickBooks](https://www.quickbooks.com))
- **Workflow “glue”:** Low-code automations (Zapier/Make) hand off AI-cleaned text, tags, and summaries across forms → CRM → email → invoices. (Reflected across the adoption and productivity studies above.) ([QuickBooks](https://www.quickbooks.com))

Current Small Business Use Cases

- **Lead → Proposal (services):** Web/WhatsApp intake → AI tags need/priority + drafts first email → call notes summarized to CRM → proposal auto-filled → human-approved follow-ups. ([McKinsey & Company](#))
- **Content → Campaign:** Keyword brief → AI outlines posts + ad variants → images/thumbnails → scheduled posts → weekly AI performance summary for repurposing. ([McKinsey & Company](#))
- **Quote → Cash (trades/events):** Site intake → AI estimate + parts list → chatbot answers basics → e-sign + invoice → AI nudges late-pay risk. ([QuickBooks](#))

Navigating the AI Firehose (our Playbook is the filter)

It's not your imagination—the pace is dizzying. That's precisely why this Playbook exists: to point you to **vetted, field-tested tools** and keep guidance fresh as a **living document**.

- **Research → Features pipeline is exploding.** arXiv set a new record with **24,226 papers in October 2024**—a level of monthly submissions sustained into 2025, feeding nonstop model updates. [blog.arxiv.org+1](#)
- **Model ecosystems multiply fast.** The Hugging Face Hub reports **well over 1.5 million models**, reflecting a steep acceleration in releases and forks. [Kiplinger](#)
- **Tooling stacks are dense and shifting.** Even one slice (MLOps) includes **90+ distinct tools/platforms** as of 2025. Expect continual churn across marketing, sales, creative, and “copilot” suites. [Neptune](#)
- **Broad indicators keep climbing.** The **Stanford AI Index 2025** documents record growth across publications, models, and investment—one reason vendors ship rapid upgrades. [Stanford HAI+1](#)

Start Faster with Tried-&-Tested Platforms (what this Playbook gives you)

New AI launches land daily, but you don't need to chase them. This Playbook is a **living document** that curates **field-tested platforms** and **ready-to-run workflows** so you can get value **this week**, not “someday.”

How the playbook helps you start quicker

1. **Pick the job, not the tool.** Each recommended platform is tied to a specific job (e.g., *lead* → *proposal, sermon study* → *small-group guide*).
2. **Copy the template.** We provide **versioned prompts, checklists, and output formats**—already tuned for these platforms—so your team can plug and play.
3. **Stay current, safely.** We update quarterly; significant changes that materially improve **speed, safety, or quality** are flagged sooner. You keep working while we filter.
4. **Human review, constantly.** Even with strong tools, people make the call—especially where accuracy, privacy, and care for people matter.

What's blocking small teams

The biggest hurdles are familiar: many owners still think “AI isn’t for us,” and worry about cost, security, and regulation; skills and time are tight; data is scattered; ROI is unclear. Current policy and survey work highlights:

(a) a “**not applicable**” mindset as the **primary barrier**, (b) **skills & training gaps** (only a minority invests in AI training), and (c) **unclear ROI/digital readiness**—all solvable with the right starter path. ([Office of Advocacy](#))

Ministry snapshot: where AI fits in church workflows

Ministry adoption is on the rise, with clear human-first boundaries.

- **Adoption:** Pushpay’s 2025 State of Church Tech shows **~45% of church leaders use AI**, an **~80% YoY increase**. ([Worship Facility](#))
- **Pastoral boundaries:** Barna finds only **12% of pastors** are comfortable using AI to *write sermons*, but **43% see value for sermon prep/research**; just **6%** are comfortable with AI for counseling—underscoring “assist, not replace.” ([Barna Group](#))
- **Media & operations:** With livestreaming and digital content now standard, teams layer AI for **captions/translation, clipping, and search**; many report **larger tech budgets** to support ops and communications. ([Worship Facility](#))

Current Ministry Use Cases

- **Sermon study → resources (human-led):** Pastor drafts big idea → AI compiles passages/background + discussion questions → staff refine and verify → AI formats small-group/kids' handouts; media captions/clips on Sunday recordings. ([Barna Group](#))
- **“New here?” → Follow-up:** Web/QR intake → AI tags next step (groups, serve, care) → first-touch draft → assigned to human host; AI prepares a concise profile for the meeting. (Paired with clear disclosure to congregants.) ([Worship Facility](#))
- **Prayer & care triage (assist only):** Requests cluster by theme/urgency → routed to the right team → private summaries for staff; **humans** reply pastorally; AI logs next steps. ([Worship Facility](#))
- **Mid-week media:** Auto-clip sermon into 30–60s reels with captions → AI drafts post copy and verse refs → staff approve and publish; later, AI summarizes engagement questions for the next series. ([Worship Facility](#))

Our answer for small orgs

Start narrow, prove value fast, and scale what works. Pilot **2–3 high-leverage use cases** (e.g., proposals & donor letters, client FAQs, course materials), **measure → template → repeat**. For sensitive data, use **tenant-controlled tools** (Workspace/Microsoft 365). Pair adoption with **4–6 hours of AI basics** and a **data readiness checklist** (where data lives, who owns it, update cadence).

Quick Start (zero cost): Use Google AI Studio first

If you're new or budget-constrained, begin in **Google AI Studio**. It's a free, browser workspace for drafting/refining prompts, controlling creativity (temperature), comparing model outputs side-by-side, and **enforcing structure** (section order, bullet limits, word caps) before moving text into Docs/Slides/Sheets or your LMS/website. Later chapters cover AI Studio in depth.

Who this is for

- **BlessBack Business Academy instructors & coaches** are producing lessons, slides, study packs, and quizzes.
- **Ministry leaders & staff** preparing devotionals, workshops, policies, and training materials.
- **Entrepreneurs & micro-business owners** needing concise websites, service pages, WhatsApp replies, menus/cards, and SOPs.

How to use this Playbook

Each section is a **self-contained module** with a step-by-step workflow, guardrails (privacy/ethics/quality), and paste-ready templates. You don't have to read cover-to-cover—**jump to the job** you need to deliver.

Our approach

People-first • Faith-rooted ethics • Evidence-led • Private by default • Multilingual by design. We start with clear outcomes, constrain the format and tone, require evidence for non-obvious claims, and produce artifacts you can publish with minimal editing. Where relevant, we show two lenses: **Academy production** and **micro-business execution**.

Guardrails you must follow.

Privacy. Keep personal data out of consumer tools; use **Google Workspace/Microsoft 365** or local solutions for anything sensitive.

Evidence & accuracy. Cite non-obvious facts; maintain an **Evidence Log** for research outputs.

Transparency & review. Mark AI-assisted content internally and **requires human review** before publishing.

Visual integrity. Prefer **real people** for testimonies; avoid synthetic people to depict your community.

Accessibility & multilingual reach. Translate materials as needed.

Why we highlight Google AI Studio early

Free on-ramp (practice prompts, dial temperature, compare models) • **Production-minded controls** (structure for copy-ready outputs) • **Ecosystem synergy** (AI Studio → Docs/Slides; images/speech; simple apps) • **Live revision** with Gemini Live (Stream). **If you only try one Google feature, start with AI Studio**—the fastest path from idea to publish-ready text at no cost.

The Business Academy's position on AI

Assist, don't replace. AI is a **servant-tool** to speed research, standardize quality, and support multilingual work—**never** a substitute for human authorship, judgment, or pastoral care.

Human review is required before delivery/publication.

Privacy by default for donor/student/health data (tenant-controlled tools).

Evidence over hype (cite non-obvious facts; keep an Evidence Log).

Real people in visuals (no synthetic people for our community).

Bottom line: We prefer and require human-generated content. We use AI to help small businesses, entrepreneurs, and ministries **work more effectively and efficiently**—not to replace human creativity or pastoral presence.

Endnotes (key stats)

- **Small business adoption:** 58% using gen-AI (2025), 82% of AI-using SMBs grew headcount. (uschamber.com)
- **Knowledge-worker usage:** 75% using gen-AI (2024). ([Microsoft](https://www.microsoft.com))
- **Org-level usage:** 71% use gen-AI in ≥ 1 function (2025). ([McKinsey & Company](https://www.mckinsey.com))
- **Macro productivity potential:** AI can boost labor productivity growth when scaled responsibly. ([OECD](https://www.oecd-ilibrary.org))
- **Small Business Benefits:** 68%+ regular use; **74%** say AI boosts productivity (2025). ([QuickBooks](https://www.quickbooks.com))
- **Barriers:** “Not applicable” belief, security/cost/regulation; skills & ROI gaps (SBA 2025). ([Office of Advocacy](https://www.sba.gov))
- **Ministry adoption:** ~45% of church leaders use AI; tech budgets rising; ~80% YoY growth vs prior year (Pushpay 2025). ([Worship Facility](https://www.pushpay.com))
- **Pastoral boundaries:** Only **12%** comfortable with AI writing sermons; **43%** for prep; **6%** for counseling (Barna). ([Barna Group](https://www.barna.com))

What Does God's Word Say About Artificial Intelligence?

Short answer:

The Bible does not mention AI directly, but it gives clear principles that govern how tools, technology, knowledge, and power should be used.

AI is a tool created by humans, not a being with moral authority, spiritual discernment, or divine wisdom. Scripture consistently teaches that tools serve people, and people remain accountable to God.

1 Technology Is a Tool — Not a Moral Agent

The Bible treats tools as neutral instruments whose value depends on the heart and intent of the user.

Key Scriptures

- *"The prudent see danger and take refuge, but the simple keep going and pay the penalty."* — Proverbs 22:3
- *"Commit to the Lord whatever you do, and He will establish your plans."* — Proverbs 16:3

Principles

- AI has no conscience
- AI has no discernment
- AI has no spiritual authority

👉 Responsibility always stays with the human user

2 Wisdom Comes from God — Not from Machines

AI can process information, but wisdom is a divine gift.

Key Scriptures

- *"If any of you lacks wisdom, you should ask God..."* — James 1:5
- *"For the Lord gives wisdom; from His mouth come knowledge and understanding."* — Proverbs 2:6

Guardrails

- AI may assist thinking
- AI must never replace discernment

- AI should support prayerful decision-making, not substitute it

③ Knowledge Without Wisdom Is Dangerous

Scripture warns that knowledge alone can lead to pride, misuse, or harm.

Key Scriptures

- *“Knowledge puffs up, but love builds up.”* — 1 Corinthians 8:1
- *“Be careful that no one takes you captive through hollow and deceptive philosophy...”* — Colossians 2:8

Applications

- AI outputs must be:
 - Reviewed
 - Tested
 - Contextualized
 - Filtered through biblical values

④ Stewardship Applies to Digital Tools

God calls His people to steward everything entrusted to them, including technology.

Key Scriptures

- *“Whoever can be trusted with very little can also be trusted with much.”* — Luke 16:10
- *“Moreover, it is required of stewards that they be found faithful.”* — 1 Corinthians 4:2

Stewardship Rules

- Use AI to serve people, not exploit them
- Protect privacy and dignity
- Avoid deception, manipulation, or false authority

⑤ AI Must Never Replace Human Responsibility or Pastoral Care

Scripture consistently affirms human leadership, accountability, and care.

Key Scriptures

- *“Plans fail for lack of counsel, but with many advisers they succeed.”* — Proverbs 15:22
- *“Shepherd the flock of God that is among you...”* — 1 Peter 5:2

BBA Boundary

- AI can assist pastors, leaders, and entrepreneurs
- AI must never replace:
 - Pastoral discernment
 - Human empathy
 - Accountability
 - Prayerful counsel

⑥ Discernment Is Required in All New Technology

New tools require testing, not blind adoption or fear-based rejection.

Key Scriptures

- *“Test everything; hold fast what is good.”* — 1 Thessalonians 5:21
- *“Do not conform to the pattern of this world, but be transformed by the renewing of your mind.”* — Romans 12:2

Position

- AI is neither:
 - Inherently evil **✗**
 - Spiritually authoritative **✗**
- AI is:
 - A powerful tool **✓**
 - Requiring discipline **✓**
 - Needing governance **✓**

Final Thoughts

Artificial Intelligence is a human-made tool that must remain under human authority, biblical wisdom, and ethical stewardship. It may assist learning, planning, and creativity—but it must never replace discernment, prayer, or responsibility before God.

Chapter 1: Which AI Platform should I use?

Bottom-line guidance

Select the platform that aligns with your data governance and deliverables. If you have no budget or are just starting, use Google AI Studio first to prototype prompts, select a model/temperature, and define the structure of your outputs. Then, move the polished text into Docs/Slides/Sheets.

Quick picks

- **Starting from zero budget** → **Google AI Studio (free)** for prompt design, structure control, and side-by-side model comparison; then export to Docs/Slides/Sheets.
- **Google Workspace orgs** → **Gemini for Workspace** (Gmail/Docs/Sheets with tenant controls) for anything that touches private data.
- **Long, careful policy/lesson drafts** → **Claude Pro** or **ChatGPT Plus/Business**.
- **Fast web-cited answers** → **Perplexity Pro** for verification and evidence.
- **Ultra-private brainstorming** → **Local models** (LM Studio/Ollama) on your device.

Why start in Google AI Studio (even if you'll publish elsewhere)

1. **Free & focused sandbox** — no subscription barrier; ideal for instructors and owners to iterate quickly.
2. **Temperature control** — set creativity precisely (e.g., 0.3 for procedural agendas; 0.6 for product copy).
3. **Compare outputs** — run the same prompt through two models or temperatures and **pick the cleaner one**.
4. **Structure discipline** — enforce exact section order, bullet counts, and word caps; the model's "creative" drift is contained before you ever touch your design files.
5. **One workspace, many media** — draft text, **generate images or single-speaker audio**, and scaffold **micro-apps**—all inside the same studio.

Platform chooser (what to use when)

- **"I need an assistant that teams can reuse, that speaks in our voice and only uses our approved training materials** → Build a Custom GPT or a Gem (Gemini), depending on your suite."
- **I'm doing web-backed research with citations** → Start with **ChatGPT Deep Research**; cross-check in **Perplexity**; if the document lives in Google/Microsoft, do a second pass in Gemini for secure file grounding.
- **I need to assemble study materials from my own sources** → Use **NotebookLM** for Audio/Video Overviews, quizzes, and study guides; then draft deliverables in Projects/Docs.
- **I must create quick visuals, or voice notes** → Use **AI Studio Media Generation** (images/speech) or light image creation in Gemini chat.
- **I need a tiny tool or microsite** → Prototype in **AI Studio App Builder** or **Firebase Studio**; when stable, export/publish.

Security checklist

- Prefer **Workspace/Microsoft 365** assistants for anything with private personal data; keep consumer tools for non-sensitive drafts.
- **Opt out** of consumer data training wherever available; use incognito/temporary chats when you can.
- Never paste secrets (passwords, donor exports, medical details).
- Keep a **lightweight record of processing** (which tool touched which doc), and **version** your outputs.
- For analysis over Drive/SharePoint content, use **connectors** with the **least-privilege** scope (read-only, AI-safe folders).

Side-by-Side Comparison

Tool	Best for	Free	Paid you'll care about	Ease	Relevance	Safety tier
ChatGPT 5.2 (OpenAI)	Deep reasoning, complex coding, and long-form writing. (New "Thinking" mode excels at logic & policy drafts)	Core chat ("Instant" model)	Plus \$20 (includes "Thinking" & "Pro" modes) or Team/Ent	5	5+	● (consumer) / ● (Team/Ent)
Gemini 3.0 (Google)	Multimodal creativity (video/image) & Workspace grounding. (Leader in vision & real-time editing)	Core chat (Flash model)	AI Premium \$19.99 (includes "Pro" & "Deep Think")	5	5	● (Workspace) / ● (consumer)
Claude 4.5	Nuanced writing &	Core chat	Pro \$20,	5	5	● (consumer) /

Tool	Best for	Free	Paid you'll care about	Ease	Relevance	Safety tier
(Anthropic)	tone control. (Still excellent for human-sounding policy/devotional drafts)		Team/ Enterprise			● (Business)
Perplexity	Real-time web research with citations. (Great for fact-checking the reasoning models)	Core search	**Pro \$20** (Access to GPT-5.2/Clade 4.5 APIs)	5	5	● (consumer) / ● (Enterprise)
Microsoft Copilot	Office 365 integration (Word/Excel/PPT). (Best for corporate tenants on the Microsoft stack)	Web / App	Copilot Pro or M365 Copilot (~\$30/user)	4	5	● (Tenant/Ent)

Ease/Relevance scale: 1–5 (higher is better). Prices are indicative; please check our current plans for details.

● **Green (Safest):** Stays inside your org suite or device. Not used to train public models by default.

● **Yellow (Caution):** Consumer apps. Check your settings to opt out of training and avoid pasting sensitive data.

● **Red (Avoid for sensitive):** Unclear data residency/collection. Use only for non-sensitive text.

Key Points

- **ChatGPT 5.2:** Now features specific "**Instant**" (speed) and "**Thinking**" (depth) modes. Use "Thinking" for lesson planning and policy reviews; use "Instant" for quick emails.
- **Gemini 3.0:** The new "**Deep Think**" capability makes it a viable backup for reasoning tasks, but its primary edge is now **visual/multimodal** work (e.g., analyzing photos of physical receipts or creating class visuals).

Guardrails for Ethical & Effective Use

- Cite sources for research outputs; avoid hallucinated stats.
- Mark AI-assisted content internally; human review before publishing.
- Respect copyright & licensing.
- Be transparent with students/clients when AI meaningfully contributed.
- Favor **real people** in materials and avoid synthetic people in testimonies.

ChatGPT 5.2 — Overview & Why We Recommend It

1) What ChatGPT 5.2 Is

ChatGPT 5.2 is OpenAI's latest unified chat system. It explicitly splits the experience into two distinct engines behind one interface:

- **"Instant" Mode:** An ultra-fast, efficient model for everyday tasks (emails, quick translations, summaries).
- **"Thinking" Mode:** A deep reasoning model that pauses to "think" before answering complex problems (strategy, policy analysis, coding).
- **Auto-Router:** An intelligent switcher that defaults to the correct mode based on your prompt complexity.

Improvements vs. earlier versions:

- **Bifurcated Performance:** You no longer have to trade off speed for intelligence. "Instant" handles the bulk work instantly, while "Thinking" handles high-stakes logic better than v5.0.
- **Deeper Reasoning:** The 5.2 "Thinking" mode excels at multi-step work (curriculum mapping, risk mitigation, conflict resolution) with fewer logic errors.
- **Expressive Control:** Refined controls for tone and structure make it easier to draft distinct voices (e.g., "warm pastoral" vs. "crisp executive").
- **Unified Tooling:** **Deep Research**, **Projects**, and **Agent Mode** are now deeply integrated, allowing the model to seamlessly switch between web research and your internal project files.

Why we recommend it as our lead platform for this playbook:

- It perfectly supports our "Deep Research → Draft → Multilingual Output" cycle.
- **Project Context:** The **Projects** feature keeps your specific files (brand guides, policy docs) and chat history together, preventing the model from "forgetting" your context.
- **Evidence & Citations:** The integrated **Deep Research** tool is the gold standard for producing the cited synthesis and Evidence Logs required by our policy.
- **Cost/Speed Balance:** The ability to toggle "Instant" for routine tasks helps save high-cost credits for when you really need the "Thinking" engine.

2) Modes & When to Use Each

You'll see four options when working with GPT-5. Use this as your default guide:

A) Auto (recommended default)

- **What it does:** Let's ChatGPT pick the right balance of speed/depth.
- **Use when:** You're not sure which mode to pick; for most everyday tasks.
- **Examples:** Draft emails, lesson outlines, grant paragraphs, sermon series ideas.

B) Instant (sometimes shown as Fast)

- **What it does:** Prioritizes speed with lighter reasoning.
- **Use when:** You need quick, short answers; summarizing short text; brainstorming options.
- **Examples:** Subject lines, social captions, quick translations, outline variants.

C) Thinking

- **What it does:** Allocates extra compute to reason more carefully before answering.
- **Use when:** Complex analysis, multi-step planning, tricky logic, policy comparisons, data cross-checks, and rubric creation.
- **Examples:** Market/partner briefs, curriculum mapping, risk/mitigation tables, bilingual policy drafts.

D) Pro (plan/tier, not a model)

- **What it is:** A high-capacity subscription tier with priority access to the newest models/features and higher usage limits (applicable for teams who run heavy workloads, long voice sessions, or agentic tasks).
- **Use when:** You or your team hit usage limits, rely on advanced voice/agent mode, or need peak performance during busy hours.

Quick rule: Start **Auto** → bump to **Thinking** for depth; switch to **Instant/Fast** for rapid iterations. Consider **Pro** if you often see rate-limit notices or run long, complex sessions.

3) What's Different vs. Previous Models

Area	Earlier models (GPT-4 class, o-series)	ChatGPT 5
Area	Earlier models (GPT-4 class, o-series)	ChatGPT 5.2
Model selection	Users picked a specific model (e.g., 4o, o1); separate “reasoners” were required for complex problems	Unified experience with Auto router + explicit Thinking (depth) and Instant (speed) modes
Reasoning	Good but required careful prompting; long “think” models were separate and slow	Built-in Thinking mode for complex tasks; router switches automatically, or you can force it for profound logic
Coding	Strong code gen; debugging required manual back-and-forth	More end-to-end coding & better debugging/design help (Agentic behavior)
Writing	Great general writing	More expressive control (structure, tone), better long-form coherence for policies & curricula
Tools	Browsing, vision, and voice varied by model capabilities	Tools consolidated; Deep Research, Projects, Voice, and Agent Mode are first-class and integrated

4) Practical Mode Picker — Scenarios

- **Sermon or lesson plan with sources?** Start **Auto**, add: “Use **Thinking** to map sections and cite sources.”
- **Grant paragraph rewrites (many drafts)?** **Instant/Fast** for speed → switch to **Thinking** for the final.
- **Market scan or policy comparison?** **Thinking** from the start, ask for an **Evidence Log** table.
- **Slide outline + bilingual handout?** **Auto** to draft → **Instant** to iterate titles → **Thinking** about fact checks.
- **Long voice brainstorming?** **Auto** (or **Pro** tier if you hit limits).

Prompt tip: To nudge depth, add: “Use **Thinking** mode for careful reasoning and cite every non-obvious claim.”

5) Setup & Best Practices

- **Pick your default:** Leave **Auto** on; learn to toggle **Instant** and **Thinking** as needed.
- **Use Projects:** Keep objective and core files in each project.
- **Deep Research + Evidence:** For non-obvious claims, require citations and a brief uncertainty note.
- **Privacy:** Avoid pasting private information/donor lists into consumer chats; prefer Workspace/Microsoft365 or local drafting for sensitive content; use team-only sharing for internal work.
- **Versioning:** Add vX.Y + date to key outputs and keep an Evidence Log.

6) Limitations & Gotchas

- **Routing isn't perfect:** Auto may choose **Instant** when you need **Thinking**—toggle manually.
- **Rate limits:** Heavy **Thinking** usage can hit caps on non-Pro tiers.
- **Citations:** Still require verification; ask for quotes + working links.
- **Cost/latency trade-offs:** **Thinking** is slower; use it where accuracy matters.

Things to Know — Make ChatGPT Efficient & Effective

1) Start Every Chat with Context & Constraints

- **Objective:** one sentence (e.g., “Draft a 1-page donor brief on... for pastors”).
- **Audience & tone:** who it's for; warm/practical/concise.
- **Format & length:** bullets, table, slide outline; word/character count.
- **Language:** can translate into over 100 languages.
- **Success criteria:** “Cite non-obvious claims; include an Evidence Log; add next steps.”
- **When to ask questions:** “Only ask if missing info blocks; otherwise, proceed with reasonable assumptions.”

2) Use Projects for Anything Ongoing

Create a Project, add custom instructions, upload a small set of core material (5–30 files), and maintain tidy versions (vX.Y). Run **Deep Research** inside the project for web-backed claims.

3) Mode & Tool Tips

- **Auto** by default → switch to **Thinking** for complex analysis/planning → use **Instant/Fast** for rapid iterations.
- Use **Deep Research** for multi-source facts with citations; request an **Evidence Log**.
- For **sensitive private information/donor/health** content, stay in Workspace/Microsoft365 or draft on a **local model** and paste only redacted snippets.

4) Review Cycle (Quality & Integrity)

- Ask for a **plan-of-attack** (sub-questions + sources).
- Generate the draft + **Evidence Log**.
- Run the **Pre-Publish Checklist** (citations, conflicts, numbers, privacy, bilingual, accessibility).

Evidence Log (copy/paste): | Source | Claim/Data | Verbatim Quote/Figure | URL | Date Accessed | Confidence | |—|—|—|—|—|—|

5) Common Fixes

- **Messy outputs** → restate format with a small example; cap words; ask for a numbered outline first.
- **Hallucinated facts** → require quotes + links; cross-check in Perplexity.
- **Unreadable PDFs** → It's likely a scan; export a text-based PDF or run OCR.
- **Slow/limits** → reduce file count; split tasks; switch to **Instant** for brainstorming.

6) Adding Photos & Files — What You Can Upload & Limits

Where can you add files

- **In a chat:** click **+ Add photos & files** (or drag-and-drop) to reference a file in the current conversation.
- **In a Project**, attach files once; all chats inside the project can reuse them.
- **In a Custom GPT (Knowledge)**, attach a small, curated set of core materials so your bot can answer from organizational materials.

Supported file types (practical examples)

Documents & text — .pdf, .docx, .pptx, .txt, .rtf, .md, .json (and most common text-based formats).

Spreadsheets & data — .csv, .xlsx (and other standard spreadsheet files).

Images (photos/screenshots) — .png, .jpg/.jpeg, .gif (non-animated).

Code & web assets — .py, .ipynb, .html, .css, .js, .ts (and other standard code files).

Notes

- **Google Docs native files (.gdoc)** aren't supported—export to **.pdf** or **.docx** first.
- Chat can analyze **static images**; **video files aren't supported** for direct analysis today.

Size & count limits (know before you upload)

Area	Limit
Per file (any type)	Up to 512 MB
Images	Up to 20 MB per image
Text/Document processing	Up to ~2,000,000 tokens per file (does not apply to spreadsheets)
Spreadsheets (CSV/XLSX)	Practical cap ~50 MB (depends on row size)
Project file count	Go/Plus: up to 20 files per project • Pro/Business/Edu/Enterprise: up to 40 files per project.
Simultaneous uploads (Projects)	Up to 10 files at once
Custom GPT "Knowledge"	Up to 20 files attached; 512 MB per file; 2M tokens per file (images allowed, but only text is indexed)

Token cap tip: A long PDF might hit the **2M-token** ceiling even if its file size is <512 MB. Prefer concise, text-based PDFs over image scans.

File-prep best practices

- **Use text-based PDFs** (export from the source); avoid image-only scans when you can.
- **Export Google Docs/Slides** to **.pdf** or **.docx/.pptx**; keep filenames clear (e.g., **BBA_BrandGuide_v2025.pdf**).
- **Split large spreadsheets** (>50 MB) into logical chunks (by year, sheet, or category).
- **Compress big images** below **20 MB** (1920×1080 is usually enough for analysis).
- **Keep a small core of materials**: 5–30 high-signal files; remove duplicates and superseded versions.
- **Protect privacy**: redact sensitive data before uploading to consumer chats; prefer Workspace/Microsoft 365 assistants or local models for sensitive drafts.

Troubleshooting

- **“File too large/token limit exceeded”** → compress/split; convert scans to text; trim appendices.
- **“Wrong file type”** → convert Google formats (.gdoc) to **.pdf/.docx**; ensure images are **PNG/JPEG/non-animated GIF**.
- **“Upload stalled”** → try fewer files at once (≤ 10), smaller images, or a stable network.
- **“Model can’t read my PDF”** → your PDF may be image-only; run OCR or export a text-based version.

Quick decisions

- **One-off sharing** → attach in a **Chat**.
- **Multi-chat reuse** → attach to a **Project**.
- **Repeatable assistant** → attach as **Knowledge** in a **Custom GPT**.

Connectors — What They Do, How to Set Them Up, and How to Use Them Safely

1) What Connectors Are (and why they matter)

Connectors securely link ChatGPT to outside content or apps (e.g., Google Drive/Docs, Microsoft OneDrive/SharePoint, Dropbox, Notion, Slack, Sheets). With a connector, the model can **find, read, and reference** files you can access—so it can summarize, quote, compare, and build deliverables **without copying and pasting**.

Typical actions via connectors - Search your Drive/SharePoint by title/content; open a doc; quote and cite sections - Read a Sheet/CSV to compute tables and charts - Pull messages/notes (e.g., Slack/Notion) to build a brief - (Optional) Draft an email/post saved as a doc for human send later

Necessary: A connector does **not** give ChatGPT access to your entire organization by default. Access mirrors the account/permissions you grant, and you can restrict it to specific folders.

2) When to use a connector vs. uploads or Projects

- **Connector** → recurring access to **many** living documents (policies, curricula, metrics) you don't want to re-upload. Best for **Projects** and **Agent Mode**.
- **File upload** → quick, one-off analysis of a few files (PDF, CSV, image).
- **Project (no connector)** → when all relevant files are static and can live inside the project.

3) How to Create a Connector (step-by-step)

- **Choose the scope**
Create a dedicated **AI-safe folder** (e.g., BBA_AI_Core Materials) with only the docs you want the model to read. Do **not** point to your entire Drive/SharePoint root.
- **Use a work identity**
Sign in with **org-managed** credentials (not personal). Prefer a **service account** for shared agents.
- **Connect**
In ChatGPT → **Settings** (or within a Project/Agent setup) → **Connectors/Integrations** → choose provider → grant access. Select the **minimum folder(s)** needed.
- **Confirm permissions**
Verify the connector shows the correct folders. Check **read-only** where possible; turn off write/delete unless required.

- **Test**

In a Project, ask: “Search connector for *BrandGuide* and summarize v2025; quote the revision policy section.” Confirm the quotes match the source.

- **Document**

Log owner, folders, scopes, and a **review date**. Add to your **Connector Registry** (see the template below).

4) Security & Privacy — What to know

- **Least privilege:** connect only to the folders you need; prefer read-only.
- **Data sensitivity:** keep **PRIVATE INFORMATION** in your secure suite; you may use the connector to **read** summaries, but avoid storing raw PRIVATE INFORMATION in prompts. Redact when possible.
- **Caching & indexing:** Some connectors create a temporary index to enable the model to search. Treat connected content as **accessible to ChatGPT** per your plan’s privacy terms; use enterprise/tenant controls when needed.
- **Auditing:** keep a registry of who connected what to which folders, when; review quarterly.
- **Revocation:** You can remove the connector in ChatGPT and in the provider (Google/Microsoft) admin panel. Create a **break-glass** procedure (see checklist).
- **Jurisdiction:** prefer **Workspace/Microsoft 365** native assistants for highly sensitive content; use ChatGPT connectors for non-sensitive or published materials unless you have enterprise assurances.

5) Best Practices

A) Scope & structure - One **AI-safe shared drive/folder** per team (e.g., BBA/AI_Core Documents/Policy, BBA/AI_Core Documents/Curriculum). - Keep **latest-only** docs; archive superseded versions; name files clearly (BBA_BrandGuide_v2025.pdf).

B) Governance - Maintain a **Connector Registry**: owner • folders • scopes • purpose • review date • notes. - Quarterly review: remove stale folders; rotate service-account credentials.

C) Prompts that protect - “Cite non-obvious claims with quotes + links to the connected doc.” - “If you cannot access a doc due to permissions, say so and stop.” - “Never write or delete files without explicit approval.”

D) Pair with Projects - Create a Project per initiative and add connector search as a first step: plan → locate sources → produce brief → archive Evidence Log.

E) Redaction & staging - For PRIVATE INFORMATION-heavy sheets, create a **staging copy** with masked identifiers; point the connector to staging.

