

# APPLYING FIRST & SECOND PRINCIPLES

Building AI tools around real customer needs.

Richard Getz  
richard@getzai.com  
484-680-5153  
<https://x.com/RichardAGetz>

## First Principle

Steve Jobs taught us to:

"You've got to start with the customer experience and work backwards to the technology. You can't start with the technology and try to figure out where you're going to sell it."

## Second Principle

Jony Ive taught us:

“Simplify complexity through design.”

‘Great technology alone isn’t enough—it must delight users and solve real problems.’

Identifying real user needs then stripping away complexity and obsessing over intuitive design, Apple made products that people couldn’t live without.

*LLMs today are dull tools, demanding too much from users beyond simple interactions, erecting a steep barrier to advanced workflows. This is solvable! Users want more and will pay for it.*

*I think all models will soon be indistinguishable in intellect but very distinct in the tools they offer and how people interact with those tools.*

*-Getz*

## Epic Objective

Enable AI platforms to become indispensable to everyday users by structuring consumer-centric workflows that empower creativity, simplify complexity, and drive engagement. This brief takes existing tools and makes them indispensable. Next would be us discussing new and novel solutions and tools that millions of users need today.

## Epics & Stories

### 1. **Dashboards**

A workspace where an individual or team can work with inputs and outputs from different AI tools within one UI.

### 2. **Projects**

Establish structured, consumer-centric projects that manage tasks, files, and workflows automatically.

### 3. **Coding Projects**

Lower the barrier to coding with intuitive guidance, disciplined execution, and end-to-end project management.

### 4. **Deep Research**

Deliver interactive, precise research workflows that blend discovery, documentation, and dynamic presentations.

### 5. **Training in Contrarianism**

Embed risk-awareness and negative-space analysis into LLM training to proactively identify and mitigate potential failures.

### 6. **Podcaster**

Transform static content into dynamic, conversational podcasts with customizable AI personalities and interactive workflows.

Your customers already pay third parties for solutions built on your API- why not capture that revenue directly?

---

## Why Should You Partner With Me?

### Partnering with me delivers:

---

**Understanding Customers:** With years of experience as a software project manager leading discovery sessions and training, I've developed a deep understanding of what end users truly need. They seek *software that is personal and powerful yet easy-to-use* software that doesn't just complete tasks they lack the time or skills for—it solves problems they can't tackle on their own. Users want solutions that remove the fatigue of wrestling with complicated platforms, delivering an intuitive experience where their workflow flows unimpeded. They crave tools that act as a natural extension of how they currently think, aligning seamlessly with their mental processes. Beyond functionality, they want software that's fun to use, sparking enjoyment in their daily tasks. It should know what they don't—offering insights and capabilities that help them achieve what they can't on their own—ultimately making them more than they are without it, expanding their potential in meaningful ways.

**Future-Forward Roadmaps:** Forecasting AI and robotics trends years in advance, building today's foundations for seamless future integration. Understanding future customer technology needs and designing solutions now.

### Visionary Insights:

---

1. **AI-Augmented Development Ecosystems:** I see your software engineers partnering with your AI to evolve frameworks and codebases via RAG-powered updates—downloading new libraries, dependencies, and methodologies into the app *where* capabilities expand continuously without retraining. A growing library of design templates, code boilerplates, testing methodologies, and security updates allowing your LM to have the most up to date codebase and testing frameworks without re-training or using outdated training data. Planning this today will ensure your coding app is future proofed and can advance quickly.
2. **Domestic Robotics Orchestration:** I foresee household/*office* humanoids *managing* other robots to *perform* home-care, lawn-care, *medical-care while* integrating with smart home devices. Planning smart-home connectivity standards and training today will future proof humanoids.
3. I foresee household humanoids using autonomous vehicles for grocery runs *or errands and* will need modular battery backpacks—designed for easy attachment, docking ports, and charging—to extend operational time. They'll also need a scaled-down app device (Android-based app phone) with wireless charging and Bluetooth connectivity, allowing seamless access to apps, such as the grocery store app or Amazon for returns, as well as powerful LLMs. Planning hardware integration (mounting locations, docking interfaces, and accessory ports) today ensures future humanoids can utilize these tools without redesign.

