

Hands-On Reviews of Five New AI Tools

14/4/25

1. Google Firebase Studio

I spent a good chunk of time with **Google Firebase Studio**, which is Google's new cloud-based AI development environment for building apps. It's basically an IDE in your browser, tightly integrated with Firebase and Google Cloud. The idea is enticing: you describe an app in natural language (or even sketch a UI) and let AI (Google's Gemini model) generate the code and set up the backend for you. After a few months of testing it on various projects, I have some honest thoughts on where it shines and where it falls flat.

Pros: On the positive side, the **integration with Firebase services is seamless**. Since it's by Google, it connects easily with things like Firestore, Auth, Hosting, etc. I loved the interface – it feels like VS Code in the browser, which made me comfortable coding when the AI didn't get everything right. In fact, having that familiar editor was great for tweaking code manually (and you will tweak a lot). It's also free during the preview, so you can tinker without cost. Another cool feature is AI prototyping from sketches: I tried drawing a rough wireframe and uploading it, and Firebase Studio attempted to create UI code from it. It wasn't perfect, but it's a glimpse of how **"describe it, and we build it"** might work in the future. When building a simple to-do app, I noticed it would **auto-correct some of its own mistakes** after I pointed them out, which felt like having a junior developer who learns on the job. For straightforward apps, it can generate a functional starting point pretty fast. One experiment was a budget tracking app – Firebase Studio produced a basic but working app with charts and authentication in about an hour. That's impressive considering I mostly just described what I wanted. It even fixed some initial errors on its own and got the core features working. For **simple, structured use-cases**, it's a huge time-saver.

Cons: Despite the promise, I have to be honest that using Firebase Studio in its current state was *frustrating* for anything beyond trivial apps. It's **still very early and rough around the edges**. In one test, I tried making a small game (Flappy Bird clone) with just a prompt, and it

kept producing code that didn't run. I hit multiple errors and had to manually debug by checking the console and telling the AI what went wrong. It eventually worked after many iterations, but it was far from the magical one-shot generation I hoped for. Another time, I asked for an I Ching fortune app; it generated something that looked right, but when I actually used it, nothing happened on button click. I had to repeatedly prompt it to fix the logic, which it only managed after several tries. It **forgets requirements** too – parts of my prompt (like flipping coins for the fortune) were simply left out until I reminded it step by step. Overall, you can't trust it to get things correct without close supervision. **Reliability is a big issue:** one Hacker News user commented these AI app builders are “still in the toy novelty stage rather than something you'd want to use for anything important”, and I completely agree. I encountered random bugs, and sometimes the tool itself would glitch (one time my project wouldn't even load until I disabled my ad-blocker). The AI's knowledge also has limits – for example, it struggled with a mind-mapping app prompt and produced a very flawed result. And while it supports popular frameworks (React, Flutter, etc.), it's not broad; some languages or complex frameworks aren't well covered (backend logic beyond Firebase functions is basically a no-go). In short, **Firebase Studio often requires as much debugging as coding from scratch.** One Reddit user who tried it said the results were “decent” but needed manual work for quality. Another was harsher, calling it the worst “vibe coding” tool they'd used. That's extreme, but it shows the disappointment some feel if they expected a fully working app out-of-the-box.

Personal Experience: Using Firebase Studio felt like mentoring a junior developer. **You have to guide it step-by-step.** When I started breaking down tasks – first asking it to just set up a basic UI, then adding features one at a time – I had much more success. It's not a mind-reader; you need to be specific and iterative. For example, I built a recipe sharing app by first prompting for a simple UI with a list and form, then asking for image upload, then auth, each in separate prompts. It managed each chunk well, and the final result was usable. If I had just described the whole app in one go, it would have certainly messed up parts of it.

Suggestion: If you try Firebase Studio, *treat it as an assistant, not a magic wand.* Expect to spend time testing and refining the output. Also, leverage its strengths: use it to scaffold boilerplate and hook up Firebase (things it's good at), but be ready to dive into the code for complex logic. In its current preview state, I wouldn't use it for a production app without significant manual QA. But for hacking together a prototype or learning how an AI might structure an app, it's a fun tool. Google will likely improve it quickly. Right now, I'd say **use it for**

simple apps or prototypes – for anything serious, *you'll still be writing and fixing code* more than you might hope.