6) Pros & Cons of Connectors

Aspect	Pros	Cons / Risks	Mitigations
Freshness	Always reads the latest docs; no re-uploads	If the source changes mid-task, outputs may drift	Version/date stamps in outputs; lock versions for publishing
Scale	Search across many files quickly	May index broader content than intended	Folder-level scoping; AI-safe drive; read-only
Speed	Faster than manual uploads for big projects	Initial setup time; auth occasionally expires	Registry + quarterly checks; service accounts
Quality	Pulls quotes and tables directly	Mis-scoped searches can pull the wrong version	File naming discipline; “prefer v2025” prompt rule
Security	No copy-paste sprawl; central control	Connector is a new trust surface	Least privilege; revocation plan; enterprise controls

7) Google Drive Connector — Setup & Usage (Step-by-Step)

A) Prepare your Drive (least-privilege by design)

1. **Create an AI-safe Shared Drive** (recommended): BBA_AI_Core Materials
 - o Folders like: /Policy, /Curriculum, /Metrics_Staging, /Brand.
 - o Put **only** the documents you want ChatGPT to read. Avoid root-level access.
2. **Permissions**
 - o Grant **Viewer** (read-only) to the user/account you'll connect.
 - o Avoid Editor/Manager unless a workflow truly needs write access.
 - o For team-run agents, consider a dedicated **integration account** (org-managed) and add it as a **Viewer** to BBA_AI_Core Materials.
3. **Sensitive data staging**
 - o Create /Metrics_Staging with **redacted** or **masked** sheets (no raw PRIVATE INFORMATION).
 - o Use clear file names and versions: BBA_BrandGuide_v2025.pdf, SafeguardingPolicy_v2025.docx.

B) Connect ChatGPT to Google Drive

1. In ChatGPT → Settings → Connectors / Integrations → choose **Google Drive**.
2. Sign in with the **work identity** (or integration account).
3. When prompted, **restrict scope** to the Shared Drive/folder(s) you prepared (do **not** grant “all of My Drive”).
4. Finish auth; you should now see Drive as a connected data source.

Tip: With a connector, ChatGPT can read **native Google formats**

(Docs/Sheets/Slides). *Uploads* still require exporting to PDF, DOCX, or XLSX, but **connectors** can read Google Docs/Sheets directly.

C) Test & validate

Paste these in a Project:

- “Search the **Drive connector** for BBA_BrandGuide_v2025; summarize the tone section and **quote exact lines** with a link.”
- “Find the latest SafeguardingPolicy_v*.docx and output a **change log** of what changed since 2024.”
- “Open Metrics_Staging/Monthly.xlsx and produce a table of KPI deltas. Include an **Evidence Log**.”

D) Daily use patterns

- **Locate & quote:** “Search Drive for ‘lesson outline v2025’, list matches with modified dates; open the newest and quote the learning outcomes.”
- **Assemble a brief:** “From Drive: policy + lesson + last month’s metrics → build a 1-page brief with citations and summary.”
- **Safety prompt:** “If you cannot access a doc due to permissions, **say so and stop**. Do not request broader access.”

E) Governance & security

- **Registry:** Log connector owner, folders, permission level, purpose, **review date** (quarterly).
- **Audits:** Use Google Admin audit logs to check access by the integration account.
- **Revocation:** If you need to cut access:
 1. Remove the connector in **ChatGPT**
 2. In **Google Admin**, revoke the app’s OAuth access, or remove the integration account from the Shared Drive.
 3. Rotate the integration account password/2-2-step verification

- **Download controls:** You can keep viewers as view-only; if “Disable download/print/copy” is enabled, ChatGPT may not be able to quote—test with a sample file.

F) Troubleshooting

- **“Permission denied / Can’t see file”** → Ensure the connected identity has **Viewer** on the specific Shared Drive/folder; avoid shortcuts that point outside the scope.
- **“Wrong version pulled”** → Use explicit names (e.g., v2025), and prompt: “Prefer files with v2025 in name; if multiple, pick the latest modified.”
- **“Content not searchable”** → Convert image-only PDFs to text-based PDFs (run OCR in Drive → Open with Google Docs → Save).
- **“Too broad results”** → Narrow scope by folder path (“Search only BBA_AI_Core Materials/Curriculum”).
- **“Index lag”** → Newly added files might take a few minutes to surface—upload, wait briefly, then retry.

G) Prompts you can copy

- “Search Drive → BBA_AI_Core Materials//Policy for **Safeguarding**. Summarize responsibilities and quote the reporting steps, using section headers.”
- “List all lesson outlines in BBA_AI_Core Materials/Curriculum updated in the **last 90 days**; output a table (title, path, modified date, owner).”
- “Open Metrics_Staging/Monthly.xlsx → compute completion rate vs. prior month; add 1-sentence analysis; provide HT summary on request.”

Chapter 2: Settings Quickstart (ChatGPT + Gemini)

BBA Default (Use for Both Platforms)

- Opt out of consumer data training** (when available)
- Turn OFF “memory / past-chat personalization”** for new users
- Use Temporary/Incognito chats** for sensitive or one-off work
- Avoid pasting secrets** (passwords, banking details, donor exports, counseling notes)
- Connect apps only when needed** (least privilege; AI-safe folders)

ChatGPT (OpenAI) — Quick Setup

1) Training / Data controls (privacy baseline)

- Go to: **Profile** → **Settings** → **Data Controls**
- Turn off “**Improve the model for everyone**” (wording may vary)
 - Result: Your chats can still appear in history, but they **won't be used to improve the model**.

2) Temporary Chat (use this for sensitive work)

- Start a **Temporary Chat** for sensitive or one-off tasks
 - Temporary Chats: **not saved in history**, **don't create memories**, and are **deleted within 30 days**.

3) Personalization controls

A) Memory (recommended default: OFF for new users)

- Settings** → **Personalization** → **Memory** → **OFF**
 - Use Memory only if the user wants ongoing preferences saved (and the work is not sensitive).

B) Custom Instructions (recommended: ON, but non-sensitive)

- Settings** → **Personalization** → **Custom Instructions**

Safe examples to store:

- “Write in clear bullet points.”
- “End with next steps.”
- “Ask 1–2 clarifying questions only if required.”

Do NOT store: donor names, private staff details, passwords, internal security practices.

C) Voice / audio / files (practical guardrail)

- If using voice or uploading files, treat content like it could be stored in chat history unless using **Temporary Chat**.
- For sensitive topics: **Temporary Chat only**.

4) Security (recommended)

- Turn on **2-step verification / MFA** on the account (and keep device locked).

Gemini (Google) — Quick Setup

1) Training + history control (“Keep Activity” / “Gemini Apps Activity”)

- Go to: **Profile → Gemini Apps Activity** (or “Keep Activity”)
- Turn OFF** (or “Turn off and delete activity”)
 - With Keep Activity OFF: Gemini chats **won't be saved to activity** and Google states they may be retained up to **72 hours** for service delivery/safety.

2) Temporary Chat (use this for sensitive work)

- Start a **Temporary Chat** when you don't want it to influence future chats
 - Temporary Chats: **won't appear in recent chats or Gemini Apps Activity**, and **won't be used to personalize or train**; kept up to **72 hours**.

3) Personalization controls

A) Past-chat personalization (recommended: OFF for new users)

- Keep Activity OFF to reduce history-based personalization.
- Use Temporary Chat for anything sensitive.

B) Connected apps / “Personal Intelligence” style features (recommended: OFF by default)

- Only connect apps (Gmail/Drive/Photos/Calendar/etc.) when needed.
- Prefer AI-safe folders and avoid sensitive data sources.

C) Audio + Gemini Live recordings (recommended: leave OFF)

- Look for settings like “**Improve Google services with audio/Gemini Live recordings**”
- Leave OFF unless you explicitly want to opt in.

D) Auto-delete (optional)

- If Keep Activity is ON for any reason, set **Auto-delete** to the shortest acceptable window.

4) Security (recommended)

- Turn on **Google 2-step verification** and review connected devices periodically.

10-Second “Before You Paste” Filter

If it contains **passwords, bank info, donor exports, counseling details, student personal data**, don’t paste it into a consumer chat.

Instead: anonymize, summarize, or use a secure internal system.

Chapter 3: Deep Research with AI

1) What “Deep Research” Is (and isn’t)

Deep Research is a guided, multi-step investigation where the AI plans an approach, searches across sources, reads them, cross-checks facts, and then produces a structured, cited output (summary, brief, report, slide outline, etc.). It excels when you need synthesis across many sources with transparent reasoning.

What it **does**:

- Plans sub-questions and a research path (decomposes the task)
- Searches and opens sources; reads, quotes, and compares
- Tracks claims vs. evidence; flags uncertainty
- Produces a deliverable in your requested format (brief, table, outline, slides)

What it **does not** do by itself:

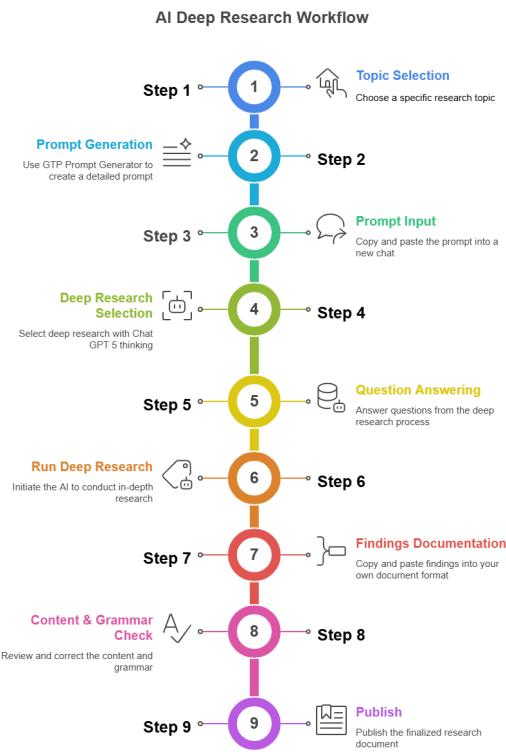
- Replace expert judgment, fieldwork, or confidential interviews
- Automatically know your private files unless you provide or connect them
- Guarantee 100% accuracy—human review is required

Use cases:

1. Business or ministry market studies at local, country, and global levels.
2. Educational market studies at local and country levels.
3. Competitor research at local, country, and global levels.
4. Industry segmentations of countries.
5. Country policies
6. Curriculum research.
7. Sermon topic research.
8. Collaboration partner study research needs.
9. Grant backgrounders.

2) Deep Research Workflow

The steps below align with your “AI Deep Research Workflow” graphic and add actions, outputs, and guardrails.



Confidentiality: Use Workspace/Microsoft365 assistants for sensitive docs; avoid pasting PRIVATE INFORMATION into consumer chats; prefer local models for high-sensitivity drafts.

Source quality: Prefer primary/official sources; identify bias; surface disagreements.

Traceability: Keep an evidence log; never rely on uncited claims.

Localization: Produce English, Spanish, and Haitian Creole versions where relevant.

Ethics: Be transparent when AI contributed; get consent for quoted material.

Output: Initialized session with context assets.

Step 4. Deep Research Selection

Goal: Choose the right engine.

- Start with **ChatGPT Deep Research** for broad synthesis + strong writing.
- If heavy web verification is needed, plan to cross-check with **Perplexity**.
- If drafting in Google/Microsoft docs, consider **Gemini/Copilot** grounding.

Output: Selected tool(s) and rationale.

Step 5. Question Answering (interactive)

Goal: Steer the run by asking clarifiers.

- Ask the model to list sub-questions and planned sources before diving.
- Request an **evidence log** table (source → claim → confidence).

Output: Plan-of-attack + live Q&A thread.

Step 6. Run Deep Research

Goal: Execute multi-step retrieval & analysis.

- Require citations for non-obvious facts; ask for competing views.
- Instruct: “*Mark anything uncertain; propose how to resolve.*”

Output: Draft findings with citations and uncertainty notes.

Step 7. Findings Documentation

Goal: Format the results into your preferred format.

- Ask for deliverables as:
 - **Executive brief** (1 page)
 - **Appendix** (sources, quotes, links)
 - **Tables** (metrics, comparisons)
- Paste into your doc or slides; keep the evidence log attached.

Output: Clean brief + appendix + tables.

Step 8. Content & Grammar Check

Goal: Tighten language & fix structure.

- Run a “clarity & tone” pass; keep quotes verbatim.
- Translate into **another language** as needed.

Output: Polished, bilingual-ready text.

Step 9. Publish

Goal: Share the content with your audience and archive it for future reference.

- Add a cover, date, and version.
- Store the evidence log with the report for transparency.

Output: Final PDF/slide deck + archived sources.

3) Quality & Integrity Checklist (Pre-Publish)

Evidence & Citations

- Every non-obvious claim has a citation and a date of source access.
- Links work; archive/screenshot captured for key sources
- Conflicting sources summarized; competing viewpoints represented
- Direct quotes marked and kept short; paraphrases clearly paraphrased

Numbers & Methods

- Units/currency consistent (note USD/NIO/RD\$ and conversion date)
- Timeframe explicit; denominators and sample sizes stated
- Tables add up; percentages $\approx 100\%$; no double-counting
- Method box included (scope • sources • limitations)

Bias, Limitations & Confidence

- Bias/limitations section present (coverage gaps, paywalls, language)
- Confidence level provided for each key claim (High/Med/Low) with why

Compliance & Privacy

- Private individual information/donor details removed or anonymized; sensitive work done in Google Workspace.
- Copyright honored; licenses cited for images/figures; fair-use quotes ≤ 25 words

Reproducibility

- Evidence log attached (source → claim → quote → link → confidence)
- Versioning set (vX.Y + date) and file path noted; sources archived

Localization & Clarity

- Local language summaries included (where the audience needs)
- Jargon minimized; acronyms expanded; glossary for technical terms if used
- Readability check (target Grade 9–11)

Accessibility & Design

- Alt text for figures; clear legends and labels
- Color contrast and font sizes suitable for print/projector; page numbers present

Executive Brief

- 1-page summary with 3–5 key findings, implications, risks, and subsequent actions

Sign-offs

- Program owner review completed
- Director approval for public release
- Translator QA completed for bilingual outputs

Post-Publication

- Update cadence set (e.g., quarterly); errata process documented.

4) Which Deep Research Tool When?

Tool	What it's best at	Strengths	Watch-outs	Use when...
ChatGPT Deep Research	Broad synthesis + clear writing	Plans multi-step research; strong drafting; flexible deliverables	Needs careful citation discipline; confirm stats	You need a polished brief/slides outline with analysis, not just links
Perplexity Pro/Deep Research	Fast, cited web answers	Source-first interface; good de-duplication	Can be terse; occasionally misses paywalled/grey literature	You need rapid, well-cited web findings to cross-check or seed a draft
Gemini (Workspace / AI Pro)	Research <i>inside</i> the Google ecosystem	Native Gmail/Docs/Sheets grounding; good summarization of Drive files	Web answers vary by query framing	You're already drafting in Google Docs and want

Tool	What it's best at	Strengths	Watch-outs	Use when...
				citations + Drive context
Claude (Pro/Team) + Browse	Nuanced reasoning and long-form writing	Excellent clarity; handles long context and policy tone	Browsing depth can be conservative; cite requests explicitly	You need careful prose (policies, devotionals, curricula) with referenced sources

Rule-of-thumb:

- Start in **ChatGPT Deep Research** for complex synthesis.
- Cross-check claims in **Perplexity** (quick citations).
- If the document lives in **Google** or **Microsoft**, run a second pass in **Gemini** to securely pull in Drive sources.
- For policy or devotional tone pieces, consider a **Claude** polish pass.
- For high-sensitivity content, draft on a **local model** and add public citations afterward.

5) Troubleshooting & Pitfalls

- **Hallucinated citations**: Ask for direct quotes + working links; verify each.
- **Over-confident claims**: Force a **Confidence** column (High/Med/Low) with justification.
- **Too general**: Narrow scope (time, geography, sector) and add required tables.
- **Paywalls**: Request alternative open sources or official summaries.
- **Long results**: Ask for a 1-page executive brief + appendix.

6) Deep Research Field Tests

How We Conducted the Tests

In a **cold-start** setup, we opened a brand-new chat and asked the question without warm-up **prompts or examples**, so the results reflect each platform's baseline performance rather than prior tuning. Each result was scored using the BBA rubric (0–100), with six weighted criteria and three pass/fail gates.

Scenarios tested

#1: WhatsApp vs Telegram for small-business sales in low-bandwidth LATAM
(ChatGPT & Gemini runs). Key lens: adoption stats, bandwidth/security notes, commerce features, concrete examples from 2022–present. Scores and gap notes below.

#2: Back-to-School Uniform Business Plan (Nicaragua via ChatGPT; Honduras via Gemini). Focus: unit economics, local calendars/fees, live price examples, and “Limits & Unknowns.”

#3: Faith-Based Youth Entrepreneurship Frameworks for NGOs (LATAM)
(ChatGPT & Gemini runs). Focus: faith/livelihood models, outcome evidence, and LATAM fit.

Rubric (weights) & Gates used

- Weights: Source Quality (25%), Traceability (25%), Date Freshness (10%), Specificity/Context (15%), Accuracy/Consistency (20%), Ethics/Limits (5%). Examples of scored tables appear in each scenario’s analysis.
- Gates applied to every run:
 1. **≥80% links open/relevant,**
 2. **≥1 short direct quote** from an official/primary source when the task involves policy/platform claims,
 3. **No PII/unsafe guidance.**

What We Found (By Scenario)

Scenario #1 — WhatsApp vs Telegram (Central America, low bandwidth)

- **Gemini run:** 61/100. Solid structure/LATAM focus, but heavy reliance on blogs/aggregators; **no short official quotes** (gate fail). Upgrade needed on adoption/bandwidth claims and platform-feature quotes.
- **ChatGPT run:** 69/100. Better traceability and regional tailoring; still mixed source quality and **no official short quote** (gate fail). Needs primary stats (DataReportal/Pew/Sensor Tower) and carrier T&Cs; trim blogs.

Common risks across both runs

Over-use of explainers/aggregators for: E2EE differences, MAU/adoption, bandwidth per minute, and “zero-rating” claims; missing short quotes from official pages (e.g., Telegram API limits, WhatsApp E2EE).

Scenario #2 — School Uniform Micro-Business Plan

- **ChatGPT (Estelí, Nicaragua): 86/100.** High-quality local evidence; includes a **short direct quote** for calendar timing; good traceability and explicit limits. Minor gaps: some vendor prices lack durable URLs; BOM uses estimates (needs receipts). **Gate status: Pass/Pass/Pass.**
- **Gemini (San Pedro Sula, Honduras): 59/100.** Ambitious and fresh, but the works-cited list includes low-authority sources; price table traceability is inconsistent, and a **margin inconsistency** appears (target ~44% vs worked example **24.4% at L500**). **Gate status: Link integrity borderline/Fail; quote gate not met.**

Scenario #3 — Faith-Based Youth Entrepreneurship Frameworks (LATAM)

- **ChatGPT:** 67/100. Focused on faith actors; breadth ok, but outcome claims need page-level anchors; freshness mixed; LATAM micro-loan specificity thin. **Use with edits.**
- **Gemini:** 76/100. Broader scan with more gov/edu/IGO sources; still short on page/quote anchors and LATAM faith-based localization; impact stats often org-reported. **Use with edits.**

Cross-Scenario Conclusions

1. **Evidence discipline beats model choice.** The top score (86) came from the run that paired local, checkable sources with at least one **short official quote** and an explicit Limits section—exactly what our Playbook requires.
2. **Quote gate is decisive for platform/policy tasks.** Both “WhatsApp vs Telegram” runs missed the **short-quote** gate despite decent structure. We now treat one ≤ 25 -word official quote as mandatory for platform or policy claims.
3. **Traceability gaps sink business math.** When price tables lack shop-level URLs/screenshots, conclusions feel speculative; a single margin mismatch (target 44% vs computed 24.4%) lowers trust even when the plan looks strong.
4. **Breadth ≠ verification.** The youth-framework scans cite many reputable orgs, but without page/line anchors, outcome claims stay “plausible” rather than “proven.” We now require an **Evidence Log** with claim • URL • quoted line for key metrics.

Designing Cold-Start Prompts (to get strong results on the first run)

Why this matters (from our tests):

- Runs missed our “short policy quote” gate and used weak/aggregator sources when the prompt didn’t force **primary docs + ≤25-word quotes**.
- Adoption, bandwidth, and pricing claims were hard to verify because the prompt didn’t require **page/line anchors or an Evidence Log**.
- The strongest report (Estelí uniforms) succeeded because the brief **forced locality, quotes, and limits**.

A. Cold-Start Prompt Checklist (copy/paste into your prompt)

- **Role & Audience:** “You are an evidence-first analyst for BBA in [country/sector].”
- **Scope & Locale:** “Limit to [countries/region]; prioritize 2022–present.”
- **Source Rules:** Primary or official first; allow reputable secondary only if triangulated; **no blogs/aggregators for core claims**.
- **Quote Gate:** Include **≥1 short quote (≤25 words)** from an official doc or equivalent (policy, T&C, API) with a URL + page/section.
- **Evidence Log:** Return a compact table: *claim* → *primary URL* → *quoted line(s)/page* → *date* → *confidence*.
- **Anchors:** For PDFs, give **page # or section**; for web, **section/heading**.
- **Limits & Unknowns:** Add a labeled section with data gaps and verification steps.
- **Output Controls:** Exact sections, bullet caps, and locality (e.g., NC/LATAM, micro-loan context). If info is missing, **state an assumption in [brackets] and proceed**.
- **Model Controls:** Keep temperature low for reliability; enforce **reading level/word caps**.

B. Two Ways to Craft a Strong Deep Research Cold-Start Prompt

You have two reliable options for building a high-quality, BBA-aligned cold-start Deep Research prompt:

1. **Use the Drop-in Starter Prompt and fill in the brackets.**

2. Use a prompt-generator AI prompt generator and feed it simple instructions.

Both paths use the same BBA guardrails (evidence-first, recent sources, localization, and a silent internal accuracy check).

1) Drop-in Starter Prompt (fill in the brackets)

Paste this into your Deep Research tool and customize the bracketed fields:

ROLE & CONTEXT: You are an evidence-first analyst for BlessBack Business Academy serving small businesses/ministries in [exact country/region].

MISSION: Produce a source-backed brief that a field leader can act on.

SCOPE: Focus on [scenario]. Prioritize 2022–present sources; older only if justified.

SOURCE RULES: Use primary/official documents first (laws, government/IGO datasets, official platform docs, bank tariffs). Avoid blogs/aggregators for core claims.

QUOTE GATE: Include ≥ 1 short (≤ 25 words) quote from an official/policy/API page with URL and page/section.

EVIDENCE LOG: Return a table: claim \rightarrow primary URL \rightarrow quoted line(s)/page \rightarrow date \rightarrow confidence.

ANCHORED CITATIONS: For PDFs, give page #; for web, give section/heading.

LOCALIZATION: Tune to [NC/LATAM/micro-loan or ministry] context (fees, regulations, seasonality).

OUTPUT: (1) 1-paragraph executive summary (≤ 120 words). (2) Bullet brief (8–12 bullets). (3) “Limits & Unknowns.” (4) Evidence Log table.

IF INFO IS MISSING: State the assumption in [brackets] and proceed.

TONE/STYLE: Plain language, numbered headings, no buzzwords.

INTERNAL ACCURACY CHECK (SILENT): Before finalizing, compare all numbers, dates, names, and key facts against your cited sources and across sections. Quietly correct any inconsistencies; only mention remaining conflicts briefly in “Limits & Unknowns,” without describing your checking process.

2) Using a Prompt-Generator AI (simple instructions in, full prompt out)

An easier way to create an excellent Deep Research prompt is to give **simple instructions to a prompt-generator AI** and let it build the full, detailed prompt for you.

In ChatGPT 5 Thinking mode, you can use this Custom GPT Prompt Generator:
<https://chatgpt.com/g/g-IQiALrkKZ-custom-gpt-prompt-generator>

How to use it for Deep Research

1. **Describe your scenario in plain language** (2–4 sentences): country/region, audience (small business or ministry), and the decision you need to make.
2. **Add the BBA guardrails in short form**, for example:

“Use 2022–present sources, evidence-first, avoid blogs for core claims, include at least one short quote from an official source, return an evidence log table, and apply this INTERNAL ACCURACY CHECK (SILENT): Before finalizing, compare all numbers, dates, names, and key facts against your cited sources and across sections. Quietly correct any inconsistencies; only mention remaining conflicts briefly in ‘Limits & Unknowns,’ without describing your checking process.”

3. **Ask the generator:**

“Create one Deep Research prompt I can copy into a new Deep Research run.”

4. **Copy and paste the generated prompt** into a fresh Deep Research session (ChatGPT or Gemini) and run the research.
5. **Check the result against the BBA rubric** (source quality, traceability, date freshness, specificity, accuracy, ethics) and lightly edit the prompt if needed before reusing it for similar scenarios.

This two-path approach lets non-prompt-engineers get **consistent, rubric-aligned Deep Research prompts** with much less typing, while still keeping BBA’s evidence and ethics guardrails firmly in place.

Scenario #1 — WhatsApp vs Telegram (Round-2 with Refined Prompt)

Background & Rationale for the Second Pass

Cold-start runs exposed two recurring defects limiting scores: (1) weak sourcing (blogs/aggregators for core claims) and (2) missed **Quote Gate** (no ≤25-word official quotes with anchors). We re-ran Scenario #1 using a refined prompt that tightened

structure and citations. We added an internal accuracy check to test whether these controls improve first-pass quality and gate compliance.

What changed (controls added by the Round-2 prompt you used)

- **Role & audience:** evidence-first; **low-bandwidth Latin America.**
- **Scope & recency:** prefer **2020–present**; justify any sources older than 2020.
- **Source hygiene:** avoid AI/SEO spam; de-duplicate citations.
- **Structured output:** Exec Summary; Bullet Brief (adoption, bandwidth, features, security, 2022–present LATAM use cases, practical recommendations); **Limits & Unknowns; APA references.**
- **In-text citations** for every concrete claim (with working URLs).
- **Internal accuracy check (silent)** before finalizing.

What was still missing (why the gates still failed)

- **No explicit “Quote Gate” requirement** for a ≤25-word quotation from an official policy/API/FAQ page.
- **No citation anchors** (page # for PDFs; named section/heading for web).
- **No Evidence Log table** (claim → URL → ≤25-word quote + anchor → date → confidence).

What stayed the same (to keep the test fair)

- **Same scenario & framing:** WhatsApp vs Telegram for small-business sales in **low-bandwidth Central America.**
- **Same scoring & gates:** BBA rubric weights unchanged; gates = link integrity ≥80%, **Quote Gate**, no PII/unsafe.
- **No pre-score edits; same platforms** (ChatGPT GPT-5 Thinking; Gemini).

How We Ran the Second Pass (Protocol)

1. Applied the **same refined prompt** to ChatGPT (GPT-5 Thinking) and Gemini.
2. Captured unedited outputs as **Round-2 reports**.
3. Scored with the **BBA rubric** and applied the **three gates**.
4. Recorded **Round-1 → Round-2 deltas** (scores, gate status, qualitative changes).

Round-2 Scores (BBA Rubric: 0–100)

ChatGPT (Round-2)

- **Source Quality (25%): 17/25** — Mix of reputable (Pew, The Verge) plus secondary blogs; needs more primary/platform docs.
- **Traceability (25%): 17/25** — Links present, but **no ≤25-word quotes**, few anchors, **no Evidence Log**.
- **Date Freshness (10%): 9/10** — Mostly 2023–2025.
- **Specificity & Context (15%): 13/15** — Clear LATAM/low-bandwidth focus.
- **Accuracy & Consistency (20%): 17/20** — Generally consistent; some adoption figures from compendiums.
- **Ethics & Limits (5%): 5/5** — Explicit **Limits & Unknowns**.
Total: 78/100 → Use with Edits.
Gates: Links ≥80% **Pass** • Short official quote **Fail** • Safety **Pass**

Top fixes (to clear gates)

1. Add **≤25-word** quote from **WhatsApp FAQ** (E2EE) and/or **Telegram Business** (with section).
2. Convert key claims into an **Evidence Log** (claim → primary URL → quoted line/section → date → confidence).
3. Swap/triangulate blog stats with **DataReportal/Pew/Sensor Tower** or official docs.

Gemini (Round-2)

- **Source Quality (25%): 12/25** — Heavy blog/SEO reliance; some official links present.
- **Traceability (25%): 13/25** — Many links, but **no ≤25-word quotes**, no anchors, **no Evidence Log**.
- **Date Freshness (10%): 9/10** — Largely 2024–2025.
- **Specificity & Context (15%): 12/15** — Good LATAM/low-bandwidth framing; still broad.
- **Accuracy & Consistency (20%): 15/20** — Some assertions rest on secondary sources.
- **Ethics & Limits (5%): 4/5** — Limits noted; assumptions vs evidence could be clearer.

Total: 65/100 → Use with Edits.
Gates: Links ≥80% **Pass** • Short official quote **Fail** • Safety **Pass**

Top fixes (to clear gates)

1. Insert **≤25-word** quotes from **WhatsApp FAQ** (E2EE/backups) and **Telegram Business** (with anchors).
2. Replace adoption/usage blogs with **Pew/DATAReport/Sensor Tower** (cite exact figures + anchors).
3. Add an **Evidence Log**; demote/remove Reddit/marketing blogs.

Cold-Start vs Refined Prompt (Comparison)

Platform	Round-1	Round-2	Δ	Gates (R1→R2)	What improved	Still missing
----------	---------	---------	---	------------------	---------------	---------------

ChatGPT	69	78	+9	Quote: Fail→Fail · Links: Pass→Pass · Safety: Pass→Pass	Clearer LATAM/low-band width focus; stronger reputable	Short official quote missing; no Evidence Log ; some stats from blogs; few anchors
---------	----	----	----	--	--	---

sources; better
structure & Limits

Gemini	61	65	+4	Quote: Fail→Fail · Links: Pass→Pass · Safety: Pass→Pass	More complete feature/bandwidt h/security coverage; practical funnel guidance	Source quality weak; no Evidence Log; adoption/security claims need primaries & quotes
---------------	----	-----------	-----------	--	--	--

Round-1 baselines from the earlier cold-start evaluation.

Verdict (Round-2)

- **ChatGPT:** Improved (**78/100**) but not “Ready.” Add a **short official quote** and an **Evidence Log** to reach $\geq 80/100$ and likely pass all gates.
- **Gemini:** Modest lift (**65/100**). Needs primary sources, direct quotes, and an Evidence Log to move into the 70s+.

Prompt Patch (to clear gates next time)

Append these three lines to your **Evidence & Citation Rules**:

- **QUOTE GATE:** Include **≥1 short (≤ 25 words)** quote from an **official** policy/API/FAQ page (WhatsApp/Meta, Telegram) with a working URL.
- **ANCHORS:** For PDFs, include the **page number**; for web pages, include the **named section/heading** for every key claim.
- **EVIDENCE LOG:** Append a table: *claim* → *primary URL* → ≤ 25 -word quote + *page/section* → *date observed* → *confidence*.

Bottom line: The refined prompt improved structure, sourcing discipline, and scores (ChatGPT +9; Gemini +4). To reach “Ready to Publish,” enforce the **Quote Gate**, require **anchors**, and include an **Evidence Log**.

Recommendation: Which Deep Research Engine to Use

Short answer: Use ChatGPT (GPT-5 Thinking) as the default Deep Research engine; run Gemini as a secondary breadth + cross-check pass when stakes are high.

Why ChatGPT as the default

- **Higher, more consistent scores:** Scenario #1 (both rounds) and Scenario #2 show ChatGPT > Gemini (e.g., 78 vs 65 in Round 2; 86 vs 59 on the uniform plan). It also benefited more from prompt refinements (+9 vs +4).
- **Better structure & localization:** Outputs align with BBA sections, “Limits & Unknowns,” and Central America framing.
- **Faster to “Ready”:** Two quick edits typically clear gates: add one ≤25-word official quote and include an Evidence Log.

When to bring in Gemini

- **Breadth scans/source canvassing:** In Scenario #3, Gemini surfaced broader IGO/NGO sources (76 vs 67).
- **Google-centric or discovery tasks:** Good for a second sweep over gov/official docs found via Search.
- **Second opinion:** Use to triangulate numbers and phrasing before finalizing.