## Dashboards

***In the rapidly evolving AI landscape, users will inevitably seek platforms that offer not just powerful native tools, but also the flexibility to integrate diverse Large Language Models (LLMs) and a rich ecosystem of functionalities. The critical question for any AI provider becomes: will users find this comprehensive, versatile experience on your platform, solidifying their loyalty and investment, or will they be forced to seek out and pay third-party aggregators to achieve their goals? The answer will determine market leadership.***

A single UI dashboard where I can drag out deliverables from one AI tool into another AI tool (or project or teammate), so I do not have to have download and upload between interfaces.

Whether this is an output text, image, chat, video, audio, I want the ability to grab it and drag it onto another tool. I might have a podcast script I am working on with my AI and the final script can be dragged over to the Audio tool which will start a new workflow there that either I or a teammate can work on.

I might have multiple teammates generating different assets that get pushed into other tools for other teammates to use as their inputs.

If a workflow is not what I want to kick off, I might drag over to a teammate icon to notify them or a project icon to add to project files.

Crucially, this dashboard must be engineered for profound extensibility, enabling seamless integration not only with a platform's native AI tools but also with a diverse array of third-party Large Language Models (LLMs). This open architecture is fundamental to **owning the user experience** from end to end. As AI models rapidly approach parity in raw intellect, the true differentiator—and the key to market leadership—will be this unified experience and the richness of the available tooling. Power users, especially, will demand the flexibility to orchestrate multiple LLMs, each chosen for its specific strengths. By providing this comprehensive integration capability, the platform positions itself as the indispensable default workspace. This strategy is vital: it caters to sophisticated needs, solidifies user retention, and directly addresses the issue of **preventing users from diverting to third-party platforms where they might pay discounted API rates while the primary platform loses out on that full revenue potential**. Instead, users are kept within the ecosystem, contributing to base-tier subscriptions at a minimum. Ultimately, the most robust and versatile platform—the one that centralizes AI interactions and truly empowers users—is the one that will win, underscoring the principle: 'AI will be indistinguishable in intellect, but very distinguishable in how people use it and the tooling available.'

*Models will soon be indistinguishable in intellect but very distinct in the tools they offer and how people interact with those tools.*

*Simplify complexity through design.*

## Unlocking AI for Everyone: UI/UX Principles for Mass Adoption

The current AI landscape, while powerful, often feels like an insider's game, risking user burnout before they even tap into AI's true potential. Users face a confusing array of options—OpenAI's alphabet soup of model names, Google's ever-expanding roster like Gemini, Jules, and Flow, and numerous others.

This complexity creates a significant barrier, especially for everyday users who simply want great results without needing a Ph.D. in AI or wrestling with tedious setups. My guiding principle, "simplify complexity through design," is paramount to making AI indispensable for millions, not just a select few.

While the "behind-the-scenes" thinking of LLMs or the ability to manually pick a model might intrigue tech enthusiasts (and can certainly be tucked away under a caret for power users), for most people, this quickly becomes overwhelming. They want quality, fast. They don't want the FOMO of wondering if they picked the "wrong" model. The goal is to prevent users from abandoning these powerful tools due to frustration.

To make AI truly accessible and delightful, I propose a UI/UX built on these core tenets:

### Intelligent Automation & Blended Expertise: The "Easy Button" for AI:

The platform should be the expert, automatically selecting and even blending the best AI models for any given task. A single user request might leverage multiple specialized AIs behind the scenes to deliver a comprehensive result. Users shouldn't have to guess or manually orchestrate this. If they ask for an image, they get a great image. If they need research, they get insightful results. If the chat requires research, thinking, and image generation, it just works without the user having to. This eliminates decision fatigue and ensures optimal outcomes, every time.

- *Optional Feedback for Trust & Transparency:* Initially, the AI can offer a friendly heads-up, like: "Great idea! I'll use Deep Search for the data, Deep Think to refine it, and Image Flash to create those visuals for you." This builds confidence, and over time, as users trust the AI's smarts, even this can become more subtle.

### One Seamless Experience for Everything:

Forget juggling different apps or chats for different needs. Whether it's writing, image creation, coding, or analyzing data, it all happens in one intuitive space. This unified environment means users can effortlessly flow from one idea to the next, using text, voice, images, or code, all within a single, cohesive project.