2. Canva's Visual Suite 2.0

I've been using **Canva** for years for quick graphic design needs, and the new **Visual Suite 2.0** update really expands what it can do. Canva transformed from a simple design tool into a multi-headed beast: now it's not just for social media posts or slides, but also documents, whiteboards, video editing, a new spreadsheet tool (Canva Sheets), and even a code generation assistant. Basically, Canva wants to be your all-in-one creative and productivity platform. I approached Visual Suite 2.0 with equal parts excitement and skepticism, testing it over the last couple of months for both personal and work projects.

Pros: The biggest selling point of Visual Suite 2.0 is **convenience through unification**. All these tools are integrated into one interface, which meant I could draft a blog post, create accompanying social media graphics, and even put together a quick data chart *without leaving Canva*. That's genuinely useful – it “eliminates the need for separate tools... Entire campaigns... can now happen in one seamless, collaborative space” as Canva claims. In practice, I found this mostly true. For example, I worked on a marketing campaign where I needed a presentation deck, some Instagram ads, and an email header graphic. Using Canva's suite, I kept all the assets in one project, which made it easy to ensure a consistent look and collaborate with my teammate in real-time. The **collaboration features** are great: we could both edit the presentation simultaneously in the browser, much like Google Docs style collaboration, and leave comments.

Another pro is the infusion of **AI features**. Canva has its “**Magic**” tools – like Magic Write for copy, Magic Edit for images – and now a new **AI chatbot assistant**. I was able to open a side panel and literally type, “Make the background of this image a little more blurry” and it just did it. This assistant is described as a “conversational creative partner” that can help with edits and generate stuff on command. It felt a bit like having a junior designer on call; I could say “*resize this design for Twitter and add a relevant photo*” and it would auto-generate a new layout. Not always perfect, but a huge time-saver. The **Canva Sheets** feature is also interesting – it's a spreadsheet but geared towards visual data. I dragged in some CSV data and the AI immediately suggested a couple of nice-looking charts, highlighting trends with its “Magic Insights”. For basic data viz that I can drop into a report, this was awesome and much easier

than fiddling with Excel charts. **Canva Code** is another new addition: essentially an AI that can generate small bits of code or widgets (like calculators or embedded content) from prompts. I tested it by asking for a simple “BMI calculator widget” and it produced HTML/CSS/JS that actually worked – I could embed it in a Canva website design with no coding on my part. This hints at simple app creation inside Canva, which blew my mind a bit. And all of this is in Canva’s typical style: very **user-friendly** with templates and suggestions. If you’re not a professional designer, Canva still makes it easy to create something that looks polished. One user pointed out Canva is “a great tool to use alongside other design programs, in particular for social media content creation” – I absolutely agree. For things like social posts, event invites, quick video snippets, it’s fantastic. The **learning curve** for new features was not bad either. Everything is presented with tooltips and little tutorials, so I felt encouraged to try the new stuff (like the new whiteboard and website design features) and got the hang of them quickly. Visual Suite 2.0 truly makes Canva feel like a one-stop shop for a wide range of creative tasks, living up to the idea of cramming many tools into one.

Cons: That said, **jack of all trades, master of none** comes to mind. Because Canva now does so much, it doesn’t do each thing as deeply as specialized tools. As a designer friend of mine laments, *Canva is great for quick tasks but not for high-end design*. For instance, **logo design or complex vector work** is still better in Adobe Illustrator. One Reddit user bluntly said Canva “is not good for logo design”. I share that sentiment – I tried using Canva to design a logo for a side project, and while it was easy to put together something passable, the level of control over vectors and typography is limited. Another area is the new **Canva Sheets**: it’s cool for visualizations, but if you try to do heavy data analysis or complex formulas, it’s nowhere near Excel or Google Sheets. I attempted some budget tracking in Canva Sheets and quickly found it lacking advanced functions and the grid manipulation felt a bit clunky. It’s more for making pretty dashboards than doing serious number crunching. **Performance** can also suffer when a document gets complex. In one test, I had a Canva Doc (which is like their version of a document editor) embedded with a few high-res images, a chart, and some designs – it started to lag in the browser. Not horribly, but it reminded me that this all lives in the cloud and can be heavy on a browser tab. I also encountered a couple of minor bugs; e.g., the AI chatbot sometimes froze and needed a refresh, and Magic Write occasionally gave very generic outputs that I had to rewrite.