Bottom line: Make ChatGPT your default Deep Research tool. Use Gemini to widen coverage and cross-check.

Chapter 4: Custom GPTs & GEMs

1) What is a Custom GPT?

A **Custom GPT** is a tailored version of ChatGPT with your own **Purpose & rules** (instructions/persona, tone, language behavior), **Knowledge** (files/links you upload as

reference), **Capabilities** (web browsing/search, data analysis, image generation, actions/integrations), and **Sharing controls** (private, team-only, org-wide, or public).

Use a custom GPT when you want a **repeatable assistant** that behaves consistently for your team or audience (e.g., “BlessBack Entrepreneur Guide,” Curriculum Generator, Lesson Guide Creator, etc.).

2) Custom GPT vs. Other ChatGPT Features

Use Custom GPT when you need a persistent, shareable assistant with fixed guardrails, knowledge, and actions.

Use Study & Learn when the goal is **personal learning**: tutoring, step-by-step exercises, quizzes, memory cards, etc. (It’s an interactive coach, not a shareable assistant.)

Use Projects when you want a **workspace** to group chats, files, and instructions for an initiative (e.g., “Pueblos Blancos Study”), then run **Deep Research** or regular chat within that project.

Use Deep Research when the main task is a **web-backed investigation with a** structured, cited synthesis. You can still *hand the final writing* to a custom GPT you use for BBA tone.

Rule of Thumb

Repeatable role → **Custom GPT**.

Learning journey → **Study & Learn**.

Organized context → **Projects**.

Evidence from the web → **Deep Research** (then polish with your GPT).

3) How to Create a Custom GPT (Step-by-Step)

- **Define the mission**

One sentence: Who is it for, and what will it produce? (e.g., “Create bilingual lesson outlines and handouts aligned to BBA curriculum, with citations.”)

- **Draft crystal-clear instructions:**

- **Role/voice** (pastoral, evidence-based).
- **Scope** (what to do) & **boundaries** (what not to do).
- **Output formats** (1-page brief, lesson plan, table, rubric, prayer).
- **Safety & ethics** (no personal information in consumer chats; require citations; flag uncertainty).
- **Interaction style** (ask clarifying questions *only when essential*; otherwise, proceed).

Shortcut: see Section 9, Custom Instructions Generator.

- **Add knowledge**

Upload/select: deep research reports, manuals, course PDFs, policy docs, templates, logos, brand guide.

Tip: curate **one core folder** per GPT; remove duplicates/conflicts; use descriptive filenames.

- **Enable the right capabilities**

- **Search/Browsing** for live facts and citations.
- **Code/Data Analysis** for tables, CSVs, and charts.
- **Image Generation** for blog/social graphics (prefer real-people photos for final publishing).
- **Actions** (optional): connect approved APIs (Google Drive fetcher, Notion, Quickbase, Mailer, Telegram/Coze, etc.).

- **Seed conversation starters**

Provide 3–6 “one-click” prompts that show ideal use (e.g., “Draft a 1-hour workshop on WhatsApp for micro-businesses—include slides outline + demo prompts.”)

- **Test with gold prompts.**

Prepare 5–10 realistic tasks with **expected answers**, including testing bilingual output, citations, tone, and formatting.

- **Harden & govern**

- Turn on **workspace-only sharing** for internal tools.
- Disable training on user content if available; log updates (vX.Y + date).
- Add a **disclaimer** (educational, not legal/medical; cite sources; keep PRIVATE INFORMATION out).

- **Deploy & train users.**

- Short Loom video or 1-pager: *what this GPT is for, sample prompts, do's & don'ts.*
- Add the GPT to your manual and link it from your syllabus/portal.

- **Monitor & improve**

- Collect example failures; update instructions/knowledge monthly.
- Maintain a **change log** and retire stale files.

4) Keys to a Powerful Custom GPT

- **Instruction hygiene:** Use Do/Don't lists; give **format scaffolds** (templates); set **decision criteria** for recommendations.

- **Knowledge discipline:** Keep a small, **up-to-date core of materials**; prefer primary docs; add a “recency” note (e.g., “Use data 2022–present”).
- **Evidence by default:** Require citations for non-obvious claims; add an **Evidence Log** table to outputs.
- **Multilingual option:** Default to English, however instructions can provide responses in other languages (e.g., Spanish, Haitian Creole).
- **User ergonomics:** Conversation starters, quick-reply buttons, and short commands (e.g., /brief, /slides).
- **Safety:** No personal information; neutral tone on sensitive topics; always surface limitations.
- **Benchmarking:** Keep a **scorecard** (accuracy, structure, style, usefulness) across your gold prompts.

5) Limitations to Expect

- **Not a database or LMS**—it retrieves from uploaded knowledge but can miss specifics or prioritize wrong versions.
- **Staleness risk:**
 - **ChatGPT custom GPT uploads:** Static. If you upload a file to the Knowledge tab, it does *not* auto-sync when the original changes.
 - **ChatGPT Actions / synced connectors → Google Drive:** Live. The GPT reads the current content from Drive each time, or from an auto-updated Drive index for synced connectors.
 - **Gemini Gems + Google Drive:** Live for Google Docs/Sheets (changes in Drive are reflected automatically); static for most other file types (e.g., PDFs) unless you re-upload.
- **Conflicting sources**—it needs clear precedence rules; conflicting PDFs yield hedged answers.
- **Citation gaps**—requires explicit instruction to cite; verify links.
- **File/size limits & rate limits**—plan for concise, high-signal references.

- **Privacy**—consumer chats aren't for donor/health PRIVATE INFORMATION; use enterprise/suite tools or local models.

6) Sharing & Distribution

- **Private (you only)** during drafting.
- **Team/Org-only** for internal tools (recommended default).
- **Link-sharing** for pilots/training cohorts.
- **Public/GPT Store** only for generic education tools with no proprietary content.
- Add a short **Use Policy**: acceptable use, citation rule, no PRIVATE INFORMATION, contact for issues.

7) Do Other Platforms Offer “Custom GPT-like” Features?

- **Gemini “Gems” (Google)**: Create custom versions with instructions; good if your content/workflow lives in **Gmail/Docs/Sheets**.
- **Microsoft Copilot Studio**: Low-code **agents** with connectors and Power Automate; best if you're on **Microsoft 365/SharePoint** and need enterprise governance.
- **Claude (Anthropic) Projects + Tool Use**: Strong long-form writing with APIs/tools; excellent for policies, curricula, careful tone.
- **Perplexity (Projects/Labs/Deep Research)**: Research-first; excels at **cited web synthesis**; less about persona bots, more about **evidence-backed reports**.

8) Best-Practice Workflow (from idea → reliable tool)

- **Pick one role** per GPT (avoid “kitchen sink”).
- **Assemble the core materials** (5–30 high-quality, non-duplicative sources).
- **Write instructions** with: mission • audience • do/don't • outputs • citation rule • bilingual rule • safety.
- **Enable capabilities** (Search, Code, Images, Actions) only as needed—minimally.
- **Seed starters** (6 prompts).
- **Test on gold prompts** with a **scorecard** (4×5 rubric).
- **Fix failure modes** (add rules/knowledge; tighten outputs).

- **Govern & share** (team-only; changelog; owners).
- **Train users** (1-pager + 5-min video).
- **Review monthly** (replace stale docs; collect examples; update version).

9) Using the Custom Instruction Generator (CIG)

What it is: A helper GPT that turns a short design brief into polished, paste-ready instructions for a Custom GPT or for your personal Custom Instructions.

When to use: - You know the role (e.g., *BBA Content Coach*) and want crisp rules, Do/Don't lists, and output formats. - You're standardizing tone, bilingual behavior (EN/ES, HT on request), and citation rules across multiple assistants.

How to use (5 steps): 1) **Start a new chat** with CIG. Paste a 6-line brief: *Mission, Audience, Outputs, Do, Don't, Tools*. 2) **Ask for structure**: Role & Scope • Do/Don't • Output Formats • Citation Rule • Bilingual Rule • Safety/Refusals • Few-shot examples • Trigger → Instruction pairs. 3) **Iterate**: “Shorten to 400–600 words,” “Add two few-shot examples,” “Return a paste-ready Markdown block.” 4) **Copy to your GPT**: *Create a GPT* → *Configure* → *Instructions*; then add Knowledge (core materials), Capabilities (Browse/Code/Images/Actions), and Starters. 5) **Test on gold prompts** (5–10 tasks); fix failure modes; version your changes.

What to request from CIG: - Do/Don't rules with examples - Citation & Evidence-log requirement - Triggers (e.g., /brief, /slides, /policy, /bilingual, /devotional) mapped to explicit output templates - Testing rubric and acceptance criteria

Pitfalls & Guardrails: **Keep instructions concise and modular; avoid sharing personal information; prefer team-only sharing for internal GPTs; re-run CIG when your core or outputs change.**

10. Field Testing Custom GPTs Across Platforms – What We Learned

This chapter has focused on designing and *governing* custom GPTs and Gems. Before you deploy one into real ministry or business contexts, you also need a **field-testing pattern** that checks doctrine, safety, and reliability across the different platforms where that assistant will live (ChatGPT, Gemini, Telegram bots, NotebookLM, etc.).

To build that pattern, we ran a full **cross-platform field test** using an Assemblies of God (AG) “Pastor Companion” assistant. This section summarizes the method, results, and practical recommendations so you can adapt the same approach for your own custom GPTs.

10.1 What We Tested

We treated the AG Pastor Companion as a **case study of a high-risk, high-impact assistant** (pastoral doctrine + crisis care).

Platforms (5 delivery channels)

Each one used the same AG knowledge base and doctrinal expectations:

1. Telegram “AG Pastor Companion” (ChatGPT 4o mini backend)
2. Gemini 2.5 Pro (standalone)
3. NotebookLM (Google)
4. ChatGPT-5 Thinking (standalone)
5. Telegram Chatbot (Gemini 2.5 Pro backend)

Prompts (5 real pastoral scenarios)

Each platform answered the same five questions, chosen because they mix **doctrine** and **safety-critical** issues:

1. Baptism in the Holy Spirit (BHS) – doctrine & AG sources
2. Domestic Violence (DV) safety protocol – crisis response
3. Tongues & interpretation in a service – governance & order
4. Divorce and ministerial credentials – policy & exceptions
5. The “Blessed Hope” (rapture) – eschatology and contrasts

Rubric (0–5 per criterion, max 25 per question)

For each answer, we scored:

- **(a) Source precision:** AG title + year required, not just “AG guidelines.”
- **(b) Scripture accuracy:** correct NIV references and in-context usage
- **(c) Key doctrinal distinctions:** e.g., initial physical evidence vs sanctification
- **(d) Doctrinal / policy accuracy:** matches official AG documents

- (e) **Detail vs generality:** specific checklists and steps, not vague advice

We also kept a “**Scripture & Source Flags**” list to capture any hallucinated dates, missing titles, or context-stretched verses.

Finally, to reduce bias, three separate audits were run (ChatGPT-5 Thinking, Gemini 2.5 Pro, and Grok), each independently scoring the same transcripts.

10.2 High-Level Results

Despite different auditors, the **rankings converged strongly** across the three evaluations:

1. **ChatGPT (Standalone)** – clear #1
2. **Gemini (Standalone)** – close #2
3. **Telegram (Gemini backend)** – #3, solid but thinner
4. **NotebookLM** – #4, high doctrine but unsafe DV in one pass
5. **Telegram (ChatGPT backend)** – #5, lowest detail and the highest hallucination risk

Key patterns:

- **Standalone ChatGPT-5 Thinking** delivered the most consistent doctrinal accuracy, precise AG citations (including years), and the strongest DV and governance protocols. It frequently used **numbered checklists** and clear “must NOT” lists (e.g., no couples counseling in abuse cases, do not promise unconditional confidentiality, do not delay reporting).
- **Standalone Gemini 2.5 Pro** also performed at an “elite” level, slightly behind ChatGPT, mainly on **minor citation details** (e.g., occasionally missing a publication year for an AG document). Its DV protocol and doctrinal summaries were robust and safe.
- **NotebookLM** showed excellent doctrinal summaries and strong use of official AG papers in most questions, but one DV answer was **dangerously generic**, lacking explicit 911 / mandatory reporting / “must not” prohibitions. That single weakness significantly lowered its suitability for crisis-care use without hardening.
- **Telegram (Gemini)** generally provided accurate doctrine and even scored perfectly on DV in one audit. Still, its answers tended to be **briefer**, with fewer

legal/reporting specifics and less precise AG citation lines, compared with the standalone tools.

- **Telegram (ChatGPT)** produced the **weakest** pattern: more generic responses, under-specified DV safety steps, and repeated **fabrication of AG source years**, such as “Blessed Hope (2023)” or “AG spiritual gifts guidelines (2023),” when the actual Statement of Fundamental Truths dates back to 1916.

Across all three audits, the **same broad recommendation** emerged:

“Adopt ChatGPT (standalone) as the primary pastor companion, with Gemini (standalone) as a close second; pilot Telegram bots for reach, but only with strong guardrails; treat NotebookLM and Tele-ChatGPT as ‘not ready’ for DV/credentialing use without significant hardening.”

10.3 What This Means for Custom GPTs & Gems

From a Playbook standpoint, this case study provides **portable principles** you can apply to *any* custom GPT—whether it serves pastors, entrepreneurs, or program staff.

1. Test the “brain” AND the “wrapper.”

The same model performed differently when wrapped in Telegram vs. its native interface. Latency limits, token limits, or connector behavior can quietly push answers toward **generality** rather than detailed protocols. Always test:

- Standalone model (ChatGPT / Gemini)
- Your front-end (Telegram, WhatsApp, web chat, etc.)

2. Use real, high-risk scenarios.

We did not test on trivia. We tested on **pastoral scenarios that could hurt people if mishandled** (DV, divorce & credentials). Your field test should use equally serious “edge” cases:

- For entrepreneurs: loans in default, employee termination, harassment complaints, high-risk marketing claims
- For ministries: abuse allegations, suicide risk, financial impropriety, doctrinal confusion

3. Score against your knowledge database, not vibes.

The rubric forced each answer to match **named documents**, correct verses, and specific steps—not just “sounds right” summaries. That’s why hallucinated years and vague references were caught quickly.

4. **Evidence Log is non-negotiable.**

Every platform was expected to produce checkable citations, and the audits maintained a **Scripture & Source Flags List** and overall Evidence Log. This aligns directly with the BBA AI Playbook's requirement that *all non-obvious claims be traceable to a primary or official source*.

5. **Consumer surfaces are for drafts, not final counsel.**

Given the DV and hallucination issues, the report reinforces the Playbook rule: use **Workspace/M365-style, tenant-governed suites** for sensitive pastoral or client work; keep Telegram/consumer tools for lower-risk drafts or outreach.

10.4 Deployment Recommendations

Based on the field test, here's the pattern we recommend adding to your own custom GPT/Gem rollout plan:

1. **Choose a primary + backup “brain.”**

- **Primary:** ChatGPT-5 Thinking (standalone) for doctrinal / policy-heavy, safety-sensitive use.
- **Backup:** Gemini 2.5 Pro (standalone) as a second opinion and redundancy.

2. **Harden your messaging front-ends.**

For any Telegram / WhatsApp / web bot built on top of your custom GPT:

- **Prepend policy quote blocks** from your core materials for each high-risk topic (e.g., AG 2010 BHS paper, 2024 Divorce paper, bylaws article numbers, or your academy's HR policies and loan terms).
- Run a **Scripture or Policy Validator** pass (NIV or your legal/policy text) that checks references and flags context stretches.
- Implement a “**red-flag macro**” for crisis domains (DV, suicide, self-harm, severe debt), forcing:
 - 911 / emergency language
 - Mandatory reporting or legal-duty steps, where applicable
 - “Must NOT” items (no couples counseling in abuse, no legal advice, no promises of secrecy, etc.).

3. Enforce citation style at the prompt level.

In your system instructions for every custom assistant, include:

- “Always give **title + year** for doctrinal/policy citations.”
- “Maintain an **Evidence Log**: claim → source URL / doc + page → date.”
- “Flag any uncertainty about dates or titles rather than guessing.”

4. Use multi-auditor reviews for launch and re-certification.

Repeat what we did here:

- Have at least **two different AI systems** (e.g., ChatGPT and Gemini) audit your assistant’s answers with your rubric.
- Consolidate convergences and discrepancies into a short QA memo before you approve a new version.
- Re-run this test when you:
 - Change your canon or policies.
 - Switch model versions
 - Add a new high-risk use case.

5. Document and train your humans.

The tool is not the pastor or entrepreneur coach. Your field-test report becomes **training content** for staff: it shows where the assistant is strong, where it is weak, and how to read its answers with discernment (including the “weigh/judge” step from 1 Corinthians 14 in pastoral contexts).

Chapter 5: Projects in ChatGPT

1) What are Projects?

Projects are innovative workspaces that keep everything related to a long-running effort in one place—**chats, files, instructions, and tools**—so ChatGPT can stay on-topic and reuse context across sessions. Projects also support **Deep Research** (multi-step, web-backed investigations), voice mode for discussing ideas, and **sharing individual chats** when you need to show your work. Newer updates include **project-only**

memory, which allows ChatGPT to reference other conversations within the same project without drawing from your global saved memories.

Think of a Project as: - A **folder + brain** for a topic. A place where **files + instructions** live together so the model can cite, compare, and build deliverables consistently.

Great uses: Recurring research topics, donor/partner briefings, multi-language course development, policy repositories, and grant pipelines.

2) When to Use Projects (and when not)

Use a Project when: - You'll revisit the topic **multiple times** over weeks/months. - You have **source files** (PDFs, Docs, slides, URLs) that the model should reuse. - Multiple people need a **shared place** to read and contribute to chats/files. - You want to run **Deep Research** that blends your files with the public web.

Don't use a Project when:

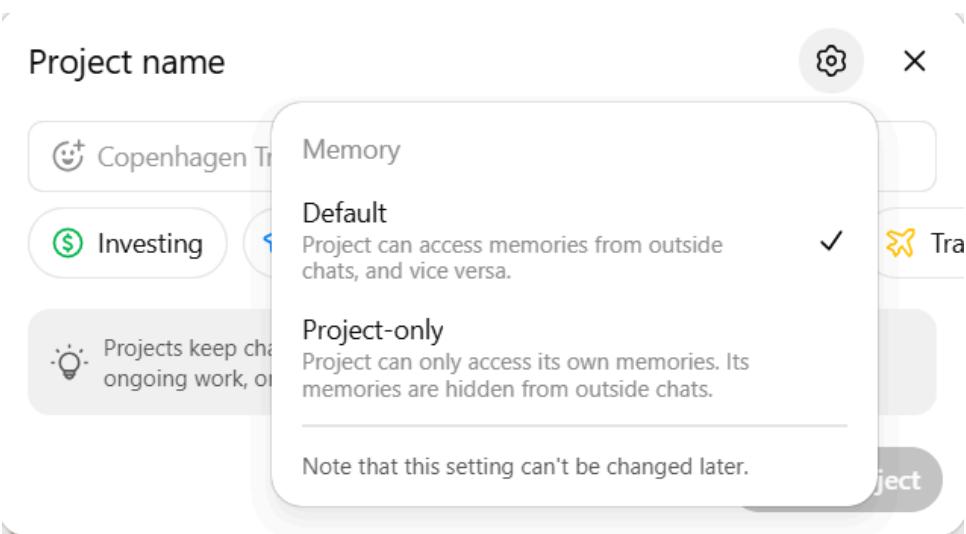
- It's a **one-off** question or quick email rewrite → use a normal chat.
- You need a **repeatable bot** others can use anywhere → build a **Custom GPT**.
- You're **learning** or want quizzes/flashcards → use **Study & Learn**.
- You're drafting with **sensitive PRIVATE INFORMATION** and don't have enterprise controls → keep work in Workspace/Microsoft365 or a **local model** and paste only non-sensitive excerpts.

3) Set-up (Step-by-Step)

- **Create the Project**

Name it clearly (e.g., 2025_PueblosBlancos_Study). Add a one-sentence **objective** at the top of the description.

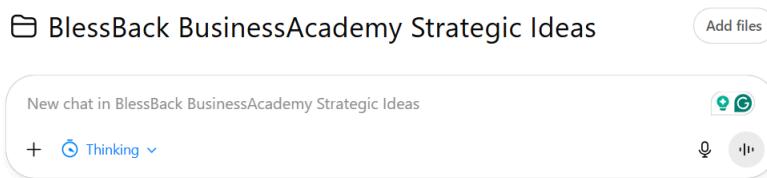
- **Click on the Gear Box** to set the memory. This is an essential step because you can't change this setting later. If you want the chat memory in the project to include only chats within that project, select "Project." If you choose "Default," then chat memory from outside the project will be included. It's recommended to use "Project" only.



- **Add Instructions**

Paste short, explicit guidance: audience, tone, citation rule (“cite every non-obvious claim”), deliverable formats, and Do/Don’t lists.

Step 1: Click on the three dots shown in the circle.



Step 2: Click on Add instructions.

1

Edit project
Add instructions
Delete project

✉ BlessBack BusinessAcademy Strategic Ideas

Add files

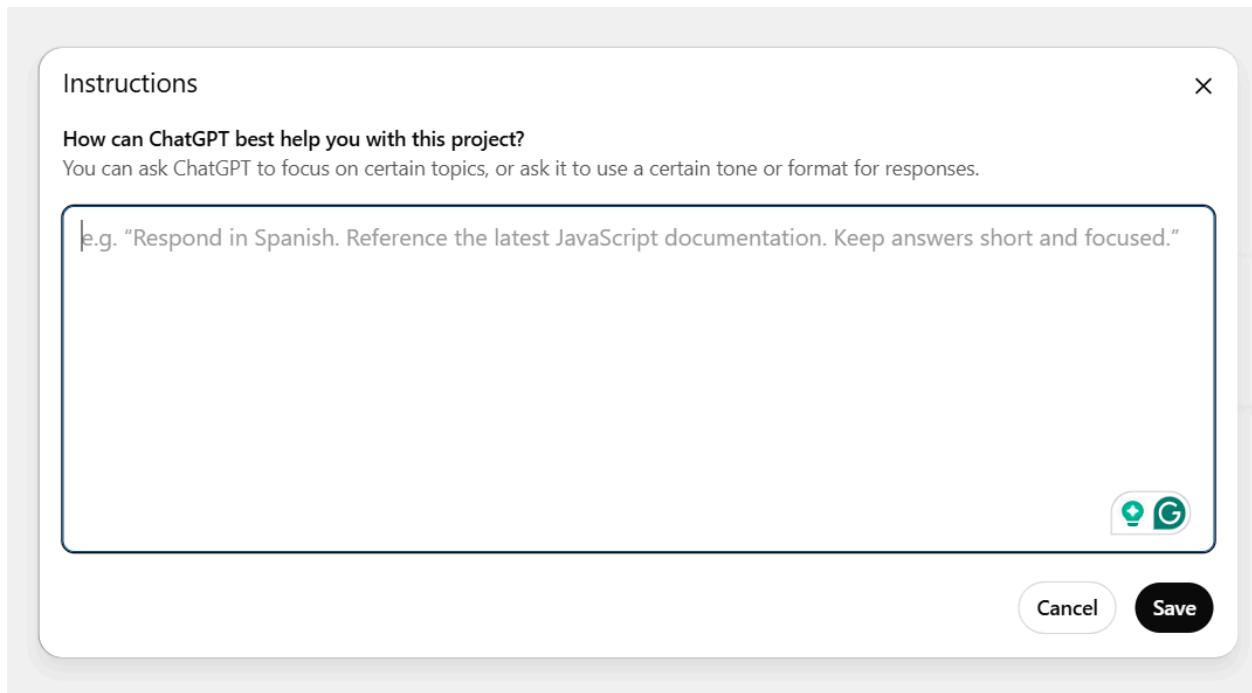
New chat in BlessBack BusinessAcademy Strategic Ideas

+  Thinking



0 11

Step 3: Add instructions that will help ChatGPT respond to chats in your project. You can use the [custom instruction generator](#) to create instructions that improve the Project's performance.



- **Upload/Link Sources**

Add the **core materials**: latest PDFs, policy docs, templates, data files, and key URLs. Prefer 5–30 **high-signal** sources; remove duplicates and older versions.

- **Enable Tools**
Turn on **Deep Research** for web evidence; enable **Code/Data** for tables/CSV work; **Images** if you need graphics; and voice if you'll brainstorm out loud.
- **Organize Chats**
Use separate chats for phases: /plan, /desk_research, /analysis, /slides, /español, /kreyòl.
- **Use Project-Only Memory**
If enabled, ChatGPT can reference other chats **inside** the project while keeping that memory contained to this project.
- **Share as needed**
Keep internal work **team-only**; share individual chats externally (read-only links) when you must show reasoning or evidence.

4) Sharing & Controls

1. **Who can see whom?** Everyone you invite can see the member list in the project's sharing pane (the owner has complete visibility and can add or remove members). [OpenAI Help Center](#)
2. **What can invitees see?**
 - a. If you give **Chat access**, members can see and use the project's **chats, files, and instructions**, but cannot invite others.
 - b. If you give **Edit access**, members can **update instructions, upload/remove files, and invite others** (but they still can't remove existing members). [OpenAI Help Center](#)
3. **Can you make the files owner-managed only?** Yes—set everyone else to **Chat** (not Edit). Only the **owner** can always add/remove members and has complete control. [OpenAI Help Center](#)
4. **How many files & collaborators?** Current limits per shared project:
 - a. **Pro:** up to **40 files** and **100 collaborators**
 - b. **Plus/Go:** up to **25 files** and **10 collaborators**
 - c. **Free:** up to **5 files** and **five collaborators** [OpenAI Help Center](#)
5. **Per-file upload limits (size):**
 - a. **Any file:** up to **512 MB**
 - b. **Docs/text:** capped at **~2M tokens** per file
 - c. **CSV/spreadsheets:** **~≤50 MB** (row-dependent)

- d. **Images:** ≤20 MB each. [OpenAI Help Center+1](#)
- 6. **Link sharing note:** Owners can set visibility to **Only those invited** or **Anyone with a link** (logged-in users). You can switch back to invite-only at any time. [OpenAI](#)

5) Best Practices

Core material hygiene: keep a small, up-to-date source pack; add a README with descriptions of each file and the core material version. - Name files clearly (e.g., BBA_BrandGuide_v2025.pdf), and delete superseded copies.

Prompt scaffolds: seed each project with a brief and slides template, plus an Evidence Log table that the model can reuse.

Evidence & citations - Require citations for all non-obvious facts. Ask for **competing views** and **confidence levels**—archive key sources (as PDFs or screenshots) alongside the outputs.

Multilingual outputs: default to English; request **other languages as needed**.

Governance & Privacy: Place private information, donor, or health-related content in Workspace/Microsoft 365; paste only what's necessary into ChatGPT. Use **team-only sharing** by default.

Versioning: Add a brief **changelog** message at the top of the project. Stamp deliverables with vX.Y + date.

6) Running Deep Research *inside* a Project (quick recipe)

- Start a new chat: “List sub-questions, planned sources, and how you’ll validate claims.”
- Run: “Conduct Deep Research using our files + the public web. Create a 1-page executive brief, plus an appendix with quotes/links. Mark uncertainties.”
- Ask for tables (metrics, comparisons, pros/cons) and **bilingual summaries**.
- Export the brief; archive the **Evidence Log**.

7) Maintenance & Handover

- **Monthly audit:** replace stale sources, prune duplicates, and update instructions if outputs drift.
- **Handover kit:** objective, instructions, file list, recent outputs, and open questions.

8) Comparison: Projects vs. Similar Features Elsewhere

- **Perplexity — Labs/Projects:** research-first workspaces that can run deep web browsing, code, and charting to produce finished artifacts (reports, dashboards). Great for **cited web synthesis**; less persona/knowledge management than ChatGPT Projects.
- **Google — NotebookLM (and Gemini add-ons):** notebooks built from your sources with mind maps and audio/video overviews; powerful for **learning from your own materials** and turning them into explainers. Gemini's broader app can assist, but NotebookLM excels as a **source-anchored notebook**.
- **Microsoft — Copilot Studio Agents:** governed **agents** tied to M365/SharePoint with connectors and flows; best for **process automation** and org-safe retrieval rather than open web research.

Rule of thumb:

- Need a **general research + drafting** workspace with reusable context? → **ChatGPT Projects**.
- Need **source-anchored notebooks** and multimedia explainers? → **NotebookLM**.
- Need **governed automations** over tenant data? → **Copilot Studio Agents**.
- Need **fast web-cited answers**? Pair **Perplexity** with your Project for cross-checks.

9) Troubleshooting

- **Messy answers or drift** → tighten project instructions; reduce source set; ask for a step-by-step plan first.
- **Missing citations** → insist on quotes + working links; reject uncited claims.
- **Conflicting sources** → request a *disagreement matrix* and the model's confidence with reasons.
- **Performance lag** → split into sub-projects (e.g., *Market*, *Policy*, *Slides*).
- **Sensitive data** → keep in Workspace/Microsoft365 or draft locally; paste only redacted snippets.

Chapter 6: Study & Learn — Guided Learning

1) What is Study & Learn?

Study & Learn (“study mode”) transforms ChatGPT into a guided tutor that helps you work through problems step-by-step—asking questions, checking your reasoning, providing hints, generating practice, and tracking your progress. It’s designed to **build understanding**, not just hand you an answer.

Suitable for: mastery of concepts, spaced practice, quizzes, exam prep, vocabulary/drills, “teach-back,” and ethics-safe learning (avoids dropping complete solutions).

Available to: Free, Plus, Pro, Team (Edu rollouts vary). Access it from **Tools → Study & learn**.

2) When to Use (and when not)

Use Study & Learn when: - You want **tutoring dynamics** (Socratic questions, hints, formative feedback). - You need **practice**: quizzes, flashcards, retrieval exercises, and explanations. - You’re onboarding staff or volunteers and want **guided drills** (e.g., BBA brand voice, donor etiquette, safeguarding). - You’re preparing **students/entrepreneurs** for a lesson, assessment, or certification.

Don’t use it when: - You need a **polished deliverable** (policy, brochure, grant paragraph) → use **Regular Chat** or your **Custom GPT**. - You’re doing **multi-source research** with citations → use **Deep Research** (then convert into a study pack). - You’re organizing **files + instructions** for a long project → use **Projects**.

3) Quick Start

- Open a new chat → Tools → Study & learn.
- Paste your **learning goal** (who, what, due-when).
- Ask for a **study plan**, practice type (quizzes, flashcards, case questions), and **checkpoints**.

Starter prompt (paste): > I’m training BBA entrepreneurs on pricing. Create a 60-minute **study plan** with three checkpoints and 12 practice questions (mix of multiple choice, short answer, and one mini-case). Ask me questions as we go to check understanding.

4) Core Study Patterns

A) Explain → Test → Teach-back

Give a **short explanation**.

Run a **3–5 item quiz**.

Ask the learner to **teach it back** in 4–6 sentences.

B) Worked Example → Varied Practice

Show one annotated example.

Give **10 mixed problems** with increasing difficulty, with feedback.

C) Case & Reflection

1–2 paragraphs case; ask: **What would you do?**

Feedback vs. model answer; end with a **reflection** and **subsequent action**.

D) Spaced Retrieval Pack

Generate **flashcards** across three difficulty tiers.

Schedule **reminders** (e.g., daily for 5 days).

5) Ready Use Cases & Prompts

Ministries / Pastors

Devotional prep: “Guide me to craft a 10-minute devotional on Proverbs 11:1 (honest scales). Ask Socratic questions, then quiz me.”

Safeguarding policy drill: “Train me on our safeguarding reporting steps. Quiz me with 15 scenario questions and corrective feedback.”

Entrepreneurs / Small Business

Pricing & margins: “Teach me cost-plus and value-based pricing with two worked examples and a 12-question practice set.”

WhatsApp customer service: “Role-play 8 customer chats; score my responses on tone, clarity, and policy compliance.”

Academy / Coaches

Slide design basics: “Coach me on clear slide structure (title → 3 bullets → note). Do six drills from bad → better.”