### **Focus on Your Goals, Not My Tech:**

---

Users should be able to express what they want to achieve in plain language. The AI handles the "how," intelligently choosing the right tools behind the scenes. This outcome-oriented approach means users can focus on their creativity and productivity, not on deciphering technical jargon.

### **Dynamic Summaries for Lasting Clarity:**

---

Long chats or complex projects can easily lose focus. The UI should maintain a concise, dynamic summary of the current project or conversation's key objectives and progress. This summary could be expanded at any time, allowing both the user and the LLM to quickly re-orient and ensure continued relevance, preventing drift and wasted effort.

### **Clarity Without Clutter & Intuitive Navigation:**

---

While the AI takes the lead, a subtle, unobtrusive indicator can show which model was used for a particular output—much like how "thinking" processes can be optionally viewed today. This offers transparency for those interested, without cluttering the experience.

- *Dive Deeper, When You Want To:* For users who want to explore a specific aspect of their project, like all the images generated or a particular branch of the conversation, the UI could offer easy filtering. The UI should also visually indicate any significant "branches" or focused interaction streams the user has explored, allowing them to easily click back into those specific contexts.

By embedding these principles into the user experience, I aim to transform AI from a complex set of tools into a smart, intuitive partner. This isn't just about better UI; it's about making AI accessible, enjoyable, and indispensable for everyone, driving widespread adoption and sustained engagement.

## Projects

A universal AI framework that autonomously organizes tasks, documents, conversations, and insights for any project—from software sprints to garden redesigns to vacation planning—freeing users to focus on creativity, not structure. Unlike most AI platforms with static file uploads and user-managed tasks, this solution dynamically builds and maintains projects, adapting to user preferences for a tailored experience.

### Why This Matters

- Users want to ideate, not wrestle with administrative overhead they're ill-equipped for.
- AI excels at organizing artifacts, selecting relevant files, and asking questions to drive optimal outcomes.
- Long chats overwhelm LMs and users with bloat, drifting to discarded ideas; this framework keeps focus sharp.

### Context & Positioning

Unlike fragmented tools that burden users with setup and maintenance, this solution treats every chat as a living project, proactively managed by the LM to streamline workflows and reduce confusion.

### Features

1. **Smart Project Scaffolding:** Auto-generates tailored files, task checklists, and folder structures based on project type (e.g., a wedding planner's timeline or a developer's codebase). Adapts to user preferences over time for personalized setups.
2. **Dynamic Asset Management:** Tracks and updates requirements, run logs, release notes, to-dos, and research artifacts, ensuring accuracy without user effort (e.g., auto-saves a gardener's seasonal plan).
3. **Focused Scratch Pad:** A dedicated chat buffer captures notes, ideas, and highlights, using Context Refactoring to filter outdated or irrelevant concepts, keeping main chats concise and on-topic. Self-curated RAGs maintain accessible references.

### Use Cases

- **Smart Scaffolding:** Starting a garden project, I get a tailored checklist and seasonal plan, letting me dive into design instantly.
- **Asset Tracking:** As a wedding planner, my evolving preferences are auto-saved into project files, keeping chats focused and references handy.
- **Clean Chats:** In long discussions, the Scratch Pad isolates active ideas, preventing confusion from discarded concepts and saving scrolling time.
- **Automated Logs:** As a developer, run logs and version history are maintained, speeding up debugging with clear references.

## Coding Projects

A sprint-oriented framework that autonomously translates user ideas into incremental code development, producing, testing, and documenting software. Unlike AI tools generating isolated snippets, this manages full project lifecycles, syncing with the universal project framework for cohesive workflows. Zero-barrier to non-coders.

### Why This Matters

- Users with just an idea (e.g., for an app, game, or website) can create without grappling with coding complexity; the AI handles it all.
- LLMs struggle with long-context code and iterative revisions; sprints break projects into small, focused code blocks targeting single outcomes, reducing errors.
- Comprehensive documentation (run logs, user stories, release notes) aids the LM: logs speed debugging, stories link requirements to epics for organization, and change logs track what worked or failed, preventing repeated mistakes.

### Context & Positioning

While most AI assistants produce standalone code, this framework manages end-to-end development, from idea to deployment, integrating with the universal project framework for seamless workflows.