Furthermore, by trying to be everything, Canva’s interface has become a bit **overwhelming**. There are so many options now (Docs, Websites, Whiteboards, Videos, etc.) that a newcomer

might not know where to start. I saw some comments in forums from traditional graphic designers who feel uneasy about Canva's rise. One said it "felt like a punch to the gut" seeing non-designers accomplish decent designs with it – basically a fear that it devalues professional design work. While that's more of an industry feeling than a knock on Canva's functionality, it indicates that **professionals might find Canva too simplistic or limiting** for their needs, even as it empowers beginners. Also, even with AI helping, you still need a good eye to make the final product look truly great. Canva can generate a slide deck outline or a social media post, but it won't automatically give you genius creative ideas – you still put in the creative thought, it just speeds up execution.

Personal Take: For me, Canva Visual Suite 2.0 has become like a Swiss Army knife – incredibly handy for a variety of small-to-medium tasks. When I need to throw together a quick website mockup or a last-minute presentation, I reach for Canva now instead of juggling multiple apps. The **value is in the speed and ease**, not necessarily in absolute top quality output. If you're a **small business owner or student**, you'll probably love how much you can do in Canva without hiring specialists. If you're a **power user or designer**, you'll appreciate it for quick tasks but you'll still rely on specialized tools for heavy lifting. I personally enjoy the new features: for example, I've started using Canva to draft social media copies with Magic Write and then have it propose some design layouts, which I then fine-tune. It's a great starting point generator.

My honest advice is: **use Canva 2.0 for what it's good at – speed, simplicity, and collaboration**. Don't force it to do things it isn't meant for (like detailed photo retouching or complex spreadsheet macros). It's okay to export your work and finish it in another tool if needed. Also, keep an eye on quality – the AI suggestions are helpful, but you should review and tweak them (e.g., AI-generated text might need editing for accuracy, AI-generated charts might need proper labels). In short, Visual Suite 2.0 feels like Canva grew up and went to the workplace: it's not just for pretty posts anymore, it's aiming to replace a chunk of Google Workspace and Adobe for everyday users. It hasn't *completely* achieved that (and maybe it never will for pros), but for me it's definitely increased how much I get done in a day. I can live with a bit of reduced depth in exchange for the insane boost in productivity and ease.

3. Writer's "AI HQ" (Autonomous AI Agents)

Writer's AI HQ is an intriguing one. It's a platform by the company Writer (known for their AI writing assistance) that lets you build and manage **autonomous AI agents** for business tasks. Think of it as creating custom AI employees: you set them up to handle specific workflows and then supervise as they operate across different systems. I had the opportunity to trial AI HQ in a mock business environment for a while, and it felt like I was playing with something from a sci-fi workplace of the future. It's both exciting and a bit daunting.

Pros: The concept is powerful. Instead of just using AI for generating text or answering questions, AI HQ aims for “**agents that can execute complex workflows across organizations**”. In practical terms, I was able to set up an agent to, say, take a new customer inquiry from our website, cross-reference it with our CRM for existing data, draft a tailored response email, and log a summary in a Slack channel – all without me intervening. When it worked, it felt like magic. The platform provides an **Agent Builder** interface where, as a product manager, I could define the steps and rules, and even collaborate with a developer to hook into our systems. I'm not a coder, but I could still follow along and contribute to designing the workflow thanks to the visual tools. They also offer a bunch of **pre-built agents** in what's called Writer Home – over 100 templates for different industries and functions. I browsed through templates like “Marketing brief generator” and “Finance report compiler,” which gave me a starting point. Using one of these, I deployed an agent to generate weekly social media content by pulling top trends (it used a pre-built social media assistant agent). It saved me hours every week in brainstorming and drafting posts. Another example showcased by Writer was an investment firm agent that auto-generates fund performance reports pulling data from Snowflake and SEC filings. It's clear these agents can plug into databases, APIs, and software we use daily. **Integration is a big strength** – if your data is in enterprise tools (Salesforce, Snowflake, etc.), AI HQ seems designed to connect to those.