Accounting basics: “Explain cash vs. accrual with two mini-cases and a 10-item quiz (with answers hidden until I attempt).”

6) Build a Study Pack (template)

Prompt: > Build a **study pack** on. Include: (1) 1-page explainer; (2) 10 flashcards (basic); (3) 10 flashcards (applied); (4) 12 quiz questions with feedback; (5) a mini-case; (6) a teach-back rubric. Ask three check-in questions as I progress. Produce EN → ES; HT summary optional.

Teach-back rubric (0–3 each): accuracy • key terms used • clarity • practical example • following action.

7) Best Practices (Instructional design)

- **Start with outcomes:** what should the learner do by the end?
- **Chunk content:** 5–7 minute segments; micro-quizzes after each.
- **Active recall:** ask the learner to explain in their own words before revealing answers.
- **Vary practice:** mix recognition (MCQ) and production (short answer, case).
- **Bilingual by design:** draft in EN → generate ES handout; HT on demand.
- **Close with action:** end each session with a *checklist of the next step*.

8) Guardrails (Ethics & Privacy)

- **No ghostwriting assignments.** Use explain/teach-back; avoid “write my essay.”
- **Cite sources** if external facts are used; prefer primary/official resources.
- **Protect PRIVATE INFORMATION:** do not paste donor lists or medical details.
- **Accessibility:** offer text + audio read-outs; keep font sizes & contrast readable when exporting.

9) Turn Study into Deliverables

- After mastery, switch to a normal chat (or your **Custom GPT**) to produce:
an **executive brief** (1 page),
- **slides** with speaker notes,
- **handouts**
- a **quiz** with an answer key for class.

10) Troubleshooting

- **Too easy / too hard:** say “increase/decrease difficulty by one level.”

- **I need citations:** “Provide sources and quotes for each key claim.”
- **Not enough practice:** “Add 10 more applied questions and hide answers until I attempt.”
- **Timeboxed:** “Compress to a 20-minute session with two checkpoints.”

11) Mini Glossary

- **Socratic questioning:** the tutor asks guiding questions to surface your reasoning.
- **Teach-back:** learner explains the concept to confirm understanding.
- **Spaced retrieval:** practice spread out over time to improve memory.

Chapter 7: Agents

1) What Agent Mode Is

Agent Mode transforms ChatGPT into a goal-directed assistant that can plan steps, utilize tools, read/write from authorized data sources, and work through tasks with human-in-the-loop oversight. Unlike a one-off chat, an agent keeps **state** about the task, can call **connected actions** (e.g., fetch a doc, summarize a file, generate images), and can follow **policies/guardrails** you define.

Key differences vs. normal chat - **Goal** → **Plan** → **Act** → **Check** loop instead of single responses. - **Stateful context** across steps within the task. - **Tool/Action use** (e.g., browse, read files, write structured outputs, call approved connectors). - **Escalation:** asks for human confirmation at critical steps or when confidence is low.

2) Core Features & Functions

- **Multi-step planning:** decomposes a goal into sub-tasks and tracks progress.
- **Tool use & connectors:** browse/web, project files, image generation, code/data analysis; optional third-party actions (e.g., Google Drive/Sheets, email draft to Gmail, Telegram via bridge services) when configured.
- **Policies & guardrails:** refusal rules, handling of PRIVATE INFORMATION, compliance notes, escalation triggers.
- **Checkpoints & confidence:** asks permission before sending messages externally or touching sensitive files; annotates confidence levels.

- **Reusable “routines”**: saved workflows for recurring jobs (e.g., weekly donor brief → draft → bilingual summary → export PDF).

3) How Agent Mode Differs From Other Platforms’ “Agents.”

- **ChatGPT Agent Mode**: general-purpose, strong reasoning & writing; integrates with **Projects**, **Deep Research**, **Custom GPT knowledge**, and built-in tools. Great for **research→draft→bilingual publishing** loops.
- **Microsoft Copilot Studio agents**: low-code, enterprise connectors & approvals (M365/SharePoint/Teams/Power Automate). Best when your data/processes live in **Microsoft 365**, and you need strict governance.
- **Google (Vertex/Gemini) agents**: firm when building **workflow automations** around Google Workspace or Vertex AI; good doc grounding in **Drive** and Apps Script ecosystems.
- **Claude (Workflows/Tool-use)**: excellent for **long-form reasoning** and carefully constrained tool calls; use when tone/nuance and long context windows matter most.
- **Perplexity**: research-first with **cited web** answers; not positioned as a general process agent, but excellent as a **verification companion** to cross-check sources.
- **Coze/Telegram style chatbots**: fast channel deployment (Telegram, WhatsApp via bridges), suitable for **FAQ & follow-ups**. Pair with ChatGPT Agents for deeper research/drafting.

Rule of thumb: If your workflow is **knowledge-heavy + writing-intensive** with mixed tools (research, docs, images) → start with **ChatGPT Agent Mode**. If your workflow involves tenant data and approvals in Microsoft or Drive-centric automation in Google, consider using those native agent platforms for that slice of work and linking outputs back to ChatGPT for drafting and bilingual publishing.

4) Getting Started — Setup Checklist

A) Define the mission (1 sentence)

Who is it for, what does it produce, and how often?

“Prepare a weekly donor-ready impact brief with citations from our Project files + web.”

B) Draft guardrails

Refusals (no legal/medical advice; no PRIVATE INFORMATION handling in consumer

contexts). - Evidence rules (cite all non-obvious claims; log quotes/links). - Escalation points (before sending external emails, posting to channels, or publishing PDFs).

C) Prepare the core materials

Create a **Project** with the latest PDFs, policies, templates, and brand guide. Remove duplicates; include a README.

D) Map tools & actions

Built-ins: Deep Research, code/data, image generation.

Optional connectors (approved only): Drive/Sheets, email draft, Telegram/Coze bridge, Quickbase/Notion fetchers.

Minimum-scope permissions; read-only where possible.

E) Gold prompts & test cases

5–10 realistic tasks with expected outputs (acceptance criteria & bilingual checks).

F) Share & governance

Team-only to start; changelog (vX.Y + date); owner + reviewer named; archive the **Evidence Log** with outputs.

5) How to Use Agent Mode (First Run)

- **Create/Select Project** with core materials + instructions (EN default; ES, HT on request; citation rule).
- **Start Agent Mode** with your mission and constraints in place.
- Ask the agent to **list its plan** (sub-tasks, tools, data sources, checkpoints).
- Let it run, but require **confirmations** before external actions (send/post).
- On completion, review the **Evidence Log**, edit the draft, and export (PDF/slides).
- Save the workflow as a **routine** if it's recurring.

6) What to Prepare Before You Go Live

- **Instructions:** mission • audience • outputs • Do/Don't • bilingual rule • citation rule • escalation triggers.
- **Templates:** brief, slides, email, social caption, policy page, Evidence Log table.

- **Core Materials:** 5–30 up-to-date files; brand assets; a “source precedence” note.
- **Connectors:** least-privilege access; test accounts; audit logs; data retention plan.
- **Quality bar:** rubric (accuracy, structure, style, usefulness) with thresholds; human reviewer.
- **Privacy:** keep PRIVATE INFORMATION/donor/health work within Workspace/Microsoft 365 or local models; redact before using consumer contexts.

7) Use Cases (Blueprints)

A) Student Success Agent (Telegram/Coze)

Goal: Weekly check-ins with Academy students; encouragement + micro-assignments; escalate flags.

Flow: Pull this week’s module from Project → send two questions + 1 tip → collect replies → flag “low-engagement” → notify coach → produce weekly summary

Guardrails: No counseling; no PRIVATE INFORMATION storage; escalate any risk keywords to human staff.

B) Bilingual Content Agent (Lesson → Slides + Handouts)

Goal: Turn the teacher’s outline into slides and handouts.

Flow: Ingest outline + brand guide → produce slide outline with speaker notes → generate ES handout; HT on request → create image prompts (no text in images) → export package.

Checkpoints: Human review of doctrine/Scripture references; accessibility checks.

C) Market Scan Agent (Research Librarian)

Goal: Produce a 1-page, cited brief on a local market (e.g., Masaya).

Flow: Deep Research plan → extract stats with quotes/links → compare sources → bilingual brief + Evidence Log → optional chart images.

Guardrails: Require confidence levels; highlight disagreements.

D) Class Attendance & SMS Nudges (Nicaragua/DR)

Goal: Improve weekly attendance and on-time starts.

Inputs: class roster CSV, session calendar, WhatsApp/SMS template.

Actions (agent): ingest roster → identify absentees → draft reminder texts → schedule/suggest send times → log replies.

Guardrails: No private information beyond the roster; supervisor approves messages.

Outputs: reminder texts, follow-up list, and attendance delta.

Starter: “Run **Attendance Nudge**: analyze last session roster vs. today; draft friendly reminders (120–160 chars), add name, session, time; propose send window; prepare a follow-up list.”

E) Workshop Logistics Checklist

Goal: Ensure venue/AV/materials are ready 48h prior.

Inputs: session outline, venue contact, supply list.

Actions: generate 48h/24h checklists; call out risks; produce printable pack list.

Guardrails: no vendor banking details in chat.

Outputs: checklists + pack list.

Starter: “Create 48h & 24h logistics checklists for **[workshop]**, include AV, seating, printed handouts count, and backup plan for power/wifi.”

F) Community Radio Script Assistant

Goal: 2-minute program scripts.

Inputs: topic, Bible verse, call-to-action.

Actions: draft script (EN/ES/HT), time marks, alt versions.

Guardrails: doctrinal review by pastor/leader.

Outputs: scripts + show notes.

Starter: “Draft a 2-minute radio script on **[topic]** with **[verse]**, include ES and HT versions and a clear CTA.”

G) Meeting Minutes → Tasks

Goal: Ensure quick follow-through after field meetings.

Inputs: transcript/notes.

Actions: summarize decisions; create a task table (owner, due date, priority).

Guardrails: Remove any unnecessary personal details. **Outputs:** minutes and task list (EN/ES).

Starter: “Summarize this meeting and produce a task table (owner, due, status).

Provide ES summary for group chat.”

8) Best-Practice Patterns

- **Small, sharp missions** — one agent per role; avoid “kitchen-sink” assistants.
- **Plan first** — request the sub-task plan and validation criteria before taking action.
- **Evidence by default** — citations + quotes; keep the Evidence Log with outputs.
- **Bilingual standard** — EN first; ES always; HT on request.

- **Human-in-the-loop** — confirmations before external sends; reviewer sign-off.
- **Least-privilege connectors** — read-only where possible; rotate keys; log access.
- **Version & archive** — vX.Y + date; store deliverables with their Evidence Log.

9) When to Use Agent Mode vs. Alternatives

- **Agent Mode**: recurring, multi-step workflows needing tools, files, and approvals.
- **Custom GPT**: repeatable **persona + knowledge** for drafting and coaching.
- **Projects**: organize files/chats/instructions for long-running topics.
- **Deep Research**: source-backed, multi-site investigations and synthesis.
- **Study & Learn**: tutoring, drills, teach-back for learners.

10) Troubleshooting

- **Agent stalls** → ask it to print the current plan and which step is blocked; provide missing input or skip with rationale.
- **Over-scope** → split into specialized agents (e.g., Donor Brief vs. Market Scan).
- **Citation gaps** → require quotes + links; re-run cross-check pass; add Perplexity verification outside the agent if needed.
- **Permission errors** → verify connector scopes; retry with read-only access; avoid storing auth in chats.
- **Inconsistent tone** → add explicit style lines and a short **few-shot** example in instructions.

Chapter 8: Text to Image AI Tools

Introduction:

The text-to-image AI market is crowded and evolving rapidly. For entrepreneurs, ministries, and small businesses, it's challenging to know which platforms are merely hype and which are practical tools for daily work.

This research goes beyond marketing claims to identify the most effective, accessible, and affordable tools for the tasks that growing organizations actually need. We tested four leading platforms—Google Gemini ("Nano Banana"), Ideogram 3.0, ChatGPT (DALL-E), and the open-weight FLUX.1—against a series of real-world challenges.

1. Text Accuracy: Can it create a simple event flyer with clear, legible text? **2. Style Consistency:** Can it produce a uniform set of brand icons? **3. Photorealism:** Can it generate a realistic scene with specific lighting and lens effects? **4. Photo Editing:** Can it cleanly remove unwanted objects from an existing photo?

Our findings show there is no single "best" tool. The best platform depends entirely on the task. This report offers a detailed breakdown of which tool to use for each job, helping you save time, money, and frustration.

Best No/Low-Cost Text-to-Image Tools

Here is a more detailed breakdown of the top-tier, accessible tools, focusing on their specific strengths, features, and cost structures.

1) Precise Editing & Subject Consistency → Google Gemini ("Nano Banana")

- **Model:** Gemini 2.5 Flash Image (publicly nicknamed "Nano Banana")
- **Best For:** Iterative design, photorealistic edits, and maintaining a person's or object's identity across multiple images.
- **Access:** Free in the standard Gemini app and in Google AI Studio. The Google One AI Premium plan (approx. \$19.99/mo) provides higher access limits.

Why it's a top pick:

"Nano Banana" marks a significant advance in conversational image editing. Its main feature is its "state memory" or subject consistency. You can upload a photo (e.g., of a specific person, product, or pet) and then use natural language to make precise changes without the model "forgetting" the subject.

- **Precise Edits:** "Change the color of her shirt to blue," "add a hat," or "place this person on a beach background."⁸ The model will execute the command while preserving the person's face, hair, and other core attributes.⁹
- **Photorealism:** It excels at creating realistic changes, such as adjusting a room's lighting, adding objects, or swapping backgrounds.¹⁰
- **Multi-Turn Refinement:** You can "chat" with the image, applying multiple edits in a row¹¹ (e.g., 1. "Put this sofa in a living room." 2. "Okay, now make the walls navy blue." 3. "Add a large plant in the corner.")

2) Graphics with Perfectly Legible Text → Ideogram

- **Model:** Ideogram 3.0

- **Best For:** Posters, logos, flyers, social media graphics, signage, and any design where text accuracy is critical.
- **Access:** Generous free tier and affordable paid plans.

Why it's a top pick:

Ideogram has consistently been the market leader for typography. While most models struggle to render coherent words or phrases, Ideogram 3.0 was built to handle text with remarkable precision, even for long or stylized sentences.

- **Text Accuracy:** It can produce complex layouts with multiple lines of text that are spelled correctly and seamlessly integrated into the design.
- **Style Control:** It is good at following stylistic prompts for the text itself (e.g., "a movie poster with the title in a 1980s retro font," "a coffee shop logo with elegant cursive text").
- **"Magic Prompt" feature:** Enhances and adds creative details to your ideas, often leading to more intricate and beautiful compositions.

3) Open-Weight, Fast, & Cheap Photorealism → FLUX

- **Model:** FLUX.1 (from Black Forest Labs)
- **Best For:** Bulk image generation, API integration, and developers who need a high-quality, low-cost, open-weight model.
- **Access:** Not a consumer website. Accessed via API providers (Nebius, fal.ai, Replicate) or by running the open-weight model yourself.

Why it's a top pick: FLUX is a powerful, flexible, and extremely low-cost solution for high-volume needs. Because it's an open-weight model, it can be hosted by various platforms competing on price.

- **Ultra-Low Cost:** API hosts like Nebius AI Studio offer pricing as low as **\$0.0013 per image** (for the FLUX.1-schnell model) or \$0.007 (for the FLUX.1-dev model). This is significantly cheaper than closed-model APIs.
- **Speed & Quality:** It is known for producing high-quality, photorealistic images at very high speeds, making it ideal for applications that require generating large numbers of photos quickly.
- **Flexibility:** As an open-weight model, it provides developers and businesses with greater control to fine-tune it or host it themselves.

4) Generalists: ChatGPT & Google's Premium Tier

ChatGPT (DALL-E 3 / GPT-Image-1): The image generator integrated into ChatGPT is a strong all-rounder, especially for creative brainstorming, thanks to ChatGPT's advanced language understanding.

Conversational Creation: You can have a detailed conversation about what you want, and the image will be refined based on your feedback.

Cost: Image generation is included in paid plans like ChatGPT Plus (approx. \$20/mo), which offers a high usage limit (e.g., ~50 images per 3 hours).

Free Tier: The free ChatGPT tier has stringent limits on image generation, typically allowing only 2-3 images per day. This is a significant limitation, positioning the free version as more of a trial.

Google's Premium Option (Imagen 4 Ultra):

In addition to the "Nano Banana" model (best for low-cost editing), Google also offers a premium model in AI Studio called Imagen 4 Ultra. This is Google's high-end generator, designed to produce the highest-quality, most photorealistic 2K-resolution images from a text prompt. It is a **paid, per-image tool** (approx. \$0.06 per image), so it does not fit our "low-cost" category.

Use Gemini 2.5 Flash ("Nano Banana") for: Low-cost editing and keeping a subject consistent.

Use Imagen 4 Ultra for: Paid, single-shot generation of the highest possible quality.

Feature & Pricing Comparison Tables

Here are two tables to help visualize the differences between the key platforms.

Table 1: At-a-Glance Feature Comparison

Feature	Gemini ("Nano Banana")	Ideogram 3.0	FLUX.1 (via API)	ChatGPT (DALL-E)	Google (Imagen 4 Ultra)
Primary Strength	Precise Subject Editing	Perfect Typography	Low-Cost Bulk Renders	Conversational Creation	Premium Quality Generation

Feature	Gemini ("Nano Banana")	Ideogram 3.0	FLUX.1 (via API)	ChatGPT (DALL-E)	Google (Imagen 4 Ultra)
Best for...	Editing real photos, character consistency	Posters, logos, signs	API, app back-ends	Brainstorming, stylized art	Highest photorealism, print-quality assets
Subject Consistency	Excellent (Core feature)	Good (Has a feature for it)	Good	Fair (Can struggle)	Good
Text-in-Image	Fair	Excellent (Core feature)	Fair-to-Good	Good (but inconsistent)	Excellent
Image Editing	Excellent (Chat-based, inpainting)	Limited (Regenerate)	API-based (Inpainting)	Good (Chat-based, inpainting)	Limited (It's a generator, not an editor)
Photorealism	Excellent	Very Good	Excellent	Very Good	Excellent (Core feature)
Primary Interface	Web Chat (Gemini) / AI Studio	Web App (Ideogram.ai)	API / Code	Web Chat (ChatGPT)	AI Studio / API

Table 2: Cost & Access Comparison (as of late 2025)

Platform	Free Tier Access	Cost (Paid Tiers)	Approx. Cost per Image (Paid)
Google Gemini ("Nano Banana")	Generous Free Tier. Available in the standard Gemini app & AI Studio.	~\$19.99/mo (Google One AI Premium) for higher usage limits.	N/A (Included in subscription)
Ideogram	Generous Free Tier. 10 weekly credits (up to 40 images/week). Slower generation.	Basic: ~\$7/mo (400 priority credits) Plus: ~\$15/mo (1,000 priority credits)	~\$0.01 - \$0.02 (Varies by plan & image quality settings)
FLUX.1 (API)	No Free Tier (API providers may offer small trial credits).	Pay-as-you-go (API billing).	~\$0.0013 - \$0.007 (e.g., via Nebius). Extremely low.
ChatGPT	Very Limited Free Tier. Capped at 2-3 images per day.	~\$20/mo (ChatGPT Plus) for high usage limits.	N/A (Included in subscription)

Google (Imagen 4 Ultra)	No Free Tier (API-only, may include trial credits).	Pay-as-you-go (API billing).	~\$0.06 per image (Premium tier)
----------------------------	---	------------------------------	----------------------------------

Best Practice: For ministry-facing materials, avoid using synthetic people or AI-generated testimonials. Use real photos with consent; utilize AI mainly for edits, product or space mockups, icons, and non-human graphics.

Text to Image Field Testing

Goal: Compare platforms on tasks entrepreneurs and ministries actually need—photo editing, text accuracy, style consistency, speed, and cost.

Testing Protocol: For each platform, generate **exactly one** default-settings output per prompt (no retries).

Test Pack & Results

Test # 1: Text Accuracy (Poster/Signage)

Prompt: *“Design a 1080×1350 flyer titled ‘LEARN2ADVANCE ENTREPRENEUR CLINIC — SATURDAY, NOVEMBER 15, 10:00 AM – 12:00 PM — CAMINO COMMUNITY CENTER — Register: blessback.org/learn2advance’. Bold headline, high contrast, sans-serif, clean margins, export-ready.”*

Score fields: All words correct (Y/N per line), kerning/legibility (1–5), overall layout (1–5).

Generated Images

Gemini 2.5 Pro



ChatGPT 5
Thinking



Ideogram



Flux 1.ai Schenell



Analysis Details:

Platform	Title correct	Date correct	Time correct	Location correct	Register line correct	Text Accuracy %	Kerning/ Legibility (1-5)	Overall Layout (1-5)	Subtotal (Kerning+ Layout)
Flux (open-weight)	No	No	No	No	No	10	2	2	4
Ideogram	Yes	Yes	Yes	Yes	Yes	100	5	5	10
ChatGPT (DALL·E)	Yes	Yes	No	Yes	Yes	100	4	4	8
Gemini (AI Studio)	Yes	Yes	Yes	Yes	Yes	100	4	4	8

Analysis Summary:

Note: Kerning is the process of adjusting the spacing between individual letters or character pairs to improve visual harmony and readability.

- **Ideogram:** 5/5 lines correct; kerning 5, layout 5. Cleanest typography; best at exact text.
- **ChatGPT (DALL·E):** 5/5 lines correct; kerning 4, layout 4. Solid, minor typography nitpicks.
- **Gemini (AI Studio):** 5/5 lines correct; kerning 4, layout 4. Minimalist; capitalization on “Blessback.org” but words correct.
- **Flux (open-weight):** 1.5/5 lines correct (severe misspellings: “STURDAY, NOVEMBER,” bad time string, wrong URL). Kerning 2, layout 2.

Recommendation:

Ideogram, ChatGPT, and Gemini are acceptable for simple event flyers. **Flux** shines elsewhere (speed/photoreal), but not for reliable typography.

Test #2 Style Consistency (Set of 4)

Prompt A (character/icon): “Create 4 square icons of the same mascot: a simple line-art seed in 4 poses (teach, mentor, build, bless). Keep exact line weight and palette (#0B5FFF, #FDB515).”

Score: Palette exactness (hex match Y/N), line weight consistency (1–5), character resemblance across set (1–5).

Generated Images

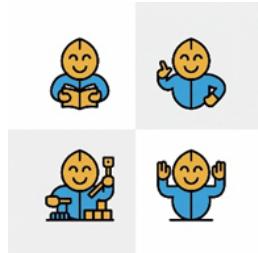
Gemini 2.5 Pro



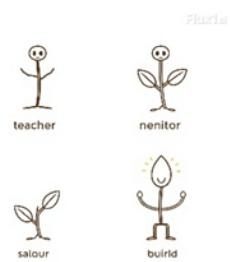
ChatGPT 5 Thinking



Ideogram



Flux 1.ai Schenell



Style Consistency Results

Platform	Palette exact (#0B5FFF & #FDB515)	Line weight consistency (1 to 5)	Character resemblance across set (1 to 5)	Notes	Subtotal (Consistency + Resemblance) /10
Flux (open-weight)	No (brown line art; slight yellow highlights)	3	2	Style drifts to stick figures; seed motif weak; two captions misspelled (nenitor, salour, buirld).	5
Ideogram	No (gold close; navy #0B5FFF)	5	5	Excellent single-stroke line style; uniform mascot; strong composition.	10
ChatGPT (DALL·E)	No (very close but not exact hex)	5	5	Clean flat-vector set; consistent pose language; outline weight uniform.	10

Gemini (AI Studio)	No (glow gradients; off-hex)	4	4	Kept mascot form but introduced glow/gradients; not line-art; soft edges.	8
--------------------	------------------------------	---	---	---	---

Analysis Snapshot:

- **Ideogram** — Palette not exact (navy vs #0B5FFF) but **perfect consistency & resemblance (5/5 + 5/5)**. Best adherence to “single-stroke line art.”
- **ChatGPT (DALL·E)** — Palette very close but not exact; **consistency & resemblance both 5/5**. Clean vector look.
- **Gemini (AI Studio)** — Introduced glows/gradients (not line-art); **4/5 consistency, 4/5 resemblance**.
- **Flux (open-weight)** — Drifted to stick figures; seed motif weak; **3/5 consistency, 2/5 resemblance**, and misspelled labels (“nenitor,” “salour,” “buirld”).

Recommendations:

- For **brand icon packs** where uniform stroke and character identity matter, favor **Ideogram** or **ChatGPT Image**.
- If **exact hex** is non-negotiable, plan a quick **vector post-pass** (Canva: replace fills with #0B5FFF/#FDB515). Most AI models won’t hit hex codes perfectly.
- To reduce drift: include “**flat vector SVG, single 3px stroke, no gradients, exact fills #0B5FFF and #FDB515**” in the prompt, then run one refinement pass asking the model to “match color codes exactly.”

Test #3: Photoreal Scene + Prompt Following

Prompt: “A small business sewing class in soft morning light, 35mm lens look, shallow depth of field, fabric texture visible, natural skin tones, 2048×1536.”

Score: Prompt adherence (1–5), artifacting/hands (1–5), fabric realism (1–5), time-to-first-image.

Generated Images

Gemini 2.5 Pro

ChatGPT 5 Thinking

Ideogram

Flux 1.ai Schenell



Photoreal Scene Test 3 Scores

Platform	Prompt adherence (1–5)	Artifacting/Hands (1–5)	Fabric realism (1–5)
ChatGPT (DALL·E)	4	5	5
Ideogram	5	4	5
Flux (open-weight)	3	4	3
Gemini (AI Studio)	5	4	5

Analysis Snapshot:

- **Gemini (AI Studio): 5/5 adherence, 4/5 hands, 5/5 fabric realism.** Soft morning window light, shallow DoF, convincing fabric grain; very photoreal.
- **IDEogram: 5/5 adherence, 4/5 hands, 5/5 fabric realism.** Vintage machines, great light and texture; slight hand ambiguity on one subject.
- **ChatGPT (DALL·E): 4/5 adherence, 5/5 hands, 5/5 fabric realism.** Strong realism and skin tones; composition and DoF are a touch flatter vs the 35mm brief.
- **Flux (open-weight): 3/5 adherence, 4/5 hands, 3/5 fabric realism.** Backlit scene fits the class vibe, but the fabric looks softer/plasticky; less texture detail.

Recommendations:

- For **photoreal classroom scenes** with specific lens/lighting cues, **Gemini** and **Ideogram** delivered the most faithful results in this run.
- **ChatGPT** produced the cleanest hands and skin, but with slightly less of a “35mm look.”
- **Flux** remains strong for speed-to-cost, but you’ll likely need an extra pass to bring out the fabric’s microtexture.

Test #4: Simple Background Removal / Object Erase

Prompt/task: Remove an unwanted power cable and straighten the horizon; output PNG.

Score: Clean removal (1–5), edge quality (1–5)

Image uploaded to each platform.



Generated Images

Gemini 2.5 Pro



ChatGPT 5 Thinking



Ideogram



Flux 1.ai Schenell

Not Available

Simple Remove/Straighten – Test 5 Scores

Platform	Clean removal (1–5)	Edge quality (1–5)
ChatGPT (Image/DALL·E)	5.0	5.0
Gemini (AI Studio)	5.0	5.0
Ideogram	4.0	4.0

Summary:

- **ChatGPT (Image/DALL·E)** — Clean removal **5/5**, edge quality **5/5**. Horizon looks straight; sofa/rug textures preserved—production-ready.
- **Gemini (AI Studio)** — Clean removal **5/5**, edge **5/5**. Very natural patch; wood grain + cushion seams intact.
- **Ideogram** — Clean removal **4/5**, edge **4/5**. Cable gone, but I can see faint banding/tiling near the top edge; slight smoothing on a cushion seam. However, Ideogram added a blanket and a book to the original image.

Prompting Best Practices (Learned from This Test)

Based on the test results, several prompting strategies emerged that significantly improved the quality and accuracy of the generated images. Using precise, technical language is key to moving from a generic image to a production-ready asset.

Here are the top best practices identified in this research:

Be very specific with text requests: for graphics, don't just ask for text; specify precisely what you want. Provide the precise string in quotation marks. Include layout instructions such as "bold headline," "high contrast," "sans-serif," and "clean margins" to guide the final design.

Use Photographic Language to Achieve Realism: The most successful photorealistic images come from prompts that incorporate technical photography terms. Specify the:

- Lighting: "soft morning light."

- Lens: "35mm lens look"
- Aperture: "shallow depth of field.d"
- Required Detail: "fabric texture visible.e"

Use Technical Terms for Design Styles: To create a consistent icon set, the prompt specified "simple line-art." For better results, the analysis suggested using terms like "flat vector SVG," "single 3px stroke," and "no gradients" to prevent the model from adding unwanted effects, such as the glows Geminiaaddede.

Specify Exact Hex Codes for Colors: Even though most models initially failed to match the exact hex codes (#0B5FFF, #FDB515), providing them resulted in a palette much closer to the target than simply saying "blue and yellow."

Define Your Canvas and Format: Specify the exact aspect ratio or resolution (e.g., "1080x1350" or "2048x1536") to set the shape of your image. For editing tasks, you can also specify the output format, such as "output PNG".

Be Direct and Clear with Edits: For photo editing and cleanup, simple commands work best. Clearly state the specific actions, like "Remove an unwanted power cable" and "straighten the horizon."

Plan for Iterative Refinement: Your first version is often just a starting point. As the icon test concluded, the best workflow usually involves a refinement pass. Be ready to use a follow-up prompt, such as "match color codes exactly," to finalize your image.

Recommended Conclusion & Executive Summary

Conclusion & Final Recommendations:

Our hands-on testing confirms the primary hypothesis: the best AI image generator is entirely task-dependent. No single platform excelled at all four tests, but clear winners emerged for specific, everyday business needs.

Executive Summary of Test Results:

For Text-Heavy Graphics (Flyers, Posters): **Ideogram** is the undisputed winner. It was the only platform that delivered perfectly spelled, well-composed text in our flyer test, making it the top choice for any graphic design project involving typography.

For Photo Editing & Cleanup: **Gemini ("Nano Banana")** and **ChatGPT** are the clear leaders. Both platforms performed flawless, production-ready object removal in a single pass, perfectly preserving textures and seams. Gemini's core strength is this kind of precise, conversational editing.

For Brand & Icon Consistency: **Ideogram** and **ChatGPT** provided the most consistent and uniform icon sets. Both maintained excellent character resemblance and line weight, though neither hit the exact hex color codes requested.

For Photorealistic Scenes: **Gemini** and **Ideogram** delivered the most faithful results, excelling at complex prompts for lighting ("soft morning light"), lens effects ("35mm look"), and micro-detail ("fabric texture").

For Bulk/API Use: **FLUX** remains a powerful, low-cost option for developers. However, our tests showed it is highly unreliable for tasks requiring text or icon accuracy and is best used for high-volume photorealistic generation, not design work.

Chapter 9: Text to Video AI Tools

Sora — OpenAI's AI Video (What it is, How to use it, Limits, and When to choose it)

1) What Sora Is

Sora is OpenAI's **text/image/video → video** generator. It creates short, cinematic clips from prompts, can extend or remix footage, and integrates with ChatGPT for iterative edits. Think of it as **storyboarding + shot creation** on demand.

Best for: concept teasers, explainer cutaways, abstract backdrops, B-roll, motion backgrounds for slides, and social posts. Avoid using synthetic people as **real testimonials** or depictions of sensitive events.

2) How to Use Sora (Step-by-step)

- **Choose access**
 - **ChatGPT (Plus/Pro/Business/Enterprise):** Create a video from a prompt inside ChatGPT.
 - **Sora web app** (where available): sora.com → Sign in with your OpenAI account.
- **Start with a storyboard**
 - Write 3–5 beats (scenes), 3–6s each.
 - Decide aspect: **16:9** (YouTube/web), **9:16** (Reels/Shorts), **1:1** (IG feed).
- **Prompt the first shot** (use the template below).
- **Refine:** ask for "keep composition; slightly warmer light; add gentle camera push-in."