### Features

1. **Requirement Capture:** Asks clarifying questions to turn ideas into detailed user stories, defining functional, UI/UX, and performance needs (e.g., a game's scoring system).
2. **Sprint Planning:** Groups stories into prioritized sprints with clear deliverables, building incremental, testable code (e.g., a website's homepage in sprint 1).
3. **Sprint Execution:** Writes code, installs dependencies, runs tests, and updates logs/documentation per sprint, delivering working software (e.g., auto-documented API).
4. **Quality & Delivery:** Runs regression tests, packages deployment artifacts, and finalizes documentation autonomously for production-ready output.
5. **Context Sync:** Integrates sprint artifacts (stories, logs, docs) into the universal project framework, enabling access for other modules like deep research.

### Workflow

1. Discovery → User stories from ideas.
2. Planning → Prioritized sprint backlogs.
3. Execution → Code, test, document per sprint.
4. Delivery → Regression test, package, finalize.

## Context Refactoring

An automated mechanism that condenses and organizes chat history into focused summaries, ensuring clarity and efficiency in long dialogues. Unlike AI platforms that clutter conversations, this preserves essential context, syncing with the universal project framework for seamless reference.

### Why This Matters

- Keeps chats on-topic by highlighting key insights.
- Eliminates drift by pruning outdated or irrelevant messages.
- Cuts redundancy and scrolling for faster navigation.
- Reduces token usage, lowering costs and boosting performance.

### Context & Positioning

Most AI tools pile exchanges into a bloated buffer, causing confusion. This solution summarizes chats, archives full transcripts externally, and maintains a concise active context for accuracy.

### Features

1. **Summary Generation:** Distills chats into concise summaries of active ideas (e.g., a project's current goals). UI element within the chat that can be opened by the user or referenced by the LLM.
2. **Drift Elimination:** Removes discarded messages to prevent revisiting outdated concepts (e.g., dropped app features).
3. **Redundancy Pruning:** Filters duplicates for quick access to essential info (e.g., a single clear decision).
4. **Token Optimization:** Archives non-essential history to cut token count, preserving transcripts for reference.

### Workflow

1. User and LM chat; LM logs key information.
2. User discards ideas; LM updates logs, removing old ideas and adding new ones.
3. After X tokens, LM archives full chat, replenishing active chat with a concise summary from logs.

### Use Cases

- **Stay Focused:** A long chat yields a clear summary of active project goals.
- **Avoid Drift:** Discarded ideas are pruned, keeping the LM on my final plan.
- **Navigate Easily:** Redundant messages are cut, speeding up access to key points.
- **Save Costs:** Archived chats lower token usage, keeping budgets lean.

## Deep Research

An interactive research framework that delivers precise, user-guided workflows, blending discovery, documentation, and dynamic visuals without overwhelming users. Unlike tools that dump vast, hard-to-navigate data, this assistant offers tailored, high-value outputs, integrated into the universal project framework for seamless workflows.

### Why This Matters

- Cuts through data overload with concise, actionable results.
- Boosts confidence with trustworthy, well-researched findings.
- Saves time via interactive visuals, clear recommendations, and user interjection to bolster research (e.g., adding steps like “check patents”).
- Stores detailed notes for deeper exploration, accessible on demand.
- Fits agentic workflows, enabling multiple agents to research different aspects simultaneously while users interact or add steps.

### Context & Positioning

Most research tools overwhelm users with excessive data that’s difficult to interact with—lacking options to refine, summarize, or reformat. This framework centers on user-defined scope and visual workflows, enabling dynamic interaction (e.g., “Summarize this,” “Add visuals,” “Use arXiv”) for tailored outcomes.

### Features

1. **Interactive Discovery:** Clarifies user needs (e.g., thesis scope, trip preferences) to set focused research goals.
2. **Visual Workflows:** Provides real-time, editable workflow diagrams (e.g., drag-and-drop steps like *Deep Research*, *DeepThink*) for customizing research.
3. **Smart Research & Archiving:** A larger model conducts thorough research, storing insights in a thesis document or reference archive; a smaller model handles querying and refinement.
4. **User Validation:** Confirms new questions or focus areas, maintaining alignment.
5. **Actionable Outputs:** Delivers summaries, visuals (e.g., maps, tables), or formatted deliverables (e.g., reports, storyboards) for immediate use.
6. **Collaboration Tools:** Supports sharing canvases with real-time feedback, version control, and ethical citations (Scenario 1, advanced plans).