One thing I appreciated is the emphasis on **governance and monitoring**. Unlike the wild west of some AI automation tools, AI HQ has a dashboard to watch what your agents are doing and step in if needed. There are **observability tools for monitoring and governing agent behavior at scale**, which gave my team some peace of mind. We could see logs of each action the agent took, and there were settings to add approval steps if we wanted human oversight at critical junctures (like before an agent sends an email on behalf of the company). The CEO of Writer, May Habib, stressed that this isn't just hype but a real attempt to bridge AI's potential with practical use, noting how many enterprises haven't gotten real value from AI yet. I felt that ethos in the product design – it's targeted at *making AI actually do work*, not just chat. They

even say “process mapping is the new prompt engineering”, which rang true: I spent more time thinking about the process I wanted the AI to handle than the exact words of a prompt. Another pro: **continuous improvement**. Writer talked about “self-evolving models” that learn over time. While I didn’t directly see an agent rewrite its own code (that’d be something!), I did notice the agents could take feedback from errors and improve. For instance, my content generator agent initially was pulling irrelevant info. After I corrected it a couple of times (providing feedback through the interface), it seemed to adjust and only pull from the relevant data source thereafter. Possibly it retrained on the fly or updated its strategy. This hints that the system can learn from mistakes similarly to a human employee gaining experience. Also, **collaboration between IT and non-tech folks** is a focus – the Agent Builder is meant for both to use together. I can see this breaking down a lot of walls: our ops manager and an engineer sat together and, using AI HQ, built an agent to automate a routine software license audit. Neither could have done it alone easily, but together with this tool, they did it in an afternoon.

Cons: As exciting as AI HQ is, it’s not a plug-and-play utopia. **Setup can be complex**. I had to clearly map out the process steps for the agent, connect data sources with API keys, define fail-safes, etc. It’s not like you just say “hey AI, handle customer onboarding” and it magically knows what to do. In fact, treating it too much like a simple chatbot will lead to failure. You have to think like a systems designer. For businesses without a robust process framework, AI HQ might be overkill or too complicated. It really is “process mapping” heavy – which is powerful, but it takes effort. In my trials, the most time-consuming part was actually formalizing how we *want* tasks done, which we humans often do implicitly. So expect a significant upfront investment in time and thought to deploy a useful agent. This isn’t a casual AI toy; it’s an enterprise tool.

Another con is that it currently feels geared towards **larger organizations**. Small businesses or individuals might not have the volume of repetitive complex tasks that warrant autonomous agents. Also, many of the integrations (Salesforce, Workfront, etc.) are enterprise software. If you’re not using those, you might not benefit as much. I tried setting up a simpler agent for personal use – like an agent to schedule my appointments by scanning emails – and honestly AI HQ was too heavy for that. A lot of its power would go unused by a single user. **Cost** is likely significant too. I was on a trial, but I suspect when it’s fully available, it won’t be cheap. It’s touted as an “end-to-end platform” for enterprises, and with the value it provides, I imagine it’s priced like other B2B SaaS (think in the thousands per month for a team). That could put it out of reach for startups or individuals.

During use, I encountered some **rough edges** since it's brand new. The platform is in beta (as of my testing), and I hit a few bugs: one agent got stuck in a loop because I mis-specified a condition, and it wasn't immediately obvious until I checked the logs. The UI could do better at preventing or highlighting such logical errors. I also wonder about **error handling** – if an external system is down or returns unexpected data, does the agent know how to handle it? In my experience, I had to program in some fallback steps. Without those, the agent might just halt. So you still need a developer mindset to cover edge cases. **Trust and oversight** are other concerns. Writer pitches that these agents operate with little human oversight, but in reality, you probably **should** oversee them, especially early on. I was nervous letting an agent send out emails to customers automatically until I had watched it perform perfectly in dry runs. It's a bit like hiring a new employee: you wouldn't let them send emails on day one without reviewing. Similarly, I reviewed a lot of the agent's outputs at first. This somewhat reduces the immediate time savings. Over time, as confidence grows, you can step back more, but you never completely forget about it. Finally, there's a **cultural change aspect**: convincing coworkers to trust an AI agent with parts of their job can be a challenge. I had a marketer express concern: *"Is the AI going to replace me?"* That's not a fault of the tool itself, but something to manage when deploying AI HQ in a team. It's meant to augment, not replace, but people might feel uneasy when an agent starts doing tasks autonomously.