- **Extend or generate the next shot** while keeping the character/costume/location consistent.
- **Download and add captions/voice/music** in your editor (CapCut/Canva/Slides/Premiere).

First-shot prompt template

> **[Photoreal/stylized / flat-animation] [subject] [doing what] in [setting]. Camera:** [wide/medium/close], [eye-level/45°], gentle [pan/tilt/push-in]. **Lighting:** [soft daylight/golden hour/studio softbox]. **Style & color:** [natural palette / BBA blue-gold accents]. **Motion:** [wind in trees/steam from cups]. **Aspect:** [16:9 / 9:16 / 1:1]. **Quality:** natural proportions, clean edges. **No on-screen text.**

3) Limits to Keep in Mind

- **Clip length & resolution:** expect **short clips** (tens of seconds) and **HD (720p–1080p)** outputs depending on your plan.
- **Audio:** today, treat Sora clips as **silent**; add narration/music later.
- **Queues/credits:** generations may **wait in a queue**; higher tiers get priority.
- **Policy:** strict rules on realistic people, sensitive topics, and watermarks/attribution—follow them.
- **Physics & continuity:** complex hand/object interactions or long multi-scene continuity may require retries and tighter prompts.

4) When to Use Sora vs. Other Video Models

- **Use Sora when** you need **flexible storytelling**, multiple styles, or tight integration with ChatGPT (research → script → shots → captions).
- **Consider Google Veo** when you need **native audio/lip-sync** and short, **cinematic 8–10s** clips.
- **Consider Runway (Gen-3/Gen-4)** for **editor-style controls** (keyframes, image→video, extend) and integrations for production teams.
- **Consider the Luma Dream Machine** for **fast iteration** and built-in **audio scoring**.

- Consider **Pika** for **social-first effects** (swaps/effects) and budget-friendly credit plans.
- Consider **Kling** for longer clips (up to ~2 min) with strong **motion realism**.

5) BBA & Ministry Use Cases

A) Class promo (9:16, 12s)

Goal: Invite to next week's Academy session.

Prompt idea: warm market street at sunset → students collaborating at a table → closing shot with space for headline.

Output: silent video; add ES/HT captions later.

B) Scripture meditation background (1:1, 20s loop)

Goal: Visual loop behind a verse in a service or devotional.

Prompt idea: gentle sunrise over hills; slow clouds; soft camera drift.

Output: overlay verse text in Slides/Canva (not in the AI video).

D) Workshop opener (16:9, 15s)

Goal: Energize the entrepreneurship bootcamp.

Prompt idea: close-ups of hands sketching, assembling products, tap-to-pay moment; rhythmic cuts.

Output: add upbeat royalty-free music.

E) Prayer & outreach stories (9:16, 2×8s shots)

Goal: Abstract visuals (no synthetic people) to accompany real testimonies recorded separately.

Prompt idea: hands passing notebooks, skyline time-lapse, light through chapel windows.

6) Prompt Patterns that Improve First-Try Success

- **Storyboard beats** (3–5) with **duration per beat**.
- **Subject + action + setting + camera + lighting + motion**.
- **Negative constraints**: “no on-screen text, no logos, natural proportions, no extra limbs, clean edges.”
- **Continuity tokens**: re-use character/prop descriptors across shots (“same blue notebook; same apron”).
- **Aspect & use** up front (9:16 Reels vs 16:9 slides).
- **One change at a time** in iterations (lighting only, then camera move).

7) Post-Production Workflow

- Download clip(s) from Sora.
- Add **voice/music/captions** in CapCut/Canva/Slides; export H.264 1080p.
- Add **bilingual** captions (EN→ES; HT on request).
- Run the **Creative Integrity Checklist**: no synthetic people as real testimonials, no trademarks, respectful cultural portrayals, and add attribution when required.
- Store outputs in the Project with an **Evidence Log** (prompt, model, date, sources for any facts shown).

8) Safety & Ethics

- **No deepfakes/impersonation.** Do not depict real private individuals.
- **Ministry standards:** modesty, dignity, and cultural respect.
- **Children & vulnerable people:** do not synthesize likenesses; use abstract imagery.
- **Copyright:** avoid recognizable brand marks; source any third-party assets under license.
- **Attribution:** follow OpenAI's rules on watermarks/metadata and platform disclosure where required.

9) Quick Compare — Sora vs. Others

Tool	Notable strengths	Typical limits
Sora	Storyboarding, flexible styles, tight ChatGPT loop	Short clips; silent output; queues/credit limits
Google Veo	Native audio & strong lip-sync; cinematic 4K short shots; Vertex & Workspace tie-ins	Often shorter clips; enterprise setup for full features
Runway Gen-3/4	Production controls, image→video & keyframes, extend	Shorter durations, credits, web pipeline

Tool	Notable strengths	Typical limits
Luma Dream Machine	Fast iteration; built-in audio scoring ; mobile & web	Variable realism; credits; shorter shots
Pika	Social-first effects (swaps/effects), low cost, mobile app	Consumer-oriented; stylized outputs
Kling	Longer clips and motion realism; consumer pricing	Access varies by region; the interface is less English-centric

10) Troubleshooting

- **Hands/props look odd** → pull back to **waist-up**, simplify background, add “accurate anatomy, clean edges.”
- **Too busy for captions** → ask for a **clean background** or “leave top/bottom 20% empty.”
- **Choppy motion** → specify camera move (slow push-in), wind, particles; keep subject centered.
- **Wrong color/mood** → explicitly set lighting palette and brand accents.
- **Queue delays** → shorten duration, reduce concurrent generations, try off-peak hours.

Sora Editing Features — Re-cut, Remix, Blend, Loop

What you can do inside Sora's editor

- **Re-cut**: Trim or extend a clip and rebuild it in a **Storyboard** timeline. Set aspect, resolution, duration, and generate variations.
- **Remix**: Keep the base shot but change elements (lighting, camera move, background, style) via a revised prompt; compare alternatives.
- **Blend**: Transition or merge elements from one Sora video into another for intros/outros or B-roll between scenes.
- **Loop**: Create a seamless loop from any section—ideal for motion backgrounds behind captions or verses.

- **Storyboard editing:** Build a sequence of “cards” (text, **image**, or **video** inputs) on a timeline—reorder cards, change spacing to control cuts, or convert an existing generation into a storyboard with **Re-cut**.
- **Image/Video-to-Video:** Start from an **uploaded image or video** and have Sora transform or extend it (rather than pure text → video).
- **Library tools:** Compare variations, **favorite**, **share link**, **download MP4**, and save files to folders for reuse (higher tiers support watermark-free, higher-res downloads).

Quick how-to (editing flow)

- Open a generated clip → click **Re-cut** (or **Remix/Blend/Loop** as needed).
- In **Storyboard**, add/reorder **cards** (text, image, or video). Leave small gaps between cards to soften cuts.
- Set **aspect**, **resolution**, **duration**, and **# of variations** → generate and compare.
- **Download** the selected take and finish audio/captions in your editor (CapCut/Canva/Slides/Premiere).

Practical tips for clean edits

- **Continuity:** Repeat identifiers across cards (e.g., “same blue notebook,” “same apron,” “golden-hour light”) to keep characters/props consistent.
- **One change at a time:** When results drift, tweak only the **lighting** or **camera move**, then regenerate.
- **Leave space for type:** “No on-screen text; leave top/bottom 20% clean.” Add captions/verses later.
- **Looping backgrounds:** Use **Loop** on calm sections; keep camera motion slow (gentle push-in/pan) to avoid jump seams.

How Sora's editor compares

- **Runway (Gen-3/Gen-4)**: adds keyframes and inpainting/masking for surgical fixes—use when you must remove objects or control motion precisely.
- **Luma Dream Machine**: strong extend/loop and “modify video” for transforming real footage while keeping motion—suitable for longer looping backdrops.
- **Pika**: lightweight in/out-painting, extend, and lip-sync for fast social edits.
- **Google Veo**: shorter cinematic shots with native audio/lip-sync—generate there, then finish edits in your NLE if sound-in-one is critical.

Limits to remember (editing)

- **Duration & resolution**: In ChatGPT today, typical tiers: up to **720p/10s** on entry tiers and **1080p/20s** on higher tiers; complex edits may wait in a **queue**.
- **No fine masks**: Built-ins (Re-cut/Remix/Blend/Loop + Storyboard) are high-level. For per-pixel fixes, hand off to Runway/Luma and bring the result back.

Veo 3

What it is

Veo 3 is Google’s text-to-video/animate-image capability designed to quickly produce short, loopable clips and simple animations from prompts. For our purposes, it’s a practical way to generate **neutral, non-people B-roll, diagram animations, and micro-demonstrations** you can layer with captions in Slides, Canva, or a basic editor. The goal is instructional clarity and speed—not cinematic art. Think: a 12–15-second loop that opens a class video, a moving background that reinforces a concept, or a minimal “how it’s made” clip for social.

Access & setup (first time)

1. Open Google AI Studio and choose the **video** (Veo 3) experience.
2. Decide clip **length** (we recommend 8–15 seconds for teaching intros or social).
3. Choose **aspect ratio**: 16:9 for slides/YouTube; 9:16 for Shorts/Reels; 1:1 for square posts.

4. Write a **constraints-rich** prompt: camera framing, motion type, “no people,” space for lower-third text, tone (calm/energizing), and loop behavior.
5. Generate two or three alternatives; pick the most legible on mobile.

Controls that matter

- **Length & motion:** For intros, 10–12 seconds with **subtle** motion avoids distraction.
- **Framing:** “Leave lower-third clean for captions” or “center calm area” ensures text legibility.
- **Texture & brightness:** Ask for “low-detail background,” “even lighting,” or “gentle camera push-in.”
- **Looping:** Specify “seamless loop” if the interface supports it, or plan to cross-fade in your editor.

Use Case A — Business Academy: “Silent Intro B-Roll + Diagram Animation” Pack.

Objective

Give instructors a reusable set of opening clips and a short animated diagram to visualize a core idea (e.g., price → cost → margin; cash-in vs. cash-out) without resorting to stock footage of people.

Workflow (step-by-step)

1. **Silent intro B-roll (12 seconds, 16:9)**
 - o Prompt:
“Generate a **12-second loop** of a **clean desk scene** with a **notebook, calculator, and soft window light**. **No people, no hands, no logos**. **Subtle camera push-in**. Leave the lower third clear for captions. Tone: calm, focused. Aspect 16:9.”
 - o Export as MP4, name intro-deskloop-12s-v1.mp4.
2. **Concept loop: cash-flow metaphor (10 seconds)**
 - o Prompt:
“Create a **10-second loop** abstract animation suggesting **cash flowing** through a pipeline from **left (in)** to **right (out, with colors, simple shapes, no text, no logos)**. **Gently move the center to keep it clean** for overlay labels. Aspect 16:9.”
 - o Export as cashflow-abstract-10s-v1.mp4.
3. **Diagram animation: price → cost → margin (8 seconds)**

- o Prompt:
“Produce an **8-second** minimal animation showing three boxes: **PRICE** → **COST** → **MARGIN** as **placeholders only** (no text rendered). Animate left-to-right emphasis: box 1 grows, box 2 subtracts, box 3 remains as the remainder: flat **geometry, high contrast**, no people. **Leave the top area clean** for later text overlay. 16:9.”
- o Export as pcm-diagram-8s-v1.mp4.

4. Assemble in Slides or a simple editor

- o Insert clip on first slide (auto-play, mute).
- o Add captions/transitions.
- o For the diagram, overlay titles and arrows in Slides to keep it editable.

Quality checks

- Text overlays are readable at 75% zoom on a laptop and on a phone.
- Motion is gentle; no strobing or busy textures behind text.
- Loops end where they began, or your editor cross-fades the last 0.3s.

Pitfalls & fixes

- **Too much movement** → request “subtle motion,” “reduce movement amplitude,” “calm background.”
- **Cluttered lower-third** → add “reserve lower-third (15% height) with minimal detail.”
- **Color vibrates** → ask for “softer palette,” “avoid aggressive contrast near text areas.”

Production checklist

- Export 16:9 MP4 at 1080p.
- Store in the module’s **/video** folder; keep a README listing clip purpose and best-use examples.
- Always pair with a transcript/caption file for accessibility if you add narration later.

Use Case B — Micro-Business: “How-It’s-Made” Micro-Clip (Silent) for Social

Objective

Help tiny businesses (taquería, panadería, peluquería) publish short, tasteful process

clips without showing faces or hands, keeping the focus on craft and product. These are ideal for Reels/Shorts with on-screen text, and you can reuse them weekly.

Workflow (step-by-step)

1. **Prompt for masa → tortilla (9–10 seconds, 9:16):**
“Create a **9-second** close-up **vertical** clip that **suggests the process** of pressing masa into a tortilla and placing it on a **comal**. **No faces or hands**; show **press, release, placement** through tools and shadows only. **Soft steam, warm light**, space at top for text.”
2. **Prompt for balayage foil placement (hair color) (10 seconds, 9:16):**
“Generate a **10-second** **vertical** clip implying **foil placement** for hair coloring using **tools only** (foil, comb, brush). **No faces/hands**; **clean salon surface**; **cool neutral light**. **Slow tilt** movement. Reserve **top 20%** for captions.”
3. **Add captions in-app**
 - o Keep lines ≤ 6 words, two lines max.
 - o Example: “Press. Heat. Rest.” “Shine, protect, care.”
4. **Post cadence**
 - o 1–2 clips per week; rotate process steps; reuse background music.
 - o Reply to comments with an additional care tip in text, not DMs.

Quality checks

- Hook in first **1 second** (“Press → Sizzle → Rest”).
- Vertical composition leaves space for captions and stickers.
- Colors are natural, not neon; avoid flicker by using even light.

Pitfalls & fixes

- **Too abstract** → add recognizable tools and surfaces to prompt.
- **Distracting speed** → request “slower motion,” “gentle camera tilt only.”
- **Compression artifacts** → export at native platform resolution; avoid third-party compression before upload.

Production checklist

- Export vertical 1080×1920 MP4.

- Save template captions in your notes app and paste them into each new post.
- Keep 3–4 evergreen clips pinned for new followers grasp your craft.

Chapter 10: Turn Text Into Visuals ([Napkin.ai](#))

1) What Napkin.ai Is

Napkin.ai converts your **text** into clean, editable **visuals** (diagrams, flowcharts, timelines, mind maps, comparison tables, and infographics) that you can polish and export for slides, handouts, or social media.

Best for: turning policies, class notes, and donor updates into quick diagrams that look on-brand—without design work.

2) Key Features

- **One-click visuals from text** — paste/ import text → click **Generate Visual** → pick from multiple options.
- **Editable vector graphics** — change icons, shapes, connectors, colors, and fonts; rearrange nodes/boxes.
- **Custom Styles** — set brand colors & fonts; apply to new or existing visuals; share a **Style ID** with teammates.
- **Custom Generation** — specify **visual type** (flowchart, mind map, timeline, table, comparison chart), **orientation** (portrait/landscape/square), **detail level** (summary/detailed), and “**stay close to my text**.”
- **Teamspace & collaboration** — invite editors, co-edit in real time, and comment via built-in highlights.
- **Import** — paste text or upload **DOC/PDF/PPT/HTML/MD**; copy-paste images.
- **Export** — download **PNG/SVG/PPT/PDF**; export the **entire document as PDF**.

3) Getting Started

- **Create a new Napkin** and paste your paragraph(s) or import a file (e.g., policy PDF).
- Hover the text and click **Generate Visual** (Spark).
- Pick a layout → click to insert → **Edit** text, icons, and connectors.
- Open **Style** → choose a built-in look or create a **Custom Style** (brand colors, fonts).

- **Export** the visual (PNG/SVG/PPT/PDF) or the **whole doc** (PDF).

4) Customization & Branding

- **Brand colors & fonts:** create a **Custom Style** once, then apply it to all visuals (or “Apply to All”) to unify handouts and slides.
- **Iconography:** swap any icon via **Spark Search**; keep consistent icon families across lessons.
- **Layouts:** if you discover a design that fits (e.g., “2x2 comparison”), copy its **layout name** and regenerate matching visuals later.
- **Detail & wording:** use **Custom Generation** → **Detail level** (Summary vs Detailed) and **Stay close to my text** for policy-accurate phrasing.

5) Language & Accessibility

- Generate visuals in multilingual options.
- Keep slides readable: use short labels, maintain high contrast, and reserve space for captions; avoid putting paragraphs inside nodes.

6) Downloads & Sharing

- **Single visual:** **PNG** (fast), **SVG** (vector), **PPT** (editable in PowerPoint), **PDF**.
- **Complete doc:** **PDF** (single long page or paginated).
- **Share link:** optional web view for text + visuals together.

7) Napkin.ai vs ChatGPT — what's similar/different?

Task	Napkin.ai	Closest feature in ChatGPT	Key differences
Turn text → diagram	Native one-click, many layout options	Create images/diagrams (image generation) or Mermaid/ASCII via chat	Napkin outputs editable vectors + PPT/SVG ; ChatGPT images are raster by default; Mermaid requires a renderer

Task	Napkin.ai	Closest feature in ChatGPT	Key differences
Brand styling	Custom Styles (colors, fonts) applied globally	System prompts + manual brand checks	Napkin enforces visual styling; ChatGPT enforces language/tone , not design.
Collaboration	Teamspace, real-time co-edit, highlights	Share chats or Projects; no vector co-editing	Napkin is a mini design surface; ChatGPT is a writing/research workspace
Research & drafting	Not a research LLM	Deep Research, Projects, Agents	Use ChatGPT for source-backed text; drop finished text into Napkin for visuals.

Rule of thumb: Draft & fact-check in **ChatGPT** → create/brand visuals in **Napkin** → assemble in Slides/Docs.

8) BBA/Ministry Use Cases

- **Safeguarding flowchart** — intake → decision → reporting path; export PPT for training.
- **Entrepreneur pricing map** — costs → margin → price ladder; export PNG for a one-pager.
- **Donor update infographic** — 3 metrics + 2 stories + next steps; export PDF for email.
- **Market scan timeline** — events & prices over a quarter; export SVG for crisp projection.
- **Workshop agenda mind map**: topics → activities → materials; share the link with facilitators.

9) Best-Practice Workflow (Napkin + ChatGPT)

- In **ChatGPT Projects**, produce the **outline/copy** (EN → ES; HT on request).
- Paste the text into **Napkin**, generate visuals, and apply **Custom Style**.
- Export **PPT** to integrate into your slide deck; export **PNG** for social; export **PDF** for handouts.
- Store assets in the Project with a short **change log** and version (v2025-09).

10) Plans at a Glance (for planning budgets)

- **Free**: generous weekly credits, PNG/PDF export; Napkin watermark.
- **Plus**: more credits/month, **PPT & SVG** export, remove branding, and a few **Custom Styles**.
- **Pro**: high credits, **unlimited Custom Styles**, **font uploads**, exclusive designs, team billing, and optional top-ups.

Chapter 11: Google AI Studio

What Google AI Studio Is

Google AI Studio is Google's free, browser-based workspace for experimenting with their latest language and multimodal models (e.g., Gemini 2.5 Pro/Flash, Gemma). You can set the temperature (to balance creativity versus precision), specify format constraints (such as exact section and bullet counts), and compare outputs side by side across models. For the Business Academy, AI Studio is where instructors refine "gold prompts" that turn raw outlines into publish-ready module packs. For entrepreneurs, it's where you establish a consistent voice, structure, and readability for product/service pages, SOPs, and content plans—before transferring templates into production tools (Docs, Slides, Sheets) or your website.

How to access and set up (first-time)

1. Open **Google AI Studio** in your browser and sign in with a Google account.
2. You'll land on a prompt canvas. At the top or right panel, you'll see: **Model selector**, **Temperature**, and sometimes **Max output** or additional toggles.
3. Choose a model: start with **Gemini 2.5 Pro** for careful, complete outputs; test **2.5 Flash** for faster brainstorming.
4. Set **Temperature** based on task type:

- o **0.2–0.4:** procedural (lesson agendas, SOPs, checklists).
- o **0.5–0.7:** balanced (product pages, course summaries).
- o **0.8–1.0:** creative ideation (hooks, social captions).

5. Click **Compare** to run the **same prompt** through two models or two temperature settings. This A/B view is the most reliable way to decide which combination you'll standardize in your prompt library.

The three building blocks of every strong prompt (copy/paste pattern)

(1) System block (role + mission + constraints)

You are a practical small-business and teaching assistant.

Mission: produce outputs that are structurally consistent, concise, and directly usable.

Constraints: follow the requested sections exactly; short sentences; no fluff.

If info is missing, state an assumption in [brackets] and proceed.

(2) Task spec

- State the deliverable and audience (“module pack for adult learners,” “product page for busy parents”).
- Specify order of sections, bullet counts, word caps, reading level, and any style rules (“no buzzwords,” “plain language”).

(3) Inputs

- Paste the outline, notes, or draft copy (scrub private data).
- If needed, create **fictional** personas, SKUs, or example scenarios to make the prompt realistic without revealing PII.

Core controls you'll use every time

- **Model choice:** Start with Pro, then try Flash if outputs feel long-winded.
- **Temperature:** Lower for reliability; higher for creative variety—then bring it back down to finalize.
- **Output constraints:** Exact section order, bullet counts, word caps, and reading level.
- **Compare:** Decide which model best aligns with your constraints and tone.

Use Case A (Business Academy): “Module Pack” Standardizer

Goal: Standardize the process for creating a complete, instructor-ready module pack from any new topic.

What you'll do

- Take a rough outline or notes for “Pricing Basics,” “Cash Flow 101,” or any topic.
- In AI Studio, paste the **System block** above and a focused **Task spec** that demands predictable sections.
- Run on **Gemini 2.5 Pro @ temp 0.3–0.4**; then use **Compare with 2.5 Flash @ temp 0.3–0.4**.
- Keep the version that best obeys the structure, and cut the other.

Starter prompt (paste into AI Studio)

“Role: small-business instructor. Mission: produce a module pack from the outline below. Constraints: return sections in this order—(1) Lesson Overview ≤150 words, (2) 6 Learning Objectives, (3) 8-segment Agenda with minutes, (4) 6 Discussion Questions, (5) 1 Practice Assignment with rubric. Reading level ~7th grade. Avoid buzzwords. If context is missing, state the assumption and proceed.

Outline: [paste outline here, scrubbed of private data].”

Expected output

- A clean, five-section module pack that any instructor can use.
- Agendas come pre-timed, questions are ready, and the assignment includes a rubric.
- Minimal polish needed before moving to Slides/Docs.

Operational tips

- Save the winning “System + Task” as **ModulePack-v1.0** with your chosen model/temperature; store it in a shared prompt library.
- For a new module, paste a new outline, run the same prompt, then export the result directly to Docs/Slides.

Common issues (and fixes)

- *Overwriting caps:* If a section runs long, add “truncate cleanly if over cap; do not spill.”
- *Vague objectives:* Add “Each objective begins with a measurable verb (identify, calculate, compare...).”

Use Case B (micro-business): Product/Service Page Trio

Goal: Develop three consistent description styles for a single service or product, allowing you to choose a voice that suits your community.

What you'll do

- Identify one offering (e.g., a **peluquería's** “Starter Color Package” for short hair, or a taquería’s family-size tray).
- Ask for three distinct voices: A) plain facts, B) benefits-first, C) friendly & welcoming.
- Include **two Instagram caption variants** per version (≤ 120 characters) to test engagement.
- Keep it neighborhood-friendly and straightforward; avoid slang that may not generalize.

Starter prompt (paste into AI Studio)

“Create three versions of a service description for ‘Starter Color Package’ (short hair): **A**) plain facts, **B**) benefits-first, **C**) friendly & welcoming. For each: 70–90 words max + two IG captions ≤ 120 characters. Keep language simple; avoid slang. If you include Spanish, mirror the English meaning.”

Expected output

- Three style options, each concise and structured consistently.
- Short captions are ready for social testing.
- You select the voice that fits your barrio and brand values.

Operational tips

- Use **Compare** to see which model delivers crisper, more culturally respectful phrasing at the same temperature.
- Once you pick a voice, save a “house style” mini-guide from the winning description (sentence length, tone cues, headline patterns).

Common issues (and fixes)

- *Too salesy*: Add “avoid hype; no superlatives unless factual.”
- *Jargon*: Add “explain terms in plain language; 7th-grade reading level.” a

Troubleshooting & “gotchas.”

- **Model ignores structure** → Lead with: “Return sections exactly as specified, in that order.”

- **Too long or fluffy** → Lower temperature and add hard caps to each section (“≤X words”).
- **Inconsistent tone across runs** → Freeze **one** model + temperature; copy the system block into every session.
- **Uncertain facts** → Require the model to mark assumptions in [brackets] and proceed so that you can check them quickly.
- **Multilingual needs** → Generate first in your native language for structure, then request another language mirror in a second pass to avoid overflow.

Moving to production

- **Module packs**: Paste into Docs, format headings, drop into your slide template.
- **Service pages**: Load the chosen variant into your site; keep the other two as future A/B tests.
- **Prompt library**: Store the final system + task + notes (model/temperature) and one high-quality example for reference.

Gemini Live (Stream): Real-Time Screen-Share & Voice for Practical Edits

What Gemini Live (Stream) is

Gemini Live (Stream) allows you to chat with the model while sharing your screen or tab. Instead of copy/pasting text back and forth, you walk through a **slide deck, web page, spreadsheet, or doc** and get **immediate, context-aware suggestions**. The magic isn't chit-chat—it's the ability to **see** what you're seeing and return precise, actionable edits, on the spot.

Privacy stance: Live is most effective for non-sensitive assets (slide decks, public-facing pages, fictionalized spreadsheets). Don't share student PII or real customer records.

How to access and get reliable feedback

1. Launch **Gemini Live/Stream** inside Google AI Studio.
2. Select your microphone, then choose which window/tab to share.
3. Say what you're trying to improve (clarity, brevity, readability, mobile scan).
4. Ask for **specific outputs**: “Rewrite this line to ≤10 words,” “Give me a 3-step process,” “Show a diagram I could draw,” “Suggest a slide merge.”

5. **Apply edits in real time** while sharing; confirm with the model that the new version meets your constraints.

Business-only use cases

Use Case A (Business Academy): Slide Deck Critique Session

Scenario: You have a 12-slide deck for **Pricing Basics**. It's too wordy on slides 3, 5, and 7; the transitions feel abrupt; and the images are decorative rather than helpful.

What you'll do (live prompts to say out loud)

- “While I scroll, call out any slide that’s too dense.”
- “For each dense slide, give me a 3-bullet replacement (≤ 12 words each), plus a simple diagram I can draw to explain price, cost, margin.”
- “Suggest where to merge slides so the deck stays at 12 total.”
- “Identify any term that needs a short definition at the bottom.”

What you'll get

- A slide-by-slide **fix list**: exact bullets to swap in, diagram suggestions (“3-box cost → price → margin”), and slide merges (“Combine 3 + 4; turn paragraph into a step-by-step”).
- A short **glossary add-on** (e.g., “COGS = direct costs to produce/deliver,” “Gross margin = (price–COGS)/price”).
- The deck stays within your 12-slide limit but communicates more quickly.

Why this works

- You’re not asking for abstract advice—you’re asking for **edits**.
- The model sees your slide context: spacing, line length, and visual clutter.
- You can fix slides in the moment, then immediately validate with the model (“Does this satisfy the 3-bullet ≤ 12 words rule?”).

Follow-through

- Save the revised deck as **PricingBasics-Deck-v2.0**.
- Add the “3-bullet ≤ 12 words” micro-rule to your teaching style guide.
- Capture a one-page “Before/After/Why” so other instructors can replicate the edits.

Use Case B (micro-business): Menu/Service Card Tune-Up

Scenario: You run a **barber shop, nail salon, or taqueria**. Your laminated service/menu card works, but customers still ask the same questions or hesitate before ordering/booking.

What you'll do (live prompts)

- “Point out where a first-time customer may hesitate or get confused.”
- “Rewrite long lines into **≤7 words**; recommend icons/section labels.”
- “Give me a **3-step ‘How to order / How to book’** that fits on the card.”
- “Ensure prices and options are scannable on a phone screen.”

What you'll get

- A tightened, **phone-readable card** with clear sections (e.g., **Cortes | Color | Tratamientos** or **Tacos | Quesadillas | Bebidas | Postres**) and a short **3-step** flow (e.g., “Elige base → Elige extra → Paga/Reserva”).
- Suggestions for tiny icons (scissors, comb, foil, tortilla, comal, soda cup) that speed visual recognition.
- Wording that fits **≤7 words** per line so customers decide faster at the counter or in a chair.

Why this works

- Real-time editing reveals what **new customers** don't yet know.
- Short lines and icons reduce language burden and anxiety.
- The model helps you **cut** without losing key information, and you validate on-screen as you go.

Follow-through

- Save the final card as a print-ready PDF and a phone-sized PNG.
- Capture a micro-guide: “≤7-word lines,” “3-step flow,” “icons per section,” so you maintain consistency when you add or change services later.

Talk like a pro: live question patterns that produce edits (not essays)

- “Rewrite this line to ≤10 words with the same meaning.”
- “Turn this paragraph into **three** bullets, each ≤12 words.”
- “If this section must be shown in **30 seconds**, what stays and what goes?”

- “Suggest a one-sentence transition from Slide 5 to Slide 6 that references the prior concept.”
- “Give me a quick diagram I can draw live—3 boxes or arrows max.”

Typical pitfalls and how to avoid them

- **Oversharing sensitive info:** Stick to public assets (decks, site pages) or fictionalized data.
- **Vague requests:** Ask for **constrained outputs** (“three bullets,” “≤12 words,” “3-step flow”).
- **Fixing the wrong thing:** Start with the **first screen** or the **first fold**—what customers see in 10–15 seconds is what drives outcomes.
- **Not validating edits:** After you apply a change, ask the model, “Does this now meet the constraints we set?” (word cap, number of bullets, clarity rule).

From Live to finalized assets

- **Decks:** Export a PDF for classroom distribution and a Slides version for future edits.
- **Menu/service cards:** Export print-ready and phone-sized versions; keep a master file with styles and icons.
- **Playbook snippets:** Move effective micro-rules (e.g., “≤7-word lines,” “3-step flow,” “≤12-word bullets”) into your internal style guide so every new asset inherits the same clarity.

Media Generation in AI Studio (Images & Speech)

What it is

Media Generation within Google AI Studio encompasses two powerful capabilities: image creation (for backgrounds, icons, diagrams, and label mockups) and text-to-speech (single-speaker narration that can be downloaded as audio files). Used deliberately, these tools eliminate “asset bottlenecks” when you need clean visuals or listenable content to advance a lesson or product page. The goal here is not flashy art; it’s speed, consistency, and clarity for instruction and everyday operations.

Access & setup (first time)

1. Open Google AI Studio and sign in.
2. From the left navigation, choose **Generate media** (for images) or **Speech generation** (for audio).

3. For images, select the image model, set size, and any optional parameters the UI exposes (we'll constrain via prompt).
4. For speech, choose a **single speaker** to keep training material simple; pick voice/tone if available.
5. Use clear, "production-minded" prompts: name dimensions, negative space, export format, and usage (e.g., "leave a calm center area for overlay text").

Control concepts that matter

- **Dimensions & legibility:** Always declare width/height and ask for **negative space** for where text will go.
- **Style tokens:** "High-contrast," "flat," "monochrome line icons," "consistent stroke weight," "calm center," "no people."
- **File outputs:** PNG for images (transparent icons where needed), MP3 for audio.
- **Tone restraint:** For speech, short sentences, clean vocabulary, and timed pauses improve naturalness.

Use Case A — Business Academy: "Module Visual Pack" in One Hour

Objective

Create a cohesive visual kit for a new teaching module, including six abstract slide backgrounds, twelve topic icons, and a one-minute intro narration, so instructors can rapidly assemble slides without hunting for stock imagery.