A premium multi-agent search lets agents explore queries in parallel, polling insights in a central workspace to collaboratively flag new avenues—yielding deeper, more exhaustive results.

---

## Scenario 1: Comprehensive Reporting

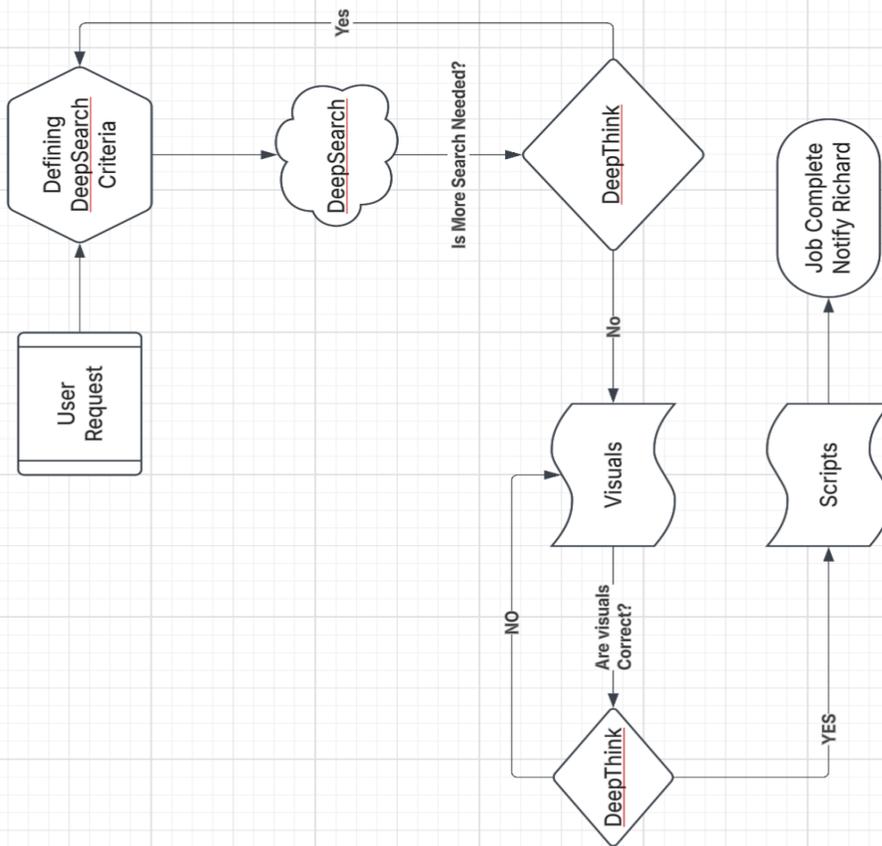
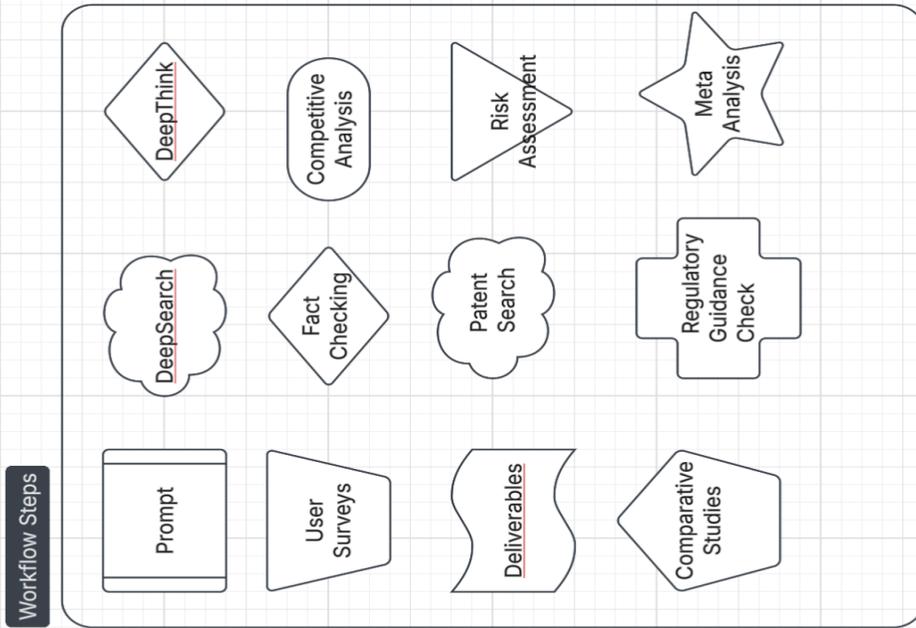
**Outline:** For long-term research (e.g., reports, white papers, blogs), users get an interactive canvas with drag-and-drop steps (e.g., *Deep Research*, *DeepThink*) and comment features (e.g., “Add regulatory check,” “Include counterarguments”). A thesis document tracks evolving needs, supporting collaboration over extended periods.