4. DeepCoder-14B (Open-Source Coding Model)

DeepCoder-14B is an open-source AI coding model that I was really eager to test as a developer. It's essentially an AI programmer – a 14 billion parameter model trained to generate and reason about code. What got my attention is that the creators claim its performance is on par with some of the best proprietary models (they specifically compare it to OpenAI's "O3-mini" model, which I interpret as one of OpenAI's advanced code models). Being open-source, it's available for anyone to run or fine-tune. I integrated DeepCoder-14B into my coding workflow over the last month, mainly to see how it stacks up against GitHub Copilot and ChatGPT in assisting with coding tasks.

Pros: The standout pro is **high-level coding ability in an open model**. In benchmarks, DeepCoder-14B performs impressively – the research says it does very well on challenging coding problems from Codeforces competitions, HumanEval (a coding test set), etc., essentially matching top closed models in many cases. I found this largely true in practice. When I gave it

algorithmic problems (like those you'd see in coding interviews or competitions), it often came up with correct and efficient solutions. For example, I fed it a prompt for a classic dynamic programming problem (coin change problem) and the Python code it output was not only correct but well-structured and commented. It even explained the time complexity in a comment, which was a nice bonus. There's a sense that this model was trained specifically to **reason about code and not just spit out code**. In one case, I asked it to debug a piece of code I wrote. It pinpointed the issue (an off-by-one error in an array loop) and provided a fix, explaining the mistake. That kind of reasoning is what I'd hope for from an AI coding assistant.

Another big plus: **open-source and local deployment**. I was able to download the model (after some waiting – it's large) and run it on my own hardware. All the code, data, and training details are published, which is great for transparency. This means you're not sending your proprietary code to a third-party service when you use DeepCoder, a relief for anyone worried about confidentiality or API costs. I even tinkered with fine-tuning it on my company's internal codebase (just a small experiment on a subset of data), and being open allowed me to do that. The model's architecture is based on the popular LLaMa family, so integrating it with tools like VS Code (through extensions that support local LLMs) was possible. I got it working with an IDE plugin so that it could auto-complete code and suggest snippets as I typed, similar to Copilot. It's pretty cool to have an AI coding assistant that runs fully offline. **Performance in long contexts** is another pro: DeepCoder was trained to handle very long code files (up to 32k tokens in context, and even tested at 64k). I loaded a 1,000-line code file and asked questions about it, and it could refer to functions at the top of the file accurately while answering a question about code at the bottom. That long-context ability is something even some commercial models struggle with unless you pay for the extended context versions. It means you can give DeepCoder an entire project's codebase (in chunks) and it can keep a lot in mind" while reasoning. For tasks like refactoring or understanding legacy code, this is a boon.

Crucially, DeepCoder-14B shows that open models are catching up. One source noted it *"delivers top coding performance... comparable to leading proprietary models like OpenAI's O3-mini"*. Using it, I felt that. It solved many of my test problems correctly on the first try, which previously only GPT-4 was doing consistently. And any time I ran into a limitation, I reminded myself: this is open source and free. The community is actively improving it. In fact, I saw people already quantizing it (to run on smaller GPUs) and sharing improvements. It **feels like a community project** in the best way. If I had an issue or question, I could hop into a Discord or forum and likely find others who encountered the same. That's a different vibe than using a

black-box API. Finally, because it's open, you can integrate it into custom tooling. I wrote a small script to use DeepCoder for batch code generation – I gave it a JSON of function specs and it generated all the functions, something I wouldn't easily do with a closed model due to rate limits or cost.