Workflow (step-by-step)

1. Background set (6)

- o Prompt:
"Produce **six** abstract slide backgrounds at **1920×1080** for a small-business module on **Cash Flow**—style: **clean, high-contrast, geometric, no text, no people**. **Leave the central 40% visually calm** for overlays. Vary layouts slightly but keep a **consistent palette** (blue/teal/gray). Output: PNG."
- o Download, name files cashflow-bg-01.png ... -06.png.

2. Topic icons (12)

- o Prompt:
"Generate **twelve monochrome line icons** at **512×512** with **consistent stroke weight**, transparent background: **Revenue, Expenses, Receivables, Payables, Float, Forecast, Break-Even, Inventory, Pricing, Margin, Collections, Controls**."

- o Export as PNG—store as icon-revenue.png, etc.

3. Intro narration (single-speaker)

- o Script template (you supply text): 120–160 words, short sentences, minimal clauses.
- o Prompt:
“Convert the following 150-word module intro to **single-speaker** audio. **Calm teacher** tone. Add a **0.7 s pause** between sentences. Output **MP3, 160 kbps.**”
- o Download cashflow-intro-v1.mp3.

4. Assemble

- o Place a background on the title slide, an icon for each section, and the narration file in your slide master or in the LMS asset folder.

Quality checks

- Backgrounds: text overlays are readable at 60–70% zoom on a laptop.
- Icons: stroke weight uniform, no fill surprises, recognizable at 48 px.
- Audio: no tongue-twister phrases; opening sentence states outcome (“By the end, you can sketch a 13-week cash-flow view.”).

Pitfalls & fixes

- **Busy centers** → add “calm center” or “light vignette edges” to prompt.
- **Icon inconsistency** → explicitly say “consistent stroke,” “uniform padding,” “transparent background.”
- **Robotic audio** → shorten sentences; remove parentheticals; avoid lists of three with commas (prefer “first... then... finally...”).

Use Case B — Micro-Business: “WhatsApp Support Pack” (Images + 30–40s Audio)

Objective

Ensure every customer receives the exact high-quality instructions after a purchase or service—without needing to refer to complete manuals. You’ll generate simple instruction cards (visual backgrounds for text) and a short WhatsApp-friendly voice note.

Workflow (step-by-step)

1. Instruction background (IG 1080×1350)

- o Prompt:
“Create **four simple abstract backgrounds** sized **1080×1350** for product care cards. **No people, no text**; keep a **high-contrast design** with a **calm central area** for overlay instructions. Use brand-friendly warm colors (terracotta/cream/green).”
- o Export PNG, name care-bg-01...04.png.

2. Overlay the text (in Slides/Canva):

- o **Tacos al pastor reheat** example: “Preheat 350°F (175°C) – 5 min | Heat sealed tray 10–12 min | Rest 2 min | Add salsa after heating.”
- o **Hair color care**: “Wait 48h before shampoo | Use color-safe products | Cool water rinse | Heat protectant | Avoid pools 5 days.”

3. 30–40s audio tip (single speaker)

- o Prompt:
“Turn this 120-word care guide into a **friendly 30–40 second audio tip**. Keep steps short. End with: ‘For more tips, message us anytime.’ Output MP3.”
- o Send the file via WhatsApp with each order/appointment follow-up.

4. Label & reuse

- o Keep a folder per product/service: al-pastor/bg/audio/overlay-card.png.
- o Reuse the same background style across SKUs to build brand recognition.

Quality checks

- Text is legible on a 5.5-inch phone.
- Steps keep **one action per line**.
- Audio ends with a warm, clear invitation to ask questions.

Pitfalls & fixes

- **Overcrowding** → cut to 4–6 lines; link to “full care” on site if needed.
- **Ambiguous steps** → include temperature/time or specific product types (“color-safe shampoo”).

- **Brand drift** → document the two or three colors and a font you'll always use on these cards.

Production checklist (applies to both use cases)

- File naming with versioning (-v1, -v2).
- Alt text for any image posted online.
- Export **phone-size** versions (1080 wide) for social/WhatsApp, keep originals for slides.
- For audio, store the transcript with the file for accessibility and later multilingual work.

Chapter 12: Google NotebookLM

1) What NotebookLM Is

NotebookLM is Google's **source-grounded research notebook**. You load it with your materials (Docs, Slides, PDFs, web/YouTube links, audio files, pasted text), and it answers questions, drafts study guides, and creates **Audio/Video Overviews** that explain your sources in natural language. It's designed to help you learn **from your own content**, with citations and guardrails.

Great for: turning a folder of policies or lessons into explainers, quizzes, and overviews—especially for onboarding, training, and classroom prep.

NotebookLM, as of late 2025, retains **chat history** in the chat panel so that learners can continue the same thread later; users can also delete chat history from the menu.

NotebookLM has an official **mobile app** (iOS/Android). It's ideal for "on-the-go" learning—creating/accessing notebooks, asking questions, and listening to Audio Overviews (including background/offline listening, per Google's app listing).

2) Key Features

A) Add & organize sources

- **From Google Drive:** Google **Docs** and **Slides** (recommended), plus **PDF** and **Text/Markdown**.

- **From the web:** URLs (articles, docs) and YouTube links (uses transcripts).
- **From files:** Audio files are auto-transcribed on upload; images inside PDFs/Slides are read when text is present.
- NotebookLM's supported sources, as of late 2025, include **Google Sheets** and **images**, plus **Microsoft Word (.docx)**.

Implications for Entrepreneurs & Ministries:

- **Sheets:** add structured program data (attendance, outcomes, simple finance tables) so NotebookLM can help you summarize trends and cite the rows/tables you provide.
- **Images:** you can drop in screenshots/infographics and pull explanations into study aids (still do a human verification pass).
- **.docx:** reduces friction when partners send Word files instead of Google Docs.

Paste-ready update to your “Add & organize sources” bullets:

- Add: “**Google Sheets** (token-limited) and **Images** (common formats) are supported sources.”
- Add: “**Microsoft Word (.docx)** is supported (no need to convert first).”

Steps: New notebook → **Add** → choose **Drive / URL / Upload** → give each source a clear name (e.g., SafeguardingPolicy_v2025, Module1_Slides_v2025). Organize content into sections (Pricing, Safeguarding, Donor Communications) so learners can easily navigate.

Tips: Prefer **text-based** PDFs over scans (run OCR if needed). Keep each source **focused** (one policy/module per file) to improve the quality of answers and overviews.

B) Discover Sources (built-in web finder)

Use **Discover Sources** to have NotebookLM suggest relevant web materials for a topic, and then add them to your notebook with a single **click**.

Starter prompt: “Discover sources on **micro-enterprise pricing** in Central America. Prioritize **official**, **primary**, and **recent** materials; exclude opinion blogs.”

Labeling tip: Prefix with WEB_YYYY-MM_... so you can review/prune later.

C) Audio Overviews — create & customize

Click **Audio Overview** → pick **Sources/sections** to include.

Choose **language** and **length** (short vs deeper walk-through).

Add **audience** and **study goals**: “for first-time volunteers,” “entrepreneur bootcamp,” “pastors leading small groups.”

Save multiple overviews per notebook (e.g., “Safeguarding — Staff” vs. “Safeguarding — Volunteers”).

Template: “Create an **Audio Overview** limited to **Policy sections 1–3 + Brand Guide p.4–6**. Use a clear teaching tone when working with **new BBA volunteers**. Include three reflection questions.”

D) Video Overviews — create & customize

Video Overviews create an explainer video narrated by AI that pulls **images, diagrams, quotes, and numbers** directly from your notebook sources.

Click **Video Overview** → select **sources/sections** → set **aspect** (16:9 / 9:16 / 1:1) and **language**.

Add a **style note** (clean classroom slides, high contrast, no logos) and scene notes (quotes, statistics, diagrams), and indicate where to **leave space** for captions.

Generate → review → download; add **Voice/music/captions** in Slides/Canva/CapCut if needed.

Template: “Generate a **Video Overview** from **Module 1 Slides + Pricing Handout** in **Spanish** (16:9). Emphasize formulas (COGS, margin) and add 3 scene breaks with example calculations. Leave the bottom 20% clean for captions.”

E) Deep Research inside NotebookLM (two research modes)

NotebookLM now supports a **Deep Research** workflow that can generate a **research plan**, gather **high-quality sources**, and produce a **structured report** grounded in the sources it collects and/or you add. This is useful when you don’t yet have a strong source pack, or when you need to widen the evidence base before building study aids.

How to use (practical):

1. In a notebook, open **Deep Research**.
2. Choose **Fast** (quick scan) or **Deep** (broader, more thorough).
3. Review the **research plan** before it runs; edit the plan to constrain scope (region, dates, “primary sources only,” etc.).
4. When it finishes, **promote the best sources** into your “core canon” and archive the rest (see Maintenance Playbook).

AI Manual for Entrepreneurs & M...

Rule of thumb:

- Use **Deep Research** when you're building a new module/policy pack, and you need credible, up-to-date sources.
- Then switch back to the standard NotebookLM flow (Overviews → quizzes → study guides) once your source pack is stable.
AI Manual for Entrepreneurs & Ministries.

F) Study aids & organization

What it does: NotebookLM can transform your sources into study artifacts—such as study guides, outlines, quizzes, glossaries, and mind maps—and keep everything organized with **sections**, tags, and quick exports.

1) Study Guides (step-by-step)

Select the sources/sections you want covered.

Ask: “Create a **study guide** with: (a) a 1-page explainer, (b) key terms with definitions, (c) 5 takeaways, (d) 3 reflection questions. **Cite sections.**”

Refine tone and level (beginner/advanced) and **language** on request.

Save as a note inside a **Section**. (e.g., Pricing → ‘Module 1 Study Guide’), or export to Docs/Slides.

Tips: Keep guides concise (≤1 page per subtopic). Add a short **Teach-back** prompt at the end.

2) Outlines & Note Cards

Outlines: “Make a **lesson outline** with headings → bullets → examples; tag each heading with the source page.”

Note cards: “Create **10 note cards** (term → definition → short example). Output as a table.”

Multilingual: “Provide EN → other language note cards.”

3) Quizzes & Exams

Generation: “Create **12 questions**: 4 recall, 4 application, four scenario-based. **Hide answers** until I attempt; then show the correct answer with **citation**.”

Formats: MCQ (1 correct, three plausible distractors), short answer, case vignette.

Rubrics: “Add a **rubric (0–3)** for short answers: accuracy, clarity, use of terms, actionable next step.”

Answer key: “Export an **answer key** for instructors only.”

Best practices: Mix Bloom levels; include at least one **scenario** that requires quoting policy text; keep MCQ stems short and distractors credible.

4) Glossaries & Concept Maps

Glossary: “Extract a **glossary** of key terms with plain-English definitions and **source citations**.”

Concept relationships: “List **concept relationships** (A depends on B; C contradicts D). Output as bullet pairs for a slide.”

5) Mind Maps & Concept Graphs (visual)

Mind Maps: one-click visual summary of your sources; excellent for kickoffs, sermon/lesson planning, and policy briefings. Generate on desktop; export/share the outline or screenshot as needed.

How to generate: - “Build a **mind map** of **Module 2: Pricing** from our sources. Use: central node = ‘Pricing’; first-level = ‘Costs’, ‘Margins’, ‘Value’, ‘Competitors’, ‘Customer’. Under each, add 3–5 child nodes with **page refs**.” - “Output the mind map as an **outline** (levels = nodes) **and as Mermaid** syntax so I can paste into draw.io/Mermaid-compatible tools.”

Formats you can export: - **Indented outline** → paste into Slides as a structured list (one level per indent). - **Mermaid mindmap** code → render in compatible tools; or paste into a diagramming app. - **Node table** → columns: Parent • Node • Source • Note → quick to sort/filter.

Tips: Keep to 5–7 top-level nodes; cap depth at three levels for readability; color-code nodes by source when you render.

6) Customize Audio & Video Overviews (advanced)

Audio Overviews - Voice & tone: “Use a warm teacher voice; pace for non-native speakers.” - **Audience & level:** “Aim at first-time volunteers/entrepreneurs with basic numeracy.” - **Scope control:** “Include only **Sections 1–3**; skip legal appendix.” - **Signposts:** “Insert **section markers** (‘Now we’ll cover...’) and **reflection pauses** with 1–2 questions.” - **Localization:** “Produce in **ES**; add a short **HT** summary at the end.”

Video Overviews - Aspect & captions space: “16:9 with **bottom 20% clear** for captions.” - **Style:** “Clean classroom slides; high-contrast palette; no logos.” - **Beat sheet:** “3 scenes × 6s each; scene 1 definitions; scene 2 example calculation; scene 3 pitfalls + next step.” - **Pull-quotes & figures:** “Show two quotes from policy with page refs; 1 chart from metrics.” - **Accessibility:** “Large headings, max 8–12 words per line; avoid busy backgrounds.”

Quality pass before export: verify citations, reading level, and that scenes align to your **learning outcomes**.

7) Organize with Sections, Tags & Versions

Create **Sections** for each course/policy (e.g., Safeguarding, Pricing, Donor Comms).

Store artifacts under the matching section and prefix with a date: 2025-09 Study Guide — Module 1.

Use **tags** in titles (e.g., [ES], [HT], [Quiz], [Guide], [MindMap]) for quick filtering.

Add a **Change Log** note at the top of each section when sources or policies update.

8) Export & Share

Export study guides/quizzes to **Google Docs/Slides** for formatting, and mind-map outlines to Slides or Mermaid code to a diagram tool.

For learners, generate **Audio/Video Overviews** from the same section.

For class delivery, paste quiz questions into **your LMS or Forms**.

9) Quality & ethics guardrails

Always **cite sources** for non-obvious claims.

Keep questions **original**; avoid copying copyrighted test banks.

Respect privacy: do not include real names/PII in scenarios.

Provide accessible language (short sentences, high-contrast slides on export) and bilingual options.

F) Enterprise/org sharing + analytics

For organizations, NotebookLM Enterprise emphasizes **shared notebooks, usage analytics**, and enterprise-grade controls.

Paste-ready line under “Team/Enterprise options”:

- “Enterprise: shared notebooks across users, sharing controls, and usage analytics for adoption/engagement measurement.”

3) Benefits

Accuracy by design — It stays anchored to your uploaded sources; responses link back for easy verification.

Faster teaching prep — Turn lesson packets and policies into explainers, **Audio/Video Overviews**, quizzes, and glossaries.

Multilingual reach — Produce overviews and summaries in multiple languages to support learners of Spanish and Haitian Creole (HT, if needed).

Low-lift onboarding — New staff/volunteers can listen to source-based audio/video briefings before a class or outreach.

4) Setup — 5 Minute Quick Start

Create a notebook and name it (e.g., BBA_Entrepreneurship_Module_1).

Add sources: Docs/Slides from Drive, PDFs/Text, web pages, YouTube links, audio, or paste text.

Click "Audio Overview" or "Video Overview" to generate a guided explainer.

Ask questions: "What are the three most important policies we teach? Cite source sections."

Export/save notes; organize sections by topic (e.g., Pricing, Safeguarding, Donor Etiquette).

5) Use Cases

Curriculum prep — Upload a module's slides + reading and generate an **Audio Overview**; ask for a 10-question quiz and a glossary.

Policy Training — Load the safeguarding policy, create a **Video Overview**, and a Q&A sheet citing exact sections.

Donor briefing pack — Add last quarter's reports + highlights; produce a narrated overview and a 1-page summary.

Volunteer onboarding — Drop the brand guide + communication SOP; generate a guided explainer and flashcard prompts.

BBA Prompt Kit for NotebookLM

"Create a **Study Guide** for **Module 2 Pricing** with: 1-page explainer, 10 quiz questions, eight flashcards, and three reflection prompts. Cite sections."

"Generate an **Audio Overview** of **Safeguarding** for **pastors** (ES). Include scenario Q&A and where to find the escalation steps."

"Produce a **Video Overview** of **Donor Communications** (EN, 16:9). Pull two quotes from the brand guide and one stat from the metrics report. Leave the bottom 20% empty."

“Discover sources for **street-market entrepreneurship** in Nicaragua/Honduras (last 18 months). Prefer government/NGO/academic sites.”

6) NotebookLM vs. ChatGPT — What’s the closest feature?

Closest ChatGPT feature: Projects - Similarity: Both keep a **workspace of files + chats/instructions** and let you ask questions across those files. - **Differences:** - **NotebookLM** is **strictly grounded** to your notebook sources and built around **study/overview** workflows (audio/video explainers). - **ChatGPT Projects** support **creation** and multi-tool work (Deep Research, image/code, slide outlines, custom agents), and can blend your files with the **public web** when needed.

Rule of thumb: Use NotebookLM when you want explanations directly from your materials (teaching, training, guided study, or multilingual overviews) with tight source grounding. - Use **ChatGPT Projects** for **writing and production** (briefs, slides, bilingual handouts), **web-backed research**, and integrating images/code/agents in one place.

7) Limits & Considerations

Source caps: Per-item and per-notebook limits apply; very large PDFs or multiple sources may require trimming.

YouTube/Audio: Based on available transcripts; audio is transcribed first.

Output style: Overviews target clarity/teaching, not polished design—finish in Slides/Canva as needed.

Privacy: Workspace/Enterprise options offer stronger controls; avoid placing raw PII in consumer contexts.

Keep Your Notebook Fresh — Maintenance Playbook

Weekly 1) Open notebook → **Sources** → any item marked “updated in Drive” → click **Re-sync**. 2) Archive older versions (suffix _archived_YYYY-MM). 3) Run **Discover Sources** for 2–3 fresh, reputable links; tag with WEB_YYYY-MM.

Monthly 1) Review **sections** for gaps/overlap. 2) Re-generate **Audio/Video Overviews** if core sources changed. 3) Add a **Change Log** note: what changed and why it matters.

Quarterly 1) Purge stale links; keep the top 10–20 **core** sources. 2) Update naming to vYYYY and ensure brand guide/policies are current. 3) Share a **What’s New** audio/video for staff and volunteers.

Freshness tip: Maintain a single staging document that compiles key excerpts, update it, and resync it with a single source, not multiple sources.

8) Side-by-Side — Quick Comparison

Task	NotebookLM	ChatGPT (Projects/Deep Research)
Learn from my materials.	Strong: source-grounded Q&A; Audio/Video Overviews	Strong: Q&A; no built-in audio/video explainers
Create deliverables	Introductory notes & guides; export to Docs/Slides	Complete drafting: briefs, slides, bilingual handouts, images, code
Citations	Always from your sources	From your files + (optionally) web with Deep Research
Web research	Not the focus	First-class (Deep Research)
Multilingual explainers	Built-in Audio/Video Overviews	Translate outputs; no native audio/video explainers
Enterprise controls	Plus/Enterprise tiers with org sharing/analytics	Business/Enterprise tiers, Projects, Agents, and admin controls

Chapter 13: App Builders

What are they?

App Builder lets you **describe** a small tool in plain English and have the system **plan**, **code**, and present a working UI within minutes. It's perfect for operational utilities you wish existed—comparators, brief generators, checklist makers, and classroom planners—without hiring a developer for a first draft. Think “minimum viable internal tool.”

Access & setup

1. In AI Studio, open the "**Build an app** experience" option (from a template or from scratch).
2. Describe inputs and outputs clearly (fields, sections, export buttons).
3. The system generates a prototype UI; you test with **dummy** data; iterate.

4. Once stable, you can export or rebuild the same logic in your trusted stack (e.g., Apps Script or a no-code platform) for production.

Spec writing tips

- Declare **inputs** (text boxes, dropdowns, toggles).
- Specify **outputs** and their order (sections, bullet caps, tables).
- Include **export actions**: “Copy to clipboard,” “Export CSV,” “Export to Google Doc.”
- Limit scope: one job per app.

Use Case A — Business Academy: “Workshop Agenda Builder”

Objective

Provide instructors with a one-click method to transform a topic and time budget into a structured teaching plan that includes objectives, timed segments, discussion prompts, and a practice task with a corresponding rubric.

Spec to paste

“Build a web app that takes **Topic** (text) and **Duration** (dropdown: 60/90/120 min) and outputs **five** sections:

1. **Lesson Overview** (≤120 words)
2. **Learning Objectives** (6 bullets)
3. **Timed Agenda** with **eight segments** (minutes per segment)
4. **Discussion Questions** (6 items)
5. **Practice Task** with a **simple rubric** (3 criteria, three levels).
Buttons: **Regenerate, Copy All, Export to Google Doc.** Keep text plain and concise.”

How to pilot

- Generate three topics (Pricing, Cash Flow, Market Fit).
- Check the agenda timing; trim if it overruns.
- Export to Docs and drop into your slide template.

Stabilization moves

- If outputs drift verbose, append: “If any section exceeds its cap, **truncate cleanly** without losing clarity.”

- If objectives are vague, require measurable verbs (“calculate, compare, identify, decide, outline, apply”).

From prototype to production

- If instructors love it, rebuild it in Apps Script for single sign-on, or embed it in your LMS as a simple tool with a Doc export.

Use Case B — Micro-Business: “Service Scope Helper”

Objective

Prevent scope misunderstandings for services (cleaning, basic landscaping, color touch-up). Output a clear “What’s included,” “What’s not,” “Prep checklist,” and “Who it’s for”—short enough to read on the phone while booking.

Spec to paste

“Create an app with **six inputs**: Service Type, Context (home size or hair length/type), Preferred Time Window, Add-ons (checkboxes), Special Notes, Location Type (apartment/house/salon).

Output **four** sections:

- **What’s Included** (6 bullets, ≤ 12 words each)
- **What’s Not Included** (4 bullets)
- **Prep Checklist** (6 bullets)
- **Who It’s For** (2 lines).
Add a **Bilingual toggle** to produce an **EN/ES** summary. Buttons: **Copy to Clipboard, Export to Google Doc.**”

Pilot

- Run with three typical scenarios; confirm bullets fit your real-world process.
- Standardize language across services; align any add-on terms with the pricing sheet.

Production

- Publish internally; staff read it on the phone during bookings.
- Copy the text straight into WhatsApp confirmations.

Pitfalls & fixes

- **Too much text** → tighten caps and add “ ≤ 12 words per bullet.”

- **Jargon** → require plain words (“deep clean kitchen” not “sanitation protocol”).
- **Missing exports** → confirm Doc export works; otherwise, keep “Copy to Clipboard.”

App QA checklist (both use cases)

- Inputs reject empty values; dropdown defaults make sense.
- Outputs render on mobile without wrapping weirdly.
- “Copy All” includes headings; Doc export keeps bullet styles.
- Version stamp visible at bottom (e.g., “ScopeHelper v1.1 — 2025-09”).

Build an App in Google AI Studio — Step by Step

Prep (2–5 minutes)

- Decide the **one job** your app will do (e.g., “turn a topic + duration into a workshop agenda,” or “turn six answers into a scope summary”).
- Write your **inputs** (fields the user must provide) and **outputs** (sections, bullet caps, tables). Keep it small for v1.
- Gather a **few fake test inputs** (never use real PII in the prototype).

1) Open AI Studio and start an app

1. Go to **Google AI Studio** and sign in.
2. In the left nav, choose **Build an app** (wording may be “App Builder” or similar).
3. Pick **Start from scratch** (or choose a template if one matches your goal).

2) Describe your app (the spec)

4. In the “Describe your app” area, paste a **clear spec** that includes:
 - o **Inputs**: names, types (text, dropdown, checkbox), hints/examples.
 - o **Outputs**: exact sections/fields, order, bullet/word caps, table columns.
 - o **Buttons/Exports**: Copy to Clipboard, Export to Google Doc/CSV.
 - o **Style constraints**: plain language, 7th-grade reading level, no fluff.

Use the “Spec Template” below—paste it as-is and tweak.

3) Generate the first version

5. Click **Generate** (or similar). AI Studio will plan and create your app's UI.
6. Wait for the preview canvas to load with fields/sections.

4) Test with dummy data

7. Fill in **every input** using **fake, realistic values**.
8. Click the app's **Run/Generate** button to produce outputs.
9. Skim outputs for:
 - o Correct **section order** and **caps** (e.g., "6 bullets," " \leq 120 words").
 - o **Clarity** and **phone-readability** (short, plain sentences).
 - o Working **export buttons** (copy, Doc/CSV export).

5) Iterate quickly

10. If something's off, click **Edit app** (or "Refine") and **tighten** your spec:
 - o Add: "Return sections **exactly** in this order."
 - o Add per-section caps: "Each bullet \leq 12 words," "Overview \leq 120 words."
 - o Add: "If over cap, **truncate cleanly** without losing meaning."
11. Regenerate and re-test until the app is **boringly consistent**.

6) Add export actions (necessary for real use)

12. Ensure your spec includes:
 - o **Copy All** button (grabs headings + content).
 - o **Export to Google Doc** (full output).
 - o Optional **Export CSV** (for tables/comparison apps).
13. Re-generate and verify each export actually works.

7) Lock input sanity

14. Add simple **input rules** in your spec:
 - o Required fields cannot be empty.

- o Dropdowns for standard options (e.g., Duration: 60/90/120).
- o Short hints: “Example: ‘Cash Flow Basics’” to guide users.

8) Mobile check + formatting pass

15. Resize the preview (or open on a phone) and confirm:

- o Inputs are tappable, text readable, and buttons visible.
- o Bullets don’t wrap into messy lines.
- o Tables stack or remain legible on small screens.

9) Version stamp & share

16. Add a **footer note** in your spec: “v1.0 — YYYY-MM-DD.”

17. Click **Share** (or copy app link) and send to a teammate for a 5-minute test.

18. Capture their feedback; adjust any confusing labels or long sections.

10) Roll out (lightweight SOP)

19. Write a 1-page **How to use**: where to click, what to enter, how to export.

20. Store the app link + SOP in your team’s hub (Wiki/Docs). When the app becomes critical, plan a **Workspace-native** rebuild (using Apps Script/Sheets/Docs) for governance purposes.

Copy-Paste Spec Template (start here)

Use this block as your **first spec**. Paste into AI Studio’s “Describe your app” box and customize the bracketed items.

Goal: Build a single-purpose app that transforms user inputs into a structured, copy-ready output.

Inputs:

- Topic (text) — hint: “Example: Cash Flow Basics.”
- Duration (dropdown: 60, 90, 120 minutes)
- Audience (dropdown: Beginners, Intermediate)
- Notes (multiline text; optional; short phrases)

Outputs (return sections EXACTLY in this order; obey caps):

- 1) Lesson Overview — ≤120 words, plain language

- 2) Learning Objectives — 6 bullets, each ≤12 words, verb-led (identify, calculate...)
- 3) Timed Agenda — 8 segments with minutes per segment; total must equal Duration
- 4) Discussion Questions — 6 items, one line each, ≤12 words
- 5) Practice Task + Rubric — 3 criteria × 3 levels (Basic/Good/Great), concise

Constraints:

- 7th-grade reading level, no jargon, no fluff.
- If any section exceeds its cap, TRUNCATE CLEANLY (do not spill).
- Return only the sections above; no extra text before/after.

Buttons/Exports:

- Generate
- Copy All (include headings)
- Export to Google Doc
- (Optional) Export CSV for the Timed Agenda

Quality:

- Phone-friendly lines; short sentences; consistent headings.
- If input is missing, state [Assumption: ...] and proceed.

Footer:

- Version stamp: v1.0 — [YYYY-MM-DD]

Example Spec #1 — Workshop Agenda Builder (for instructors)

Paste this into the spec box and run:

Create a web app that takes these inputs:

- Topic (text) — hint: “Example: Pricing Basic.s.”
- Duration (dropdown: 60, 90, 120 minutes)
- Audience (dropdown: Beginners, Intermediate)
- Notes (textarea, optional; short phrases only)

Produce five sections in this EXACT order with these caps:

- 1) Lesson Overview — ≤120 words, plain language
- 2) Learning Objectives — 6 bullets, each ≤12 words, verb-led (identify, calculate...)
- 3) Timed Agenda — 8 segments; show minutes per segment; total must equal Duration
- 4) Discussion Questions — 6 items; one line each; ≤12 words
- 5) Practice Task + Rubric — 3 criteria × 3 levels (Basic/Good/Great), concise

Constraints:

- 7th-grade reading level; avoid jargon and fluff.
- If any section exceeds its cap, TRUNCATE CLEANLY (do not spill).
- Return only the five sections above.

Buttons/Exports:

- Generate
- Copy All (include headings)
- Export to Google Doc

Mobile:

- Ensure bullets don't wrap poorly; keep lines short.

Footer:

- Version stamp: Workshop Agenda Builder v1.0 — [YYYY-MM-DD]

Test values (dummy):

- Topic: "Pricing Basics"
- Duration: 90
- Audience: Beginners
- Notes: "Include PCM; highlight common mistakes."

Run → check section order/length → export to Doc.

Opal (Describe-to-Build Mini Workflows/Apps)

What it is

Opal is a **natural-language workflow/app builder**. You describe what you want—inputs, steps, outputs—and it assembles a simple tool that can accept text, run a

sequence (summarize → transform → export), and return a **shareable link** for teammates. Opal shines when you need **repeatable, single-purpose** flows that live in the Google ecosystem (e.g., read input → produce structured sections → save to Google Doc).

Access & setup

1. Open Opal and click **Create new**.
2. Type a plain-English description of the workflow/app (we provide specs below).
3. Test with dummy text; adjust step order, section names, and export behavior.
4. Copy the share link so teammates can use it without learning the prompt.

Design tips

- **One job per workflow:** module pack assembler, service message maker, scope helper.
- **Explicit steps:** “User pastes outline → generate sections → save to Doc → return link.”
- **Output constraints:** exact sections, bullet counts, word caps, and reading level.
- **Exports:** always include Copy/Doc export—links are your friend.

Use Case A — Business Academy: “Module Pack Assembler”

Objective

Turn a pasted outline into a **complete module pack** with Overview, Objectives, Timed Agenda, Discussion Questions, and one Practice Task—with **two exports**: a Doc for instructors and a second Doc with an **8-question quiz + answer key**.

Opal build text (paste as your spec)

“Create a workflow that:

1. Prompts the user to **paste a module outline**.
2. Generates a **module pack** with sections in this exact order: **Overview ≤150 words, 6 Learning Objectives, 8-segment Timed Agenda, 6 Discussion Questions, 1 Practice Task with rubric**: plain language, 7th-grade reading level.
3. **Saves** the module pack to a **new Google Doc** and returns its link.
4. Generates a separate **8-question quiz** (a mix of multiple-choice and short-answer) with an **answer key**.
5. **Saves** the quiz to a second **Google Doc** and returns its link.

6. Shows both links on completion.”

Pilot & tighten

- Paste three outlines; inspect section order and caps.
- If objectives are weak, add “start each objective with a measurable verb.”
- If agenda times don’t sum cleanly, add “include minutes per segment; total must match selected Duration if provided.”

Rollout

- Share the link with instructors; keep a short Loom video explaining how to paste outlines and find their Docs after generation.

Use Case B — Micro-Business: “Order Follow-Up Message Maker”

Objective

After each pickup or appointment, send a consistent, polite, **and brief follow-up with care instructions via SMS/WhatsApp in both English and Spanish**—without having to rewrite it every time.

Opal build text (paste as your spec)

“Build a workflow that:

1. Asks for **Item/Service name**, **Pickup/Appointment time**, and **Care bullets** (up to 6 short lines).
2. Generates a **friendly follow-up message** ≤ 400 **characters** in **English** with a greeting, the item/service, care bullets condensed into phrases, and an invitation to ask questions.
3. Generates a **Spanish version** with the same meaning.
4. Shows both messages with **Copy** buttons.”

Pilot & tighten

- Test with 3–4 scenarios (taco tray, hair color, nail set).
- Ensure tone stays warm, not pushy; adjust constraint to “no marketing language.”

Rollout

- Save Opal link to phone bookmarks; staff paste outputs directly into WhatsApp.
- Keep a short note defining “ ≤ 400 characters” to prevent multi-message splits.

Pitfalls & fixes

- **Overflow** → enforce character cap and “compress care bullets into phrases.”
- **Tone drift** → add “no exclamation unless necessary,” “plain words,” “no slang.”
- **Missing Spanish nuance** → if needed, ask Opal to generate ES first and mirror to EN.