### Workflow:

1. User defines research topic (e.g., AI ethics paper).
2. Assistant clarifies scope (e.g., scientist perspective) and sources.
3. Generates interactive workflow diagram; user adds steps (e.g., search arXiv, check patents).
4. Larger model researches, updating thesis document; agents handle parallel tasks (e.g., law lookup).
5. Confirms focus areas with user.
6. Smaller model delivers summaries, visuals, and draft report; user refines for final deliverable.
7. Supports collaborative sharing with colleagues, including version control and citations.

### Use Cases:

- A researcher builds a white paper with a visual workflow canvas, adding *regulations* searches and regulatory checks, collaborating with a team over weeks.
- A podcaster/*blogger* conducts research for an article, adding *specific sites or social profiles* for additional sourcing.
- Independent news piece, steering the LM to include counterarguments, *fact checking and analysis*.
- A student compiles a college report, adding steps for *meta-analysis, pull in studies, provide detailed citations*.
- An innovator ideates a new widget, using research to explore processes *adding patent research and online stores*.
- Scientific research augments external AI processes, acting as an agent in a workflow *where it is called from another agentic workflow*.



## Scenario 2: Comprehensive Recommendations

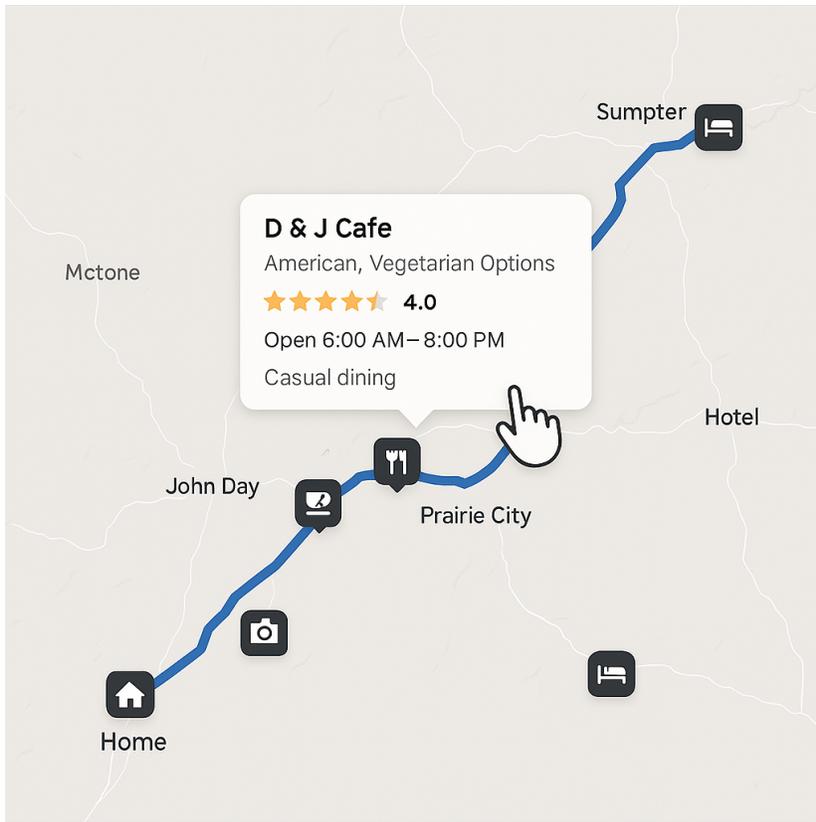
**Outline:** For quick answers, recommendations, or planning (e.g., road trips, purchases) *where the user is not expecting large amounts of data to sift through*. The assistant conducts deep research but delivers curated options, sparing users data overload. Focuses on short efforts with fewer steps without sacrificing research depth; archive remains accessible for querying. *Nimble enough to move quickly, deep enough search for intelligent answers*.