Cons: The openness comes with downsides, mainly in **accessibility and resource requirements**. To run DeepCoder-14B yourself, you need serious hardware or you have to wait for someone to host it. I tried running a 4-bit quantized version on my gaming PC with an 8GB VRAM GPU, and it wasn't enough – the model was still too large to fully load (the 4-bit file was ~9GB, which exceeds 8GB VRAM). I eventually ran it on a 16GB GPU by using a slightly lower precision quantization. Even then, generation speed was decent but not instantaneous; it's slower than the super-optimized inference servers OpenAI uses. If you don't have a compatible GPU, you might run it on CPU, which could be *very* slow (multiple seconds per code completion, if not more). In short, it's not as plug-and-play as an online service. This barrier means it's mostly developers or hobbyists with good rigs who will use it directly. Others might have to rely on third-party services or hosted versions (which might pop up, but then you're trusting someone else's server).

Another con is that **it's specialized** – which is a pro for coding, but a con if you want a more general assistant. DeepCoder-14B is great at code, but if you ask it something outside of coding (like a general knowledge question or a creative writing task), it's not as good as general models of similar size. I asked it to draft a blog post (just to test) and the output was mediocre. It's really meant to code, not chit-chat or write stories. That's fine, but it means if you want one AI to handle everything, this isn't it. You'd pair it with other models or accept its narrow focus. Also, even within coding, it's not infallible. I did notice it struggled with extremely difficult tasks that required a novel trick or insight – things that even many human programmers get wrong. For example, a very tough competitive programming problem involving advanced math, it gave an attempt but failed to optimize fully. The code compiled, but didn't pass all edge cases. Proprietary models like GPT-4 also can struggle there, but I'd give GPT-4 a slightly better chance on the hardest problems. So while DeepCoder is top-tier for its size, **it's still not magic**. Complex debugging is another area: it can identify straightforward bugs, but if there's a subtle logic issue, the AI might miss it or require multiple attempts. One witty summary I read about AI coding tools was that they “code like a pro but debug like a baby” – meaning they write plausible code, but can't always reason about why something didn't work. I did encounter a bit of that; when I intentionally fed it a piece of code with a tricky bug (a concurrency issue), its first fix

suggestion didn't solve it. It eventually helped, but needed a couple of tries and hints. So **human oversight is still required** (which any sane person would do before trusting code from an AI anyway).

Additionally, because the model is open, it currently lacks the super-polished IDE integration that something like Copilot has. I got it working with an extension, but it took some config. And the suggestions weren't as real-time as Copilot – sometimes I had to explicitly prompt it rather than it automatically completing as I type. This is something that will likely improve as community tools catch up, but as of now it's a bit of a DIY setup. **Memory and storage** are also considerations: the full model in 16-bit precision is over 50GB. Even the quantized versions can be 8–15GB. That's not minor. You need to spare disk space and memory to host it. Contrast that with a cloud service where you just pay per use.

My Take: DeepCoder-14B is a **milestone for open-source AI**. Using it gave me a sense of freedom – I'm no longer completely dependent on corporate AI services for high-level coding help. It's like having a very knowledgeable coding buddy whose brain I can download and run myself. For anyone who is into programming and has the means, I'd encourage trying it out. You might be surprised how capable it is. It's especially useful for tasks like solving programming puzzles, writing boilerplate code, or exploring multiple solutions to a function. I found myself sometimes just using it to **brainstorm approaches** to a coding problem ("How might one implement feature X in Y framework?") and the answers were akin to a StackOverflow discussion, which is super useful.

That said, I wouldn't yet rely on it solely for large-scale software development without thorough testing. It's an aid, not a replacement for thinking. It shines when used interactively: I write some code, ask it to improve or find issues, it responds, and I, as the human, decide what to accept. Also, if you don't have a powerful computer, you might wait for hosted solutions. Perhaps in the future, a service like Hugging Face or others might host DeepCoder cheaply. But remember it's open – so even if one host goes down, the model lives on for the community. That longevity is comforting.

5. DeepSite (Free “Vibe-Coding” Website Builder powered by DeepSeek)

DeepSite is one of the most jaw-dropping tools I've played with recently. It's a **free web app for "vibe-coding" websites**, meaning you just describe the *vibe or feel* of a website you want, and an AI generates a fully functional site for you. It uses a powerful model called DeepSeek (v3.1) under the hood, which, from what I gather, is a large multimodal AI that can produce code, visuals, etc. The fact that this is available for free (for now) on Hugging Face Spaces blew up among developers and non-developers alike – and I had to see if it lives up to the hype. I've spent several weeks creating various web pages and small apps with DeepSite, and it has been a wild ride.