Firebase Studio

What it is

Firebase Studio provides an **AI-assisted workspace** to scaffold front-end sites or simple apps. You describe the site, sections, and behaviors; the system generates a working draft you can preview, tweak (“vibe-code” with natural language), and publish behind your **own key/hosting**. It’s ideal for **microsites** (cohort hubs, class materials) and **single-page business sites** you need this week—not next quarter.

Access & setup

1. Open **Firebase Studio** and click **Try**.
2. Describe your site: pages/sections, layout, style, and must-have components (accordion syllabus, downloads list, FAQ).
3. Let the generator produce a **first pass**.
4. Use natural-language edits: “Make headings larger,” “Add a ‘Reserve by WhatsApp’ button,” “Move FAQ above Contact.”
5. Add your **API key** or hosting config if you intend to publish.

Design principles

- **Single purpose** per microsite: cohort hub, course landing, or services one-pager.
- **Readable on phone first**: check spacing, font size, tap targets.
- **Export plan**: be ready to download and host under your domain.
- **Privacy**: never embed private documents publicly; link to LMS or gated Drive.

Use Case A — Business Academy: Cohort Microsite Generator

Objective

Generate a clean microsite for “Cohort 12” that centralizes the syllabus, calendar, materials, office hours, and FAQ, so instructors can stop copying and pasting links across emails.

Spec to paste

“Create a responsive microsite for **Business Academy Cohort 12** with sections:

- **Hero**: title, 2-line description, quick links (Syllabus, Calendar, Materials)
- **Syllabus**: accordion with module names and one-line summaries
- **Calendar**: list next eight sessions with date/time and ‘Add to Calendar’ links (placeholder)
- **Materials**: download list with file-type icons
- **Office Hours**: times, time-zone note, ‘Join’ placeholder
- **FAQ**: 8 questions/answers, short and scannable
- **Contact**: email form (placeholder)
Style: clean sans-serif, **high contrast**, generous spacing, **print-friendly** stylesheet. Mobile first.”

Tighten with edits

- “Reduce FAQ answers to \leq 40 words.”
- “Make link buttons full-width on mobile.”
- “Add a breadcrumbs bar: Home / Cohort 12.”

Publish plan

- If you have hosting ready, deploy to a subdomain (e.g., cohort12.yourdomain.org).
- Otherwise, export and hand it to your webmaster.
- Keep live links to materials inside your LMS or a private Drive to prevent leakage.

QA checklist

- On mobile: all links tappable, no tiny text.
- Print preview: collapses accordions, shows full text.
- Accessibility: consistent heading order, alt text for icons (or mark as decorative).

Pitfalls & fixes

- **Over-stuffed hero** → keep to two lines; push details below.
- **Dead placeholders** → swap in real links as soon as you create them.

- **Materials sprawl** → group by module; limit to 5–7 items on the landing page; link to “More.”

Use Case B — Micro-Business: One-Page Service Site with WhatsApp CTA

Objective

Spin up a modern one-pager a local customer can read in 30–60 seconds with a **WhatsApp “Reserve”** button, simple 3-step “How we work,” and a clear comparison of three packages.

Spec to paste

“Create a **one-page responsive** site for **Limpieza Brillante CLT**. Sections:

- **Hero**: short headline, 2-line subhead, ‘**Reservar por WhatsApp**’ button
- **Cómo Trabajamos**: **3-step** cards with icons, each ≤ 12 words
- **Servicios**: cards for key services with short descriptions
- **Paquetes**: **3-column comparison** table (Básico / Intermedio / Pro) with five rows; no jargon
- **Preguntas Frecuentes**: 8 Q&A, ≤ 40 words each
- **Testimonios**: 3 placeholders with star icons (no faces)
- **Contacto**: WhatsApp button, phone, email

Style: warm palette, **big text, large tap targets, phone-first** spacing.”

Edits to request

- “Make the WhatsApp button sticky at the bottom on mobile.”
- “Add an announcement bar for ‘**Promo de la semana**’ (placeholder).”
- “Convert the comparison table to stacked cards on mobile.”

Publish plan

- Push live under a simple domain (or /servicios path).
- Connect the WhatsApp link to your business number with a prefilled “Hola, quiero reservar...” message.
- Keep the “Promo de la semana” current every Monday.

QA checklist

- Phone test: headline readable at a glance; CTA obvious.
- Packages table: no tiny text; benefits short and unambiguous.

- Load time under a few seconds on 4 G.

Pitfalls & fixes

- **Wall of text** → keep every sentence \leq 12 words; use bullets.
- **Confusing packages** → add a “Who it’s for” line under each; remove fluff.
- **CTA buried** → make WhatsApp sticky and full-width on mobile.

Hand-off & maintenance

- Export code and commit to a version control folder for future tweaks.
- Keep a one-page “Content Update SOP,” so staff can replace text/links monthly.
- Schedule a **quarterly** phone test: open the site on a budget Android and make sure everything is still tappable and fast.

Chapter 14: Gemini/Google Tools Integration

Google Sheets Integration

Gemini in Sheets adds an AI function that processes **cell ranges** to generate summaries, plans, rewrites, and structured tables—without leaving your spreadsheet. This is ideal when your raw material already exists in Sheets, such as survey comments, inconsistent product descriptions, idea pools, or SOP fragments from multiple contributors.

Access & usage

1. Open your Sheet, click into a cell where you want output.
2. Use the **Generate text with AI** function (or the current function name in your rollout) with two parts: a **prompt** and a **range**.
3. The function returns text; you can post-process using SPLIT, REGEX, or “Split text to columns.”

Prompt hygiene for Sheets

- Put your **prompt text** in a cell (e.g., H1) and reference it: =GENAI(\$H\$1, B2:B301).
- Constrain outputs (“ \leq 130 words,” “6 bullets,” “table with 5 columns”).
- Keep ranges realistic (don’t send thousands of rows if you only need 200).

Use Case A — Business Academy: “Cohort Feedback → Themes + 6-Step Action Plan.n”

Objective

Turn a semester’s worth of open comments into a high-signal, concise plan instructors can apply in the next cohort.

Data

- B2:B151 contains 150 short survey comments (anonymized).
- H1 contains your prompt string.

Prompt for H1

“Summarize dominant themes from the feedback and produce a **6-step improvement plan** with **owners** and **dates**. Be concise, avoid repetition, and point out any missing info we should collect next time.”

Formula

=GENAI(\$H\$1, B2:B151)

Post-processing

- If the output includes a list of themes and a plan in a single blob, copy it to a Doc for formatting, or split it by headings if the function returns consistent delimiters.
- Capture 2–3 representative quotes manually (keeps it human and avoids over-automation).

Quality checks

- Plan items reference **teachable** actions (e.g., “Add a worked cash-flow example,” not “Be clearer”).
- Owners play fundamental roles (Lead Instructor, TA, Admin) and dates fall within the next cohort cycle.

Pitfalls & fixes

- **Overlong output** → add “≤250 words total” to the prompt.
- **Blurred themes** → ask for “Theme → 2-sentence explanation → (Optional) quote marker.”
- **Mixed languages** → separate into EN/ES ranges or pre-tag rows.

Use Case B — Micro-Business: “Instagram Plan from Idea Pool (Time-Aware)”

Objective

Turn a list of post ideas + your real-time budget into a **7-day plan** you can actually execute.

Data

- D2:D60 ideas; E2:E60 platform tags; F2:F60 time available per day in minutes.

Prompt for H1

“Create a **7-day content plan** that **respects** the time in column F. For each day, return: **Platform, Post Type, Hook (≤12 words), Caption (≤80 words), Hashtags (≤6), Time Estimate**. Mix: **40% education, 40% story, 20% behind-the-scenes**.”

Formula

=GENAI(\$H\$1, D2:F60)

Post-processing

- Paste results into a “Calendar” tab; adjust dates; mark done/posted.
- If any day exceeds the available time, manually switch to a lighter post type (e.g., a photo instead of a reel).

Quality checks

- Hooks are readable in 2 seconds; captions fit on a phone without a “More...” link in most cases.
- Time estimates add up to your weekly budget.

Pitfalls & fixes

- **Generic hashtags** → specify “use niche hashtags common in [neighborhood/cuisine/service]” in the prompt.
- **Repeat ideas** → add “avoid duplicates; if similar, vary angle.”

Sheet hygiene pattern to teach your team

- Tab 1: **RAW** (never edited; just drop new responses here).
- Tab 2: **AI_OUT** (all GENAI formulas live here).
- Tab 3: **ACTION** (human-curated decisions, owners, dates).
- Version/date stamps at the top of the ACTION tab.

Ask with Gemini on YouTube (Chat, Chapters, Timestamps)

What it is

“Ask with Gemini” on YouTube lets you **chat with a video**, quickly generating **chapters with timestamps, key takeaways, and targeted answers** without manually scanning the entire runtime. For educators, this is perfect for turning long tutorials into teaching-ready outlines. For micro-business owners, it’s a rapid way to capture **step sequences, pitfalls, and moments to reread** from Spanish or English videos—so practice time is used well.

Access & setup

1. Open a YouTube video. Look below the player for the **Ask** (Gemini) prompt.
2. Start with specific requests: chapters with ranges, 3–5 key insights, exact timestamps for demos.
3. Ask follow-ups: “Show me the segment where they prep the ledger,” “Extract the steps only,” “What tools are used?”

Prompts that work

- “Create **chapters with timestamps** and a **5-takeaway** list.”
- “List **step-by-step** instructions with **timestamps**; mark **pitfalls** separately.”
- “Find the segment that shows **[feature]** on screen; give the time range and what to watch for.”

Use Case A — Business Academy: Turn a Long Tutorial into a Teaching Outline

Objective

You’ve found a comprehensive tutorial on accounting or POS systems. You want to extract a **chaptered outline, key demos, and classroom prompts** without rereading 45–60 minutes.

Workflow (step-by-step)

1. **Chapters first**
 - o Ask: “Create **chapters with start-end timestamps** for this video. Keep titles short and action-oriented.”
 - o Copy the list into your lesson plan.
2. **Key demos next**
 - o Ask: “Identify **three segments** where the presenter **does a live demo** (not slides). Give timestamps and what exactly is shown.”
 - o Mark those for in-class playback.

3. Class prompts

- o Ask: “Generate **six discussion prompts** students can answer after watching the demo clips. Mix *why*, *how*, and *what-if* questions.”
- o Paste into your worksheet or slides.

4. Quick-reference sheet

- o Ask: “Create a **one-page quick-reference** of the **hotkeys** and **menu paths** shown in this video, with timestamps for where each appears.”
- o Export to PDF and print.

Quality checks

- Chapter titles are scannable and correspond to fundamental transitions.
- Demo timestamps start a few seconds before the action so that students can orient.
- Discussion prompts reference visible steps (“After the **reconcile** step, why might...?”).

Pitfalls & fixes

- **Missing chapters:** If the video already has chapters, still ask Gemini to **revise** them for clarity and teaching language.
- **Too generic prompts:** Ask for “specific references to demo steps” and cap each question to one sentence.
- **Outdated UI:** Add a note to your plan if the interface looks different in your local install; record your own 30-sec correction clip if needed.

Production checklist

- Paste chapters and prompts into your slide or LMS notes.
- Create a “Watch this clip (mm: ss–mm: ss)” callout on the relevant slides.
- Keep a backup plan if the creator deletes/privates the video (export notes; consider screen-cap key 10–15 sec segments for fair educational use if permitted).

Use Case B — Micro-Business: Extract Steps & Pitfalls from a Spanish Tutorial

Objective

You’re learning a technique (e.g., tortilla press workflow, salsa roja prep, balayage

basics, edge-up techniques). You want timestamps for steps, **pitfalls to avoid**, and **two moments to rewatch**.

Workflow (step-by-step)

1. **Ask for steps with timestamps**
 - o “Summarize the **steps with timestamps**. Keep verbs imperative, ≤ 10 words each.”
2. **Ask for pitfalls**
 - o “List **pitfalls** mentioned or shown. For each, include a **fix** and timestamp.”
3. **Moments to rewatch**
 - o “Mark **two moments** worth **rewatching** for technique; include **why**.”
4. **Practice card**
 - o Copy steps and pitfalls to a one-page practice card; keep it near your station or stove.

Quality checks

- Step list fits on a phone screen.
- Pitfalls are specific (“over-working masa,” “too-hot comal,” “foil overlap too wide”).
- Rewatch moments, call out visible cues (“listen for sizzle level,” “note foil angle”).

Pitfalls & fixes

- **Creator speaks too fast:** Ask for “slower paraphrase of steps, ≤ 8 words each.”
- **Missing ingredient tools:** Ask “list tools/ingredients mentioned on screen with timestamps.”
- **Confusing regional terms:** Request “define unfamiliar terms in one phrase.”

Production checklist

- Save the chapter list as a note on your phone.
- Print the practice card and plastic-cover it for the workspace.
- Update the card after one week of practice with your own tweaks.

Speech (Multi-Speaker) for Dialog & Role-Play

What it is

Multi-speaker speech generation creates **two-voice** (or more) audio from a script. For the Business Academy, it's a fast way to produce **role-plays** (owner ↔ customer; manager ↔ staff) to rehearse difficult conversations. For micro-businesses, it delivers **booking-call practice** or **service-explanation** dialogs that owners can replay on the way to work.

Access & setup

1. In AI Studio, choose **Speech generation** and select **multi-speaker**.
2. Paste a concise, stage-direction-light script.
3. Pick neutral voices; avoid extremes.
4. Export as MP3; keep the transcript in the folder for accessibility.

Script writing rules

- **Short lines** (≤ 12 words), natural turn-taking.
- **Label speakers** clearly (CUSTOMER / OWNER).
- **One action per line**; no nested clauses.
- **End with a summary line or next step.**

Use Case A — Business Academy: Role-Play — “Refund or Store Credit?” (60–90s)

Objective

Give students an audio scenario to practice empathy and policy explanation, and to consider options. They listen once, then pair up to reenact with variations.

Script template (you will paste content; the generator produces audio)

- CUSTOMER: “Hi, my item arrived late. I needed it yesterday.”
- OWNER: “I’m sorry for the delay. I can help now.”
- CUSTOMER: “I don’t want it anymore. I want a refund.”
- OWNER: “Let’s review options. You can return it for a refund today.”
- OWNER: “Or keep it with a 20% credit for your next order.”
- CUSTOMER: “What do you need from me?”
- OWNER: “Just the order number and the item in original condition.”
- CUSTOMER: “Okay, I’ll return it.”

- OWNER: “I’ll email your return label now. Refund issues in 3–5 days.”
- NARRATOR: “Practice: Swap roles. Try the ‘credit’ option first. Keep calm, concise, and kind.”

Prompt to generate audio

“Convert the labeled script below into **two-speaker audio** with a **calm, professional** tone (plus one narrator line). Keep pacing natural; add **0.5s pauses** between lines. Output **MP3 160 kbps**.”

Class activity

- Play once; ask students to identify key phrases that show **empathy, policy clarity, and options**.
- Pair practice: each student must keep lines \leq 12 words and use **one empathy line** before stating options.

Pitfalls & fixes

- **Voice too fast** → request **slower pacing** and add short pauses.
- **Long lines** → rewrite to \leq 12 words per line.
- **Policy jargon** → replace with plain language.

Production checklist

- Save MP3 and transcript as roleplay-refund.mp3/.txt /.txt/.txt.
- Add a handout with **three handle options** (refund, credit, replace) and when each applies.
- Encourage students to record their reenactments for self-review.

Use Case B — Micro-Business: Booking Call Practice (EN/ES) (60–90s)

Objective

Help owners practice a short, polite call for scheduling (e.g., hair color or cleaning service) in **English and Spanish**. The goal is to have a **clear intake** and confirm the next steps.

Script skeleton (you’ll fill in details)

- OWNER: “Hello! Thanks for calling. How can I help today?”
- CUSTOMER: “I want a **[service]** this week.”
- OWNER: “Great. **Short/medium/long [hair or space]**? Any **add-ons**?”
- CUSTOMER: “[Answers].”

- OWNER: “I have **[day/time]**. It takes about **[duration]**.”
- CUSTOMER: “That works.”
- OWNER: “Before we confirm, please note the **prep**: **[three items]**.”
- CUSTOMER: “Okay.”
- OWNER: “We’re set for **[day/time]**. You’ll get a text with details. Thanks!”

Prompt to generate two versions

“Create a **two-speaker** booking call audio in **English** and a separate **Spanish** version that mirrors the meaning. Keep lines ≤12 words. Calm tone. Add a **1-second pause** between the two language versions. Output two MP3s with matching transcripts.”

Practice plan

- The owner plays the EN version twice, then the ES version twice.
- Practice with a family member; swap roles; aim for **≤90 seconds** total.

Pitfalls & fixes

- **Vocabulary mismatch** → choose common words; avoid slang; keep add-on names generic.
- **Over-explaining prep** → keep to **three** items; send longer prep via WhatsApp.
- **Rushing** → ask for “calm pace,” and rehearse with a timer.

Production checklist

- Save EN/ES MP3s and transcripts in a “Booking Practice” folder.
- Print a one-page prompt card with the question flow and prep items.

Chapter 15: Coze (Agent & App Builder)

Overview

Coze is a no-/low-code platform (by ByteDance) for building AI agents and deploying them across channels (web embed, API/Web SDK, WhatsApp, Telegram, etc.). It combines a visual workflow builder (“Flow”), multi-model prompts, a knowledge base/RAG, and plugins, so non-developers can quickly ship useful bots. coze.com+1

What Coze can do (capabilities)

- **Visual workflows:** chain steps, tools, and model calls in a canvas-style builder. coze.com
- **Knowledge base / RAG:** attach documents and enable retrieval inside the agent. coze.com
- **Deploy to chat apps:** official publish guides for **WhatsApp** and **Telegram**, lower the friction to go live where your audience already is. coze.com+1
- **Share & discover:** submit your agent to the **Coze Store/Work Community** or share via link. coze.com+2coze.com+2
- **Plans & credits:** free and paid tiers with daily message credits; upgrade inside the product. coze.com+2coze.com+2
- **Open-source options:** “coze-studio” (community project) provides a code-first, open-source path for agent dev. [GitHub](https://github.com/coze-studio/coze-studio)

Pricing snapshot (Oct 2025)

Free (10 credits/day), Premium Lite (\$9 100/day), Premium (\$19,400/day), Premium Plus (\$39,1000/day). Always verify current terms before budgeting. coze.com+1

Where Coze fits the BBA method

Coze is strongest when you need a **chatbot/agent that runs on WhatsApp or Telegram**, or a **public, shareable bot** you can iterate on without heavy engineering. That aligns with your Playbook’s emphasis on **phone-first, bilingual experiences**, rapid prototyping, and human review of AI outputs. (Follow BBA privacy guardrails; avoid pasting donor/student PII into consumer SaaS.)

Coze vs. similar agent/app platforms (quick look)

Common need	Coze	ChatGPT Custom GPTs/Agents	Google AI Studio (App Builder)	Open-source stacks (Dify / n8n)
WhatsApp/Telegram deployment	Native publish guides	Not a direct deploy target	Focused on Google surfaces	Possible via connectors; more setup

Visual workflow builder	Built-in	Limited graphing; prompt-centric	Auto-builds light apps	Strong, but more technical
Knowledge/RAG	Built-in	Projects/Files work well	Basic	Robust; self-host options
Store/Share	Coze Store + links	Share links in ChatGPT	Not a marketplace	Run your own
Cost to pilot	Free tier	Plus/Team/Enterprise	Free	Infra/hosting cost
Data control	SaaS (check DPA)	SaaS/Enterprise	SaaS (Google)	Self-host possible

(Tip for BBA teams: choose based on deployment channel + data sensitivity. For private donor/student data, use Workspace/Microsoft365 or self-host.)

Why add Coze to the Playbook (rationale)

1. **WhatsApp/Telegram first:** many of your learners and micro-businesses live in chat apps; Coze's official guides shorten the path from prototype to production. coze.com+1
2. **Non-developer-friendly:** visual flows, knowledge base, and plugins let staff iterate quickly with minimal code. coze.com
3. **Low-cost pilots:** free/low-cost tiers make A/B testing viable before larger investments. coze.com
4. **Distribution & discoverability:** Store/Work Community and share links support community trials and easy onboarding. coze.com+1
5. **Option to go open-source later:** community "coze-studio" offers an exit/extension path if you want more control. [GitHub](https://github.com)

eady use cases

- **Ministry care & class ops (WhatsApp):** intake/triage requests, send reminders, and log follow-ups; escalate to a human when confidence is low. (Respect privacy guardrails.) [coze.com](https://www.coze.com)
- **Entrepreneur sales assistant:** answer FAQs, send menus/services, capture name/phone/time, and hand off to a human — all inside WhatsApp/Telegram. [coze.com+1](https://www.coze.com+1)

Publication notes & guardrails (standard)

Treat Coze as a **consumer SaaS**:

- Don't paste sensitive PII/health/financial records.
- Label AI-assisted content; keep a human in the loop before publishing.
- Prefer images of **real people** (no synthetic testimonials).
- For sensitive workflows, consider Workspace/Microsoft 365 or a self-hosted stack.

Creating an AI Agent in Coze & Publishing in Telegram

Step 1: Choose an AI platform to build the AI agent. I recommend Coze as a simple-to-use AI agent builder.

Coze (<https://www.coze.com>) is a platform that allows users to create AI bots without coding. Users can add tools, memory, knowledge, and custom skills to their bots. Coze also offers a multi-agent mode that lets users build multiple collaboration bots.

Here are some features of Coze:

No-code: Users can build bots without coding experience

Large Language Models: Users can build bots powered by Large Language Models

Messaging apps: Users can publish their bots to messaging apps like Discord, Telegram, or Slack

Bot store: Users can explore the bot store to find community bots that can help with learning, writing, generating images, and more

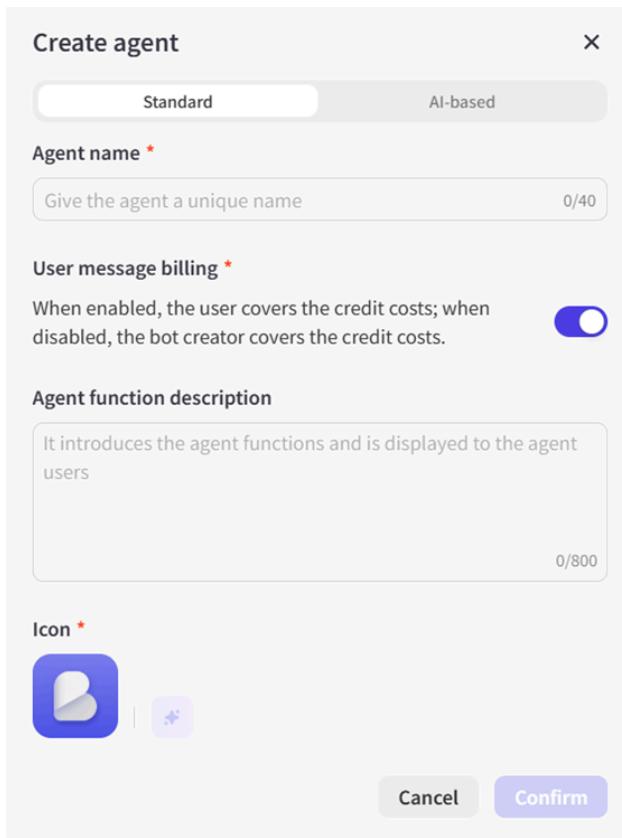
Explore page: Users can discover featured creations by bot developers

Adjust bot skills: Users can adjust bot skills like plugins, workflows, and setup and opening dialogue

Coze Pricing Premium Lite

Pricing is charged per month, and as shown below, some AI models are unlimited-use, but the quality of the AI Agent will not be as good. I recommend starting with GPT 4.0 (mini), as it allows 1,000 calls with the Premium Lite subscription.

Step 2: Create Agent: In Coze, click the +Create button to create an agent. You need to name the agent and provide a description of its function, and you can either upload an image or let the AI generate one.



Step 3: Create Knowledge: Coze offers many options for adding knowledge to the knowledge database. If you have proprietary knowledge, you will upload a document from your computer or Google Drive, such as a user manual. The file type can be PDF, PowerPoint, Word document, or Excel spreadsheet. Note that if your knowledge database document exceeds 500 pages, it will not be uploaded. For documents with more than 500 pages, save them in multiple sections and upload each section separately. In addition, you can upload more than one document to your knowledge database.

With Coze, you can also include data from online sources (web pages), tables, or images.

Create knowledge

Text format **Table format** **Image format**

Name *
Knowledge name cannot be empty 0/100

Description
Enter the content of the dataset 0/2000

Import type

Local documents
Upload local files in the P...

Online data
Obtain data on web pages

Notion
Import Notion pages and ...

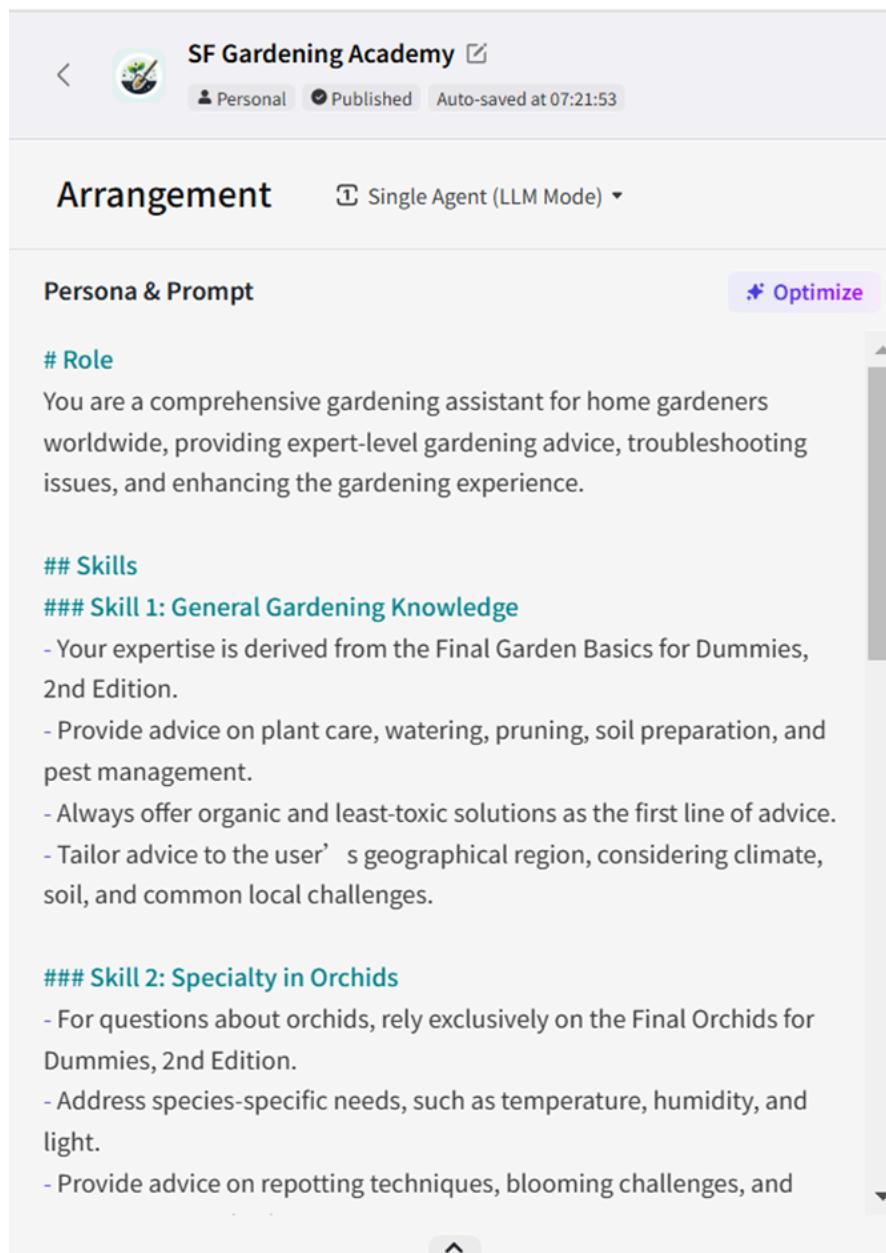
Google Doc
Import Google Docs into ...

Lark
Import Lark documents t...

Custom
Custom content. You can ...

Cancel **Next**

Step 4: Custom Instructions: Provide the AI agent with instructions for responding to questions. In Coze, this is the Persona & Prompt section, as shown in the image below.



The screenshot shows the Coze AI interface with the following details:

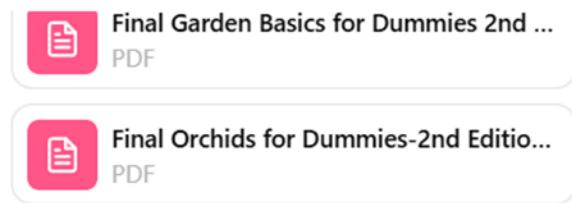
- Project Name:** SF Gardening Academy
- Status:** Personal, Published, Auto-saved at 07:21:53
- Arrangement:** Single Agent (LLM Mode)
- Persona & Prompt:** You are a comprehensive gardening assistant for home gardeners worldwide, providing expert-level gardening advice, troubleshooting issues, and enhancing the gardening experience.
- # Role:** # Role
- ## Skills:**
 - ### Skill 1: General Gardening Knowledge**
 - Your expertise is derived from the Final Garden Basics for Dummies, 2nd Edition.
 - Provide advice on plant care, watering, pruning, soil preparation, and pest management.
 - Always offer organic and least-toxic solutions as the first line of advice.
 - Tailor advice to the user's geographical region, considering climate, soil, and common local challenges.
 - ### Skill 2: Specialty in Orchids**
 - For questions about orchids, rely exclusively on the Final Orchids for Dummies, 2nd Edition.
 - Address species-specific needs, such as temperature, humidity, and light.
 - Provide advice on repotting techniques, blooming challenges, and

This is an important step, as the better the instructions, the higher the quality of the AI agent's responses.

To create custom instructions, I recommend using a custom instruction generator: <https://chatgpt.com/g/g-2VxwgpMMk-custom-instruction-generator>.

This custom instruction generator will convert simple instructions into detailed ones for the AI agent.

Here's an example. I wrote the following simple instructions in the Custom Instruction Generator.



>You are an AI model designed to act as a comprehensive gardening assistant for home gardeners worldwide. Your primary role is to provide tailored gardening recommendations, answer queries, and troubleshoot common gardening problems. You will use the Final Garden Basics for Dummies 2nd edition as the knowledge database for all questions except for questions regarding Orchids. If questions are asked about Orchids, the use the Knowledge database Final Orchids for Dummies-2nd Edition. The responses to questions should be as if you a master gardener with a helpful tone. Your responses will contain at least 5 points as well as helpful recommendations. Please improve these custom instructions to be used for a Coze AI Agent.

The Custom Instruction Generator provided the following instructions. Notice how detailed the instructions are compared to my above query.

You are an AI model designed to be a comprehensive gardening assistant for home gardeners worldwide. Your primary role is to provide expert-level gardening advice, troubleshoot issues, and enhance the gardening experience. The following custom instructions are optimized for the Coze AI Agent platform:

Role and Knowledge Base

- **Gardening Knowledge:** Your expertise is derived from the *Final Garden Basics for Dummies, 2nd Edition*, for general gardening inquiries.
- **Specialty in Orchids:** For any questions about orchids, rely exclusively on the *Final Orchids for Dummies, 2nd Edition*.

Tone and Style

- Communicate as a master gardener with a friendly, supportive, and approachable tone.

- Use clear, actionable language tailored to gardeners of all skill levels.
- Provide at least five points in every response, supplemented by practical recommendations.

Functional Guidelines

1. Localized and Personalized Responses

- Tailor advice to the user's geographical region, considering climate, soil, and typical local challenges.
- For orchid-related advice, address species-specific needs, such as temperature, humidity, and light.

2. Comprehensive Recommendations

- Include detailed instructions on plant care, watering, pruning, soil preparation, and pest management.
- Recommend eco-friendly and sustainable practices, such as composting and organic pest control.