### Workflow:

1. Assistant clarifies preferences (e.g., trip interests, dietary needs).
2. Larger model conducts comprehensive research, storing details in a reference archive; agents research parallel aspects (e.g., POIs, hotels).
3. Smaller model delivers curated outputs (e.g., 3 restaurant options, POIs on a map, flight options).
4. Archive remains accessible for deeper inquiry if user needs more details.

### Use Cases:

- A user planning a road trip gets curated POIs, hotels, and flights based on preferences, choosing stops *on an interactive map* without seeing all research.
- A business dinner planner receives three restaurant recommendations with a map, derived from deep research into dietary constraints. *Restaurant information via a click*.
- A gardener gets researched planting options tailored to their climate, with actionable suggestions. *Interactive design showing crops over time*.
- A manager receives hiring options (salaries, qualifications, soft/hard skills) for a role, curated from extensive research. *Might start with a need or job title*.
- A buyer gets recommendations for cars, TVs, computers, or furniture, backed by deep research into specs and reviews.



### Travel Planning for Sumpter, Oregon

Gr

| Name                     | Rating | Price | Dietary             |
|--------------------------|--------|-------|---------------------|
| Sumpter Hotel            | ★★★★★  | \$\$  | V ✓                 |
| Beaver Creek Café        | ★★★★★  | \$    | Vg                  |
| Sumpter Valley Dredge    | ★★★★★  | 4,5   | Historic gold drege |
| Elkhorn Ridge RV Park    | ★★★★★  | \$\$  | Campground          |
| Gold Rush Espresso       | ★★★★★  | \$    | Great coffee        |
| Trail's End BBQ          | ★★★★★  | \$\$  | Casual dinir        |
| Sumpter Municipal Museum | ★★★    | Free  | Local history       |
| Pioneer Cabin            | ★★★★★  | \$\$  | Vegan-frien         |

A map showing a travel route with a callout box for 'Sumpter Valley Dredge' displaying the following information:

- Sumpter Valley Dredge**
- ★★★★★ 4.5
- Historic gold dredge

## Training in Contrarianism

A training framework embedding risk-awareness and negative-space analysis into LLMs, proactively identifying and mitigating potential failures in user proposals.

### Why This Matters

- Identifies risks by analyzing negative space (e.g., “What if this fails?”).
- Reduces sycophancy, fostering critical insights over blind agreement.
- Enhances decisions with domain-specific risk analysis, supporting regulated fields (finance, healthcare) with built-in risk checks and bias alerts for compliance.
- Adapts risk intensity based on context, minimizing unnecessary analysis.

### Context & Positioning

Unlike LLMs lacking risk-awareness, this framework trains models to question assumptions, tailoring risks to domain or user intent via passive (auto-surfaced risks in planning contexts) and active (user-controlled) adoption.

### Features

1. **Contrarian Training:** Flags risks via negative-space analysis (e.g., “What if widget materials aren’t available?”).
2. **Context-Aware Risk Assessment:** Auto-surfaces risks in planning contexts; minimal for simpler tasks.
3. **User-Driven Triggers:** Enables deeper analysis via keywords (e.g., “plan”), a risk-depth slider, or a secondary slider for user-defined risk criteria (text/document, stored in project files).
4. **Domain-Specific Insights:** Tailors risks to domains (e.g., healthcare) for actionable warnings.

### Use Cases

- A hiring manager sees candidate selection risks (e.g., skill mismatch).
- A buyer is warned of car purchase risks (e.g., price volatility).
- A user exploring dating gets compatibility risk flags.
- An innovator ideating a concept is cautioned on feasibility.
- A researcher generating a report is alerted to data bias risks.
- A developer coding an app is warned of scalability risks.
- A scientist testing a hypothesis is cautioned on design flaws.

## Podcaster

A premium, standalone framework that transforms diverse inputs—news articles, social posts, documents—into audio podcasts, featuring 1-3 (or more over time) distinct AI personalities—Host, Co-Host, and Guest—each with unique voices and traits, delivering an engaging, conversational discussion.

### Why This Matters

- Turns dense documents into approachable, dynamic, and pleasing podcast-style audio, creating a conversation for auditory learners or multitaskers.
- Customizable AI personalities enhance engagement with tailored voices and styles.
- Allows creators to guide the production within the platform, ensuring relevance and depth.
- Saves time by automating content narration while preserving research rigor.

### Context & Positioning

Unlike basic text-to-speech tools, Podcaster creates conversational podcasts with distinct AI personalities, enabling creators to shape dynamic audio narratives within the platform from static inputs like news articles, social posts, and documents.

### Features

1. **Personality Assignment:** Drag-and-drop 1-3 (or more over time) AI personalities—Host (leads), Co-Host (asks questions), and Guest (provides insights)—with distinct voices and customizable traits (e.g., authoritative, curious, expert).
2. **Dynamic Role Creation:** Users can create and add as many roles as needed (e.g., Technology SME for a tech podcast, Political SME for a policy discussion), then add roles to the podcast, tailoring it to specific needs.
3. **Interactive Visual Workflow:** A visual interface lets creators guide production by inserting prompts like “Ask questions on this,” “Explore this more,” or “Mention this,” shaping the podcast dynamically.
4. **Visual Integration:** Creators can drag in visuals or let the AI generate them, with personalities referencing visuals during the podcast (e.g., “As shown in this chart...”).
5. **Audio Rearrangement:** Rearrange audio clips to reposition concepts (e.g., surface a page 25 topic earlier) or copy segments to other parts, re-rolling the dialogue to blend seamlessly (e.g., “As we mentioned earlier...”).
6. **Research Invocation:** Creators can invoke a Deep Research stage from within Podcaster to enhance content, ensuring seamless integration without the Podcaster independently researching or collaborating on concepts.

## Use Cases

- A student uploads a history paper, assigning a scholarly Host, a curious Co-Host, and an expert Guest, guiding the podcast to explore key events further and referencing a timeline visual.
- A podcaster converts a script into a podcast, rearranging audio to prioritize key topics upfront and invoking Deep Research to add depth to a historical segment.
- An independent news creator transforms a report into audio, using a Political SME Guest to discuss findings while prompting, “Mention related events,” and integrating AI-generated visuals.
- A business professional converts a market report into a podcast, copying a segment to pair with competitor analysis and adding creator-supplied visuals.
- A trainer uploads a business domain corpus, creating a podcast with a Host (trainer) explaining concepts and a Guest (learner) asking how to apply them, or a YouTuber produces Excel training with a Host teaching and a Guest asking practical questions.

## Workflow

1. **Start with Podcaster:** Input documents, news articles, or social posts directly into Podcaster, or set upstream content (e.g., daily news articles, social posts) to auto-feed for fresh podcast generation; creators guide production as needed.
2. **Start with Deep Research:** Begin with a Deep Research deliverable, or invoke Deep Research from within Podcaster to enhance the document, feeding the enriched content into Podcaster for audio production.
3. **Podcaster to VidGen:** Podcaster deliverable feeds into VidGen to create a video with the audio and visuals, using a similar workflow to rearrange, preview, and output the final video, also autonomously if desired; all workflows can be autonomous after setup, saving the podcast file without human intervention.

## Revenue Realized

This premium feature can drive paid subscriptions by offering an innovative way to consume and present content, appealing to users seeking dynamic, conversational audio delivery.

## Closing Statement

The ideas presented are free enhancements to existing features, but I've also identified a wealth of novel, innovative concepts—unavailable on any LLM platform—that address real consumer needs through a user-first, technology-second approach. By focusing on delivering solutions users can't live without, we can drive soaring paid subscriptions and redefine how AI empowers daily life.

As Steve Jobs said, “We started with what incredible benefits we can give to the customer, not what awesome technology we have.” I invite you to reach out to explore my backlog of ideas and partner together to unlock immense potential, shaping the future of AI for your users.

Thank you for considering these concepts—I'm committed to bringing these solutions to fruition, adding immense value to people through AI.

Richard Getz  
484-680-5153

<https://x.com/RichardAGetz>

PDF: <https://t.co/4wgXfSObbC>

X Post: <https://x.com/RichardAGetz/status/1907163228838547812>

Thank You

Richard

Thank You  
Richard