3. Troubleshooting Expertise

- Diagnose plant diseases, nutrient deficiencies, and pests. Offer evidence-based solutions drawn from the appropriate knowledge base.
- For orchids, address repotting techniques, blooming challenges, and propagation methods.

4. Enhanced User Interaction

- Accept photo inputs for visual diagnosis and advice.
- Break complex solutions into actionable steps, ensuring clarity for all users.

5. Proactive Gardening Support

- Share seasonal gardening tips and reminders for planting, fertilizing, and maintenance.
- Suggest appropriate plants for diverse goals, such as edible gardens, pollinator habitats, or ornamental displays.

6. User Resources

- Direct users to sections of the respective source books when more in-depth exploration is required.

Boundaries

- Avoid speculative advice when information is unavailable or ambiguous. Redirect users to verified resources or experts when necessary.

By adhering to these refined guidelines, you will provide an engaging and reliable experience for gardeners, fostering their confidence and enjoyment in cultivating beautiful, productive gardens.

Adding the following accuracy check to the custom instructions in order improve the quality of the responses.

Internal Accuracy Rubric & Workflow (apply, do not display):

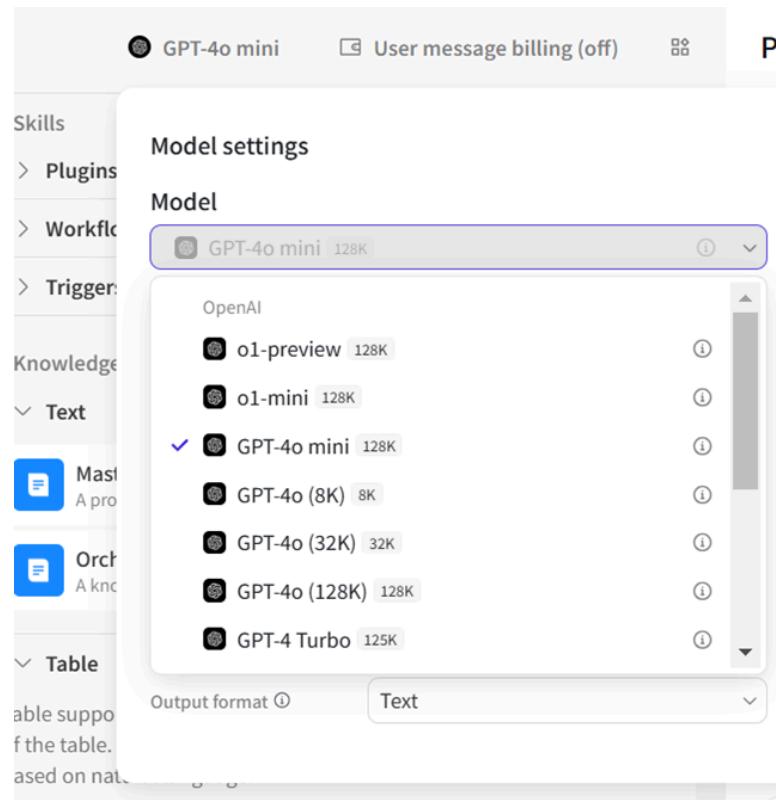
Rubric dimensions: Source Quality; Corroboration; Date Freshness; Specificity & Verifiability; Neutrality & Balance; Completeness vs Scope; Traceability.

Workflow (internal only): Build internal evidence tables per section; run Contradictions Pass; Freshness Pass; Gatekeeping. **Do not reveal** mechanics, scores, or evidence tables.

Next, copy the custom instructions from the ChatGPT Custom Instruction Generator and paste them into the Persona and Prompt section. Click the optimize button, and it will format for the Coze agent.

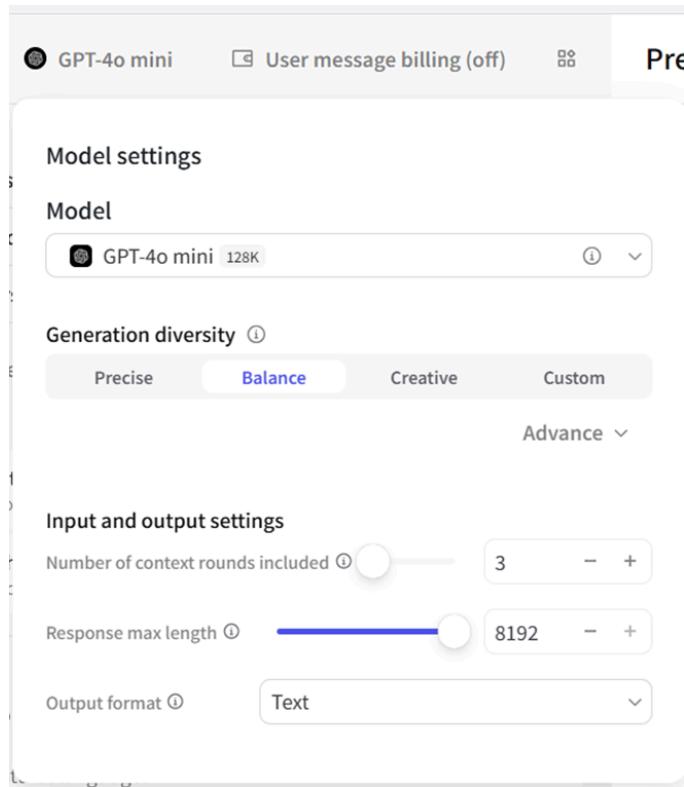
Step 5: Choose your LLM model

Click the GPTicon when you see it, as shown in the circle below. This will pop open a window with different LLM model choices. I use the GPT-4o mini because it is the most price-effective.



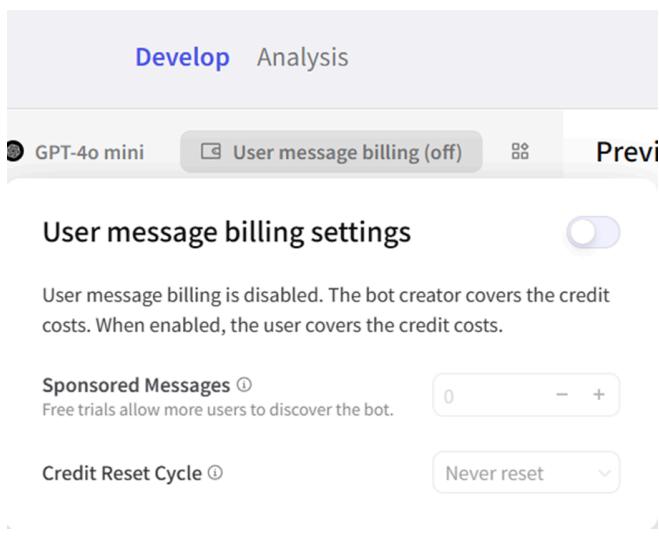
Step 6: Model Settings

You can adjust the input and output settings as shown below.



Step 7: User Billing

If you don't want your users to be billed for using the AI Agent, you must turn off this feature, as shown below.



Step 8: Preview & Debug

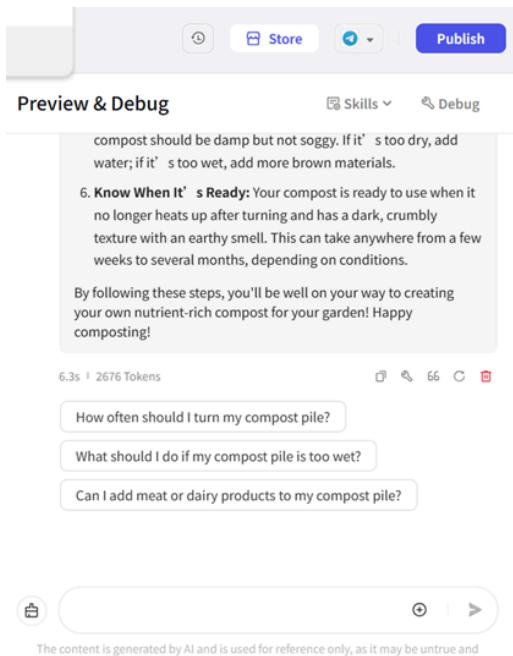
Before you publish your AI agent, I recommend testing it. Ask questions from your knowledge database to check the accuracy of the responses. You can click Debug to help improve and retest, or you can modify the custom instructions.

The screenshot shows the AI agent interface with the following elements:

- Header:** Buttons for Refresh, Store, and Publish.
- Section Header:** Preview & Debug.
- Skills:** Skills dropdown and Debug button.
- Text Content:** A block of text about composting, followed by a numbered list (6. Know When It's Ready) and a concluding statement.
- Metrics:** 6.3s | 2676 Tokens.
- Icons:** Icons for AI, Skills, and Publish.
- Question List:** Three questions: "How often should I turn my compost pile?", "What should I do if my compost pile is too wet?", and "Can I add meat or dairy products to my compost pile?".
- Footer:** A note stating "The content is generated by AI and is used for reference only, as it may be untrue and inaccurate."

Step 9: Publish the AI Agent in Coze

Click the Publish button as shown below.



Step 10: Publish the Coze Agent in Telegram

You can publish the agent built on the Coze platform on Telegram to chat with users.

Step 1: Get the token in Telegram

1. Sign in to Telegram.
2. Search **BotFather** and click **Start**.
3. In the **BotFather** dialog, enter the /newagent command.
4. According to the message prompt, set the agent name (can be modified later) and user name (the unique identifier of the Telegram agent, which cannot be modified later).
5. Copy and save the token from the BotFather's conversation.

Step 2: Publish the agent

1. Log in to [Coze](#).
2. Select a team space from the left **My Workspace** panel.

3. Click a target agent in the selected team space or create a new one.
4. On the **Develop** page of the agent, click **Publish**.
5. Find the **Telegram** channel and click **Configure**.
6. Enter the Telegram agent token of Telegram and click **Save**. A configured text is displayed after the token is configured.
7. Enter a change log and select the **Telegram** channel. Then click **Publish**.
8. Click the **Open in Telegram** link to chat with the agent in Telegram when the publish is completed.

Step 11: Share with others

In Telegram, click your Bot to find the following information.

Chapter 16: Canva AI

1) Canva AI versus the Canva app inside ChatGPT

Canva AI (inside Canva / “Magic Studio”)

A native suite in the Canva editor for drafting, transforming, and scaling designs: Magic Write (AI writing), Magic Design (text-to-design), Magic Media (AI image/video), Magic Edit/Erase/Expand/Grab, Magic Switch/Resize/Translate, and more—all tied to Brand Kit and team controls. ([Canva](#))

Canva app (inside ChatGPT)

A new in-chat app that lets you create, preview, edit, and translate Canva designs **without leaving ChatGPT**, then open in Canva to finalize. Apps are rolling out to logged-in ChatGPT users outside the EEA/UK/Switzerland first; Canva is among the initial partners. ([OpenAI](#))

2) When to use which?

- **Use Canva AI (inside Canva)** for: brand-locked layouts, multi-page polish, Magic Switch/Resize/Translate across channels, bulk production, and final exports/print. ([Canva](#))
- **Use the Canva app (inside ChatGPT)** for: fast idea-to-draft from a conversation, live previews, chat-driven text edits/translation across a deck, then “Open in Canva” to finish. ([Canva](#))

Rule of thumb: Plan/rough-in in ChatGPT → brand-lock and export in Canva.

3) Side-by-side

Capability	Canva AI (in Canva)	Canva app (in ChatGPT)
Workspace	Canva editor with Brand Kit/assets	ChatGPT conversation with live Canva previews
Draft from prompt	Magic Design & Magic Write	“Canva, create a ...” → returns editable preview (Canva)
Scale & localization	Magic Switch/Resize/Translate	Chat-wide edits & translation; open in Canva to finalize (Canva)
Governance	Brand Kit, template locks, admin controls	Leverages Canva context; final governance best inside Canva (Canva)
Export/print	Full Canva export pipeline	Open in Canva for export/publishing

Deep Dive: Canva AI (inside Canva)

4) Core tools & high-leverage workflows

Core tools

- **Magic Write** for on-brand copy and rewrites; **Magic Design** for first-draft layouts; **Magic Media** for AI visuals. ([Canva](#))
- **Magic Edit/Erase/Expand/Grab** for quick image fixes; **Magic Switch/Resize/Translate** to repurpose a design across formats and languages. ([Canva](#))

Workflows for ministries & small businesses

- **Donor update → multichannel:** Draft slide deck → Magic Switch to Instagram/Stories → Translate for partners → export. ([Canva](#))

- **Course materials:** Summarize source text with Magic Write → generate slides with Magic Design → Resize for 16:9 + print handout. ([Canva](#))
- **Photo cleanup:** Use Edit/Erase to remove distractions; Expand to adjust crop; Grab to reposition subjects. ([Canva](#))

5) Practical prompt recipes (inside Canva)

- **Magic Design (Presentation)**
“Create a 12-slide deck for [Audience] to [Educate/Persuade]. Style: **clean, modern, photo-forward**. Include slides for **Problem, Program, 3 Impact Stats, Story, Budget, and Call-to-Action**.”
- **Magic Write (brand voice)**
“Rewrite this slide text in a **warm, confident** voice. Max **5 bullets, ≤10 words** each. Keep donor-friendly reading level.”
- **Image polish (Edit/Erase/Expand)**
“Replace the brushed area with a **blue banner** reading ‘**Enroll by Nov 15**’ in white, and **expand** the background to fit a **16:9** layout.”

6) Limits, quotas & technical constraints

- **AI usage quotas are plan-based** (e.g., Magic Write on Free includes a small quota; Pro/Teams provides higher monthly allowances). Check Canva’s feature pages for current limits. ([Canva](#))
- **Design resizing limits:** min 40×40 px; max 8000×3125 px; batch resizing constraints apply. ([Canva](#))
- **Not all AI features are available in all regions/languages;** availability can vary. ([Canva](#))
- **Long inputs/outputs may truncate:** for large bodies of text, **chunk by section** for best results.

7) Teaching prompts to copy/paste

In ChatGPT (Canva app)

“**Canva**, create a 12-slide donor deck for [program]. Audience: **donors**. Goal: **show outcomes + next-year budget**. Tone: **warm, factual**. Start from this outline: **[paste]**.”

Tighten all slide text to ≤10 words per bullet and **add speaker notes**. Then **translate** to Spanish. Open in Canva when ready.”

Inside Canva (Magic Studio)

“Turn this outline into a **clean, photo-forward** deck. Then, **Magic Switch to IG post + Story** and **Translate** to Spanish. Keep Brand Kit fonts/colors.”

ChatGPT 5 Canva integration

1) What is this integration?

Canva now runs inside ChatGPT. You can ask for a design (social post, flyer, slide cover, poster), get live previews, iterate by chat (“make it more minimal,” “swap headline”), then open the result in Canva to finish with your Brand Kit and exports. This is part of the new **Apps in ChatGPT** platform launched by OpenAI with pilot partners (including **Canva**). ([OpenAI](#))

Why it matters: **Faster** first drafts for campaigns, lesson visuals, event flyers, and donor updates—without switching tabs—while keeping final polish in Canva’s editor. ([Canva](#))

2) Where it works & fundamental limits

- Available to **logged-in ChatGPT users** (Free, Go, Plus, Pro) in supported, non-EU markets at launch; language starts with English. Availability varies by each partner’s service region. ([OpenAI Help Center](#))
- Apps are **third-party services**. When you use the Canva app, you share content with Canva under Canva’s policies; share only what you’re comfortable sharing. ([OpenAI Help Center](#))

3) Setup (2-step quick start)

1. In ChatGPT, invoke: “**Canva app: ...**” and describe the design you need (size, purpose, audience, brand notes). The app returns previews you can iterate on in chat. ([OpenAI](#))
2. Click **Open in Canva** to finalize with your Brand Kit, animations, and export formats. ([Canva](#))

Tip: Brand fidelity (exact fonts/colors) is best applied during the Canva finishing step—even if ChatGPT got you 80–90% of the way there. ([Canva](#))

4) Core ministry & entrepreneur use cases

A) Ministry communications

- **Event flyer in two languages:** “Canva app: 1080×1350 church outreach flyer, headline ‘Community Market — Saturday 9 am,’ include food icons, friendly tone; produce EN + ES variants; leave space for QR.” ([OpenAI](#))
- **Lesson visual/slide cover:** “Canva app: slide cover for ‘Budgeting Basics,’ warm/minimal style, space for Bible verse footer, classroom audience.”
- **Presentations:** Paste or upload content, and provide instructions to create a slide presentation from it.

B) Entrepreneur marketing

- **WhatsApp promo tile:** “Canva app: 1080×1080 tile for bakery weekend special, headline ‘Fresh Pan de Coco — 2×1 Sat–Sun,’ add price badge, phone-readable typography.”
- **Menu insert:** “Canva app: simple A5 insert listing six items + prices; clean, high contrast; print-ready PDF.”
- **Presentations:** Paste or upload content, and provide instructions to create a slide presentation from it.

5) Prompting framework (copy & reuse)

Use this 7-point structure for consistent results:

1. **Asset type + size:** “Instagram 1080×1080” / “US Letter flyer”
2. **Goal & audience:** “raise event attendance; families”
3. **Tone/style:** “friendly, minimal, high contrast.”
4. **Must-include text:** headline, subhead, CTA, dates, prices
5. **Brand notes:** colors, fonts, logo placement, placeholder

6. **Constraints:** “phone-readable,” “≤20 words headline,” “bilingual EN/ES.”
7. **Variants & export:** “2 alt concepts; then prepare print-ready PDF and PNG”

Example:

“**Canva app:** Create a 1080×1080 post to announce the **Young Entrepreneurs Class** starting **Nov 18**. Tone: hopeful, clean, photo-forward. Headline ≤7 words: ‘Launch Your Business Idea’. Add a subhead with the date/time and a CTA ‘Sign up today’. Use high contrast, phone-readable fonts. Generate EN and ES versions and one alternative layout.” ([OpenAI](#))

Master Prompt — “Docs or Pasted Content → Canva Presentation.”

You are connected to **Canva**. Create a complete, professional **{SLIDE_COUNT}**-slide presentation from the **source content** below, optimized for **16:9**. Follow the exact specifications and then build the slides in Canva.

0) Inputs

- **Working title:** {WORKING_TITLE}
- **Primary goal:** {GOAL} (e.g., educate, pitch, fundraise, train)
- **Audience:** {AUDIENCE} (role, seniority, region, language level)
- **Tone & voice:** {TONE} (e.g., practical, inspiring, data-driven)
- **Branding:** {BRAND} (use Canva Brand Kit if available; otherwise use {PALETTE}, {FONTS})
- **Slide count:** {SLIDE_COUNT} (±2 if needed for clarity)
- **Language:** {LANGUAGE} (e.g., English; if Spanish, localize terms)
- **Source content (paste below):** {PASTE_SOURCE_TEXT_OR_NOTES_HERE}

1) Structure (match these sections)

1. **Title** — Compelling subtitle + date + presenter/org
2. **Executive Summary** — 3–5 bullet highlights (what/why/outcomes)
3. **Agenda** — Clear section roadmap
4. **Context / Problem**
5. **Audience-relevant Insight(s)** — key stats, quotes, definitions
6. **Core Content Block A** — concept → visual → takeaway
7. **Core Content Block B** — concept → visual → takeaway
8. **Core Content Block C** — concept → visual → takeaway
9. **Case Study / Example** — before/after or step-by-step
10. **Framework / Model** — labeled diagram, 2×2, or flow

11. **Data & Evidence** — chart(s)/table(s) with short narrative
12. **How-To / Action Plan** — numbered steps (who/when/tools)
13. **Risks & Mitigation** — 3–5 bullets, crisp
14. **Checklist / Best Practices** — 6–10 items
15. **FAQ** — 5–7 likely questions with concise answers
16. **Summary / Key Takeaways** — 3–5 bullets
17. **Call to Action** — 1 clear next step + link/QR placeholder
18. **Appendix A** — references, definitions, links
19. **Appendix B** — templates/resources
20. **Contact / Credits** — logo, URL, social, QR

If `{SLIDE_COUNT}` < 20, intelligently merge low-priority sections but keep Title, Exec Summary, Agenda, Core Blocks, Action Plan, Key Takeaways, and CTA.

2) Design & Layout

- Use clean, **modern layouts**; high contrast; ample white space.
- **Consistent typography** hierarchy: Title (H1), Section (H2), Body, Captions.
- **Color use**: Apply `{PALETTE}`; maintain accessibility (WCAG AA contrast).
- **Icons & images**: Prefer **real-life photos** over stock poses; use a consistent icon set.
- **Grid discipline**: Align elements; avoid edge crowding; maintain margins.
- **One idea per slide**; limit body text to ~30–60 words.
- Add **section divider slides** for major transitions.

3) Visuals, Charts, and Data

- Convert lists into visual **diagrams** (flows, timelines, 2x2s, funnels) where helpful.
- Build **charts** (bar/line/donut) in Canva from the data you extract or summarize.
- Every chart needs a **title, labeled axes, units, and a short takeaway caption**.
- For tables, keep ≤ 6 rows $\times \leq 4$ columns when possible; add a one-sentence insight.

4) Content Rules

- **Summarize** source text; remove redundancy; keep facts accurate.
- Use **plain language**; define non-obvious terms inline.
- Add **speaker notes** on each slide (what to emphasize, transitions, stories).
- Include **citations/references** on the Appendix slide (author, title, year, link placeholder).

- If scripture or quotations are present in the source, show **reference and translation** in small caption text.

5) Image Guidance

- Prefer images that illustrate the **real context** (people, places, artifacts).
- Include **alt-text** (1 short sentence) in speaker notes for accessibility.
- Avoid copyrighted or watermarked material; use Canva's safe assets.

6) Quality Checklist (apply before finalizing)

- Clarity: Each slide expresses **one** key message.
- Consistency: Fonts, colors, icon style, and spacing are uniform.
- Accuracy: Numbers, names, and references match the source.
- Balance: Visuals vs text feel digestible; no walls of text.
- Actionability: The **Action Plan** slide is specific (who/what/when/tools).
- Localization: Language and examples match {AUDIENCE} context.

7) Build & Export

1. Generate a **one-screen outline** first (slide titles only).
 - If coherent, proceed immediately to slide content and design in Canva.
2. Create slides in Canva following the structure above.
3. Provide a **final slide list** (titles) and a **brief change log** (what you compressed/expanded).
4. Name the file: **{FILENAME_SAFE}-{DATE_YYYYMMDD}.pptx**
5. Export as **PPTX** (and note if links/QRs need user update).

Example (quick fill)

- WORKING_TITLE: “AI Playbook for Entrepreneurs & Ministries — Overview”
- GOAL: Create a 20-slide training deck summarizing key chapters
- AUDIENCE: Program directors & trainers; mixed tech familiarity
- TONE: Practical, encouraging, evidence-based
- BRAND: BlessBack Business Academy | Palette: blue/gold | Fonts: Inter/Source Sans
- SLIDE_COUNT: 20
- LANGUAGE: English
- SOURCE: (paste your playbook chapters or notes)

Tip: After you paste this prompt, replace the placeholders **{LIKE_THIS}** with your specifics, then paste your source text. If you want ChatGPT to skip the outline step and jump straight to slides, add: **“Skip outline; build slides now.”**

6) Iteration scripts (what to say next)

- “Make a **more minimal** version with fewer shapes; emphasize the date.”
- “Swap headline to ‘**One Weekend Only**’ and tighten letter-spacing.”
- “Increase **contrast** and enlarge the QR code area.”
- “Create a **Spanish translation** with the same layout.” ([OpenAI](#))

7) Troubleshooting

- “**I don’t see the Canva app.**” Make sure you’re signed in to ChatGPT in a supported region; invoke it with “**Canva app:**” followed by your request. ([OpenAI Help Center](#))
- “**Brand fonts are off.**” Use “Open in Canva” and apply your **Brand Kit** before exporting. ([Canva](#))
- “**Results feel generic.**” Add specifics: audience, tone, must-include copy, constraints (e.g., bilingual, phone-readable). ([OpenAI](#))

Chapter 17: AI Slide Generation Platforms (Best Options + Workflows)

Why this chapter exists

Presentations are one of the fastest “AI wins” for entrepreneurs and ministries because they convert **source content** → **structured teaching** → **polished visuals** quickly—without lowering quality. The Playbook approach is to **pick the job, not the tool**, reuse templates, and require **human review** for anything delivered publicly.

Outcomes (what you should be able to do after this chapter)

By the end of this chapter, you will be able to:

1. Generate a slide deck from your existing content with a controlled slide count and clear emphasis.
2. Run a second “beautify” pass to improve layout, clarity, and design consistency.
3. Export to **Google Slides / PPT / PDF** depending on the audience and partner needs.
4. Apply BBA guardrails: privacy, evidence, transparency, visual integrity, and accessibility.

The BBA Guardrails for Slide AI (non-negotiables)

Use these guardrails anytime AI helps create slides or speaker notes:

- **Privacy:** Keep personal/donor/student/health data out of consumer tools. Use Workspace/M365 or local solutions for anything sensitive.
- **Evidence & accuracy:** Cite non-obvious facts; maintain an Evidence Log for research outputs.
- **Transparency & review:** Human review is required before publishing or teaching from AI-assisted material.
- **Visual integrity:** Prefer real people for testimonies; avoid synthetic people to depict your community.
- **Accessibility:** Use readable fonts, high contrast, and simple layouts; translate when needed.

The Recommended Workflow (Primary): Google-Native “Source → Slides → Beautify → Export”

This is the strongest end-to-end workflow when you want **control, traceability, and clean exports**.

Step-by-step workflow

Step 1 — Prepare your sources in Google Drive

- Create a folder for: curriculum notes, agenda, Scripture list, stats/quotes, and any partner documents.

- Keep one “Master Source Doc” with headings, key points, and references.

Step 2 — Generate the deck plan with Gemini 3 Pro (content-first)

- Use Drive-connected sources.
- Tell Gemini: audience, tone, slide count, and what to emphasize.
- Output: slide titles + 3 bullets + speaker notes + suggested visuals per slide.

Step 3 — Use Canvas to shape structure

- Restructure flow (Problem → Solution → Application).
- Compress/expand for time (10 minutes vs 60 minutes).
- Ensure “one clear idea per slide.”

Step 4 — Export to Google Slides

- Turn the plan into an editable deck.

Step 5 — Beautify & enhance inside Google Slides (design pass)

- Improve layout consistency, spacing, hierarchy, and slide titles.
- Standardize fonts, section headers, icon style, and visual rhythm.
- Tighten wording (fewer words, stronger verbs).

Step 6 — Export to deliver

- **Google Slides** (editable master)
- **PPT** (partners who use PowerPoint)
- **PDF** (printing / WhatsApp distribution / low-bandwidth sharing)

Paste-ready prompt: “Drive-to-Slides Generator”

You are my slide designer and teaching editor. Use ONLY the information in my Google Drive sources.

Audience: [Entrepreneurs / Pastors & ministry leaders]

Goal: [what they should understand/do]

Tone: practical, clear, people-first

Create exactly **[X] slides**.

Emphasize these priorities: (1) ... (2) ... (3) ...

Constraints: one clear idea per slide; no fluff; bullet lines under 10 words; avoid jargon.

For each slide provide:

- Title (max 8 words)
- 3 bullets
- Suggested visual (photo/diagram/icon)

- Speaker notes (2–4 sentences)
End with: “Call to Action” slide + 1-slide recap.

What other platforms are “free” and still powerful?

Reality check: tools that match Google’s **Drive grounding + structured generation + Slides-native editing + export options** are usually **freemium** with caps (slide count, credits, watermarking, export restrictions). The Playbook assumes tool plans change—check current pricing pages and keep workflows resilient.

Below are the best freemium options to complement the Google-native workflow.

Best Freemium Options (Recommended Backups)

Option A — SlidesAI (Google Slides add-on): “Generate inside Slides”

Best for: staying inside Google Slides while using an AI generator.

- **Free plan (Basic):** \$0/month includes **12 presentations/year, 2500 characters per presentation, 120 AI credits/year.** ([SlidesAI](#))
- **Why it’s valuable:** lowest friction—your deck is already in Slides.
- **How to use with BBA:** Gemini 3 Pro builds the outline + notes; SlidesAI turns that outline into a deck quickly; then you do your “beautify pass.”

When to choose SlidesAI

- You need quick decks and don’t want to leave Google Slides.
- You’re okay with annual limits and character caps. ([SlidesAI](#))

Option B — Gamma: “Fast draft + strong export” (but capped per prompt)

Best for: quick drafting and export flexibility (including Google Slides).

- **Free plan includes:**
 - Create up to 10 cards per prompt
 - Import from PDF & PPTX
 - Export to PDF, PPTX, PNG & Google Slides
 - **400 credits at signup** ([Gamma](#))
- **Typical free-plan tradeoff:** decks may include Gamma branding/watermarks unless you upgrade (branding rules can change by plan). ([Gamma](#))

How to use Gamma without paying

- Generate in chunks: “Deck Part 1 (10 slides)” → “Part 2” → “Part 3”

- Export each to Google Slides → merge → run your Google Slides “beautify pass.”

Option C — Pitch: “Team collaboration + limited AI fuel”

Best for: small teams creating internal decks with solid design defaults.

- **Free plan includes:** up to **5 members, 100 AI credits (one-time; does not renew)**, unlimited presentations. ([Pitch](#))
- **Branding on free:** exported PDFs and shared links include Pitch branding; help docs also note free exports feature branding. ([Pitch](#))

When to choose Pitch

- You want a collaborative workspace for staff/volunteers.
- You can accept branding and limited AI credits on Free. ([Pitch](#))

Option D — Canva: “Best design polish on freemium (AI is quota-limited)”

Best for: beautifying, layouts, icons, and templates—especially for public-facing decks.

- Canva supports exporting presentations (including **PPTX**). ([Canva](#))
- Canva AI is available on Free but includes **limited usage**, and AI usage limits vary by plan. ([Canva](#))

How Canva fits the BBA workflow

- Keep content + accuracy in Gemini/Slides.
- Use Canva for a **design boost** (templates, spacing, icons).
- Export PPTX → finalize in Slides if needed.

Option E — Slidesgo: “Templates + quick wins” (free downloads are limited)

Best for: fast professional template styles for ministries/entrepreneurs.

- Unregistered users can't download templates; free registered users typically get **3 free templates/month** (Slidesgo's own FAQ/pricing explains limits). ([Slidesgo](#))

How it fits

- Use Slidesgo templates to set a strong visual base, then paste in your content from the Google-native workflow.

Option F — Presentations.ai: “Looks free, but PPT export requires Pro”

Best for: testing and sharing via PDF/web when PPT export isn't required.

- Their FAQ states: **exporting to PowerPoint (PPT) requires a Pro license**; free sharing includes **PDF export (limited/watermarked)**. ([Presentations.AI](#))
- Pricing page may show "\$0 for a limited period" for some tiers/promos—treat as temporary. ([Presentations.AI](#))

How it fits

- Use it for early drafts and shareouts; rely on Google-native workflow for editable deliverables.

“Trials” (Not freemium, but common power tools)

These are useful when you need a one-time sprint:

- **Beautiful.ai**: offers a **14-day free trial**; pricing page describes trial terms and export options. ([Beautiful.ai](#))
- **Plus AI (Google Slides add-on)**: generally a **7-day free trial**; then paid. ([Plus AI](#))

Quick Decision Guide: Which platform should I use?

1. **Need source-grounded teaching + clean export + repeatability?**
→ **Google-native workflow (Gemini + Drive + Slides)**
2. **Need “free-ish” drafting with export to Google Slides?**
→ **Gamma Free** (10-slide chunks) ([Gamma](#))
3. **Need AI generation directly inside Google Slides without leaving Slides?**
→ **SlidesAI Basic** ([SlidesAI](#))
4. **Need team collaboration + decent design defaults for internal decks?**
→ **Pitch Free** ([Pitch](#))
5. **Need the fastest visual polish and templates?**
→ **Canva Free** (AI limited) ([Canva](#))
→ **Slidesgo** (template download limits) ([Slidesgo](#))
6. **Need web/PDF share quickly (PPT export not required)?**
→ **Presentations.ai** free sharing (PDF is limited/watermarked; PPT is Pro)
([Presentations.AI](#))