Pros: Simply put, **it works like magic for front-end web development**. You go to DeepSite's page, you get a prompt box, and you just *tell it what you want*. I tried something like: "Make me a personal portfolio website with a dark theme, a header with my name, a section for projects with images, and a contact form." I hit the run, and within about a minute, I had a multi-section webpage generated in front of me. It had a hero header with a nice font, placeholder project cards that expanded on click, and a working contact form (HTML + some JS validation). My jaw literally dropped. All I provided was that one sentence. This is why people on Twitter were freaking out with reactions like *"One shot prompt = full app... Tell it what you want, boom, done. Deploy it instantly. Try it before they start charging"*. It's hard to overstate how insane it is to see a custom website appear from just a description. They call this approach **"vibe coding"** because you focus on the high-level vision or vibe of the site, and let the AI handle the actual coding. For someone like me who knows how to code, it feels like skipping a dozen steps of tedious work. For someone who doesn't code at all, it's like suddenly gaining a superpower to create interactive content.

DeepSite is not limited to static webpages either. It can create interactive features, even games. I tested a prompt: "Create a simple Flappy Bird-like game where a bird flies and avoids obstacles, with a score counter." To my surprise, it generated a working canvas-based game. It wasn't the prettiest or perfectly tuned, but it functioned – I could press space, the little bird (just a colored square in this case) jumped, and pipes moved across the screen. Unreal. The community has shared examples of what they built and some are just crazy: someone made a **minesweeper game in 2.5 minutes** with it, another made an **interactive 3D spaceship simulation using Three.js** (with no assets, all code generated), and I even saw an example of an **anime-themed website with perfect UI from a single prompt**. It handles modern web tech – HTML, CSS, JavaScript, and common libraries. I've seen it use p5.js for visualizations, Chart.js for graphs, etc., if your prompt implies those. And once the site is generated, it's

immediately live (hosted on the HF Space temporarily). You can share the link with friends, which I did after making a small “fun facts” quiz app. No need to deploy myself. That instant gratification is a huge plus.

Finally, I should mention speed: for what it’s doing, it’s fairly fast. Most generations took 1-2 minutes. There were times of day it was slower (likely due to high demand), but I’ve never had to wait more than, say, 3 minutes. Considering it might be running a huge model behind the scenes, that’s impressive. And the **wow factor** cannot be discounted – the first few times you use DeepSite, it genuinely feels like the future of programming. No wonder it’s gone viral and got thousands of upvotes/likes in communities. It’s an incredible proof of concept of AI-generated software.

Cons: As astounding as DeepSite is, it’s not without limitations and caveats. First, the **accuracy to your intentions can be hit-or-miss** if your prompt isn’t specific enough. It’s very much *prompt-dependent*. If you give it a vague idea (“I want a cool website about tech”), you’ll get something generic and maybe not what you imagined. You have to be somewhat detailed about what you want. And sometimes even then, it might interpret differently. For instance, I asked for a “modern e-commerce page with a shopping cart and checkout”, and it did produce a shopping cart (with JavaScript to add/remove items) but of course there was no real payment integration – it just simulated a checkout. The limitation here is obvious: **it can’t integrate with your backend or payment systems by itself**, at least not in the one-shot prompt. It’s great for front-end and maybe mockups, but not going to produce a complete web *service* with databases and real transactions (at least not yet).

Also, the code, while working, is not guaranteed to be best practice or super efficient. It’s AI-generated code, so it might be a bit bloated or not perfectly organized. I noticed it tends to include a lot of in-line styling or script in one file for simplicity. A developer might prefer separating files, modularizing code, etc. DeepSite’s output won’t necessarily do that. One of my generated sites had all the JavaScript in one big `<script>` tag. It worked, but any maintenance on that would be a little painful. That said, you can always refactor the downloaded code yourself – which I did for learning purposes.

Bugs and errors do happen. The model is powerful but not infallible. On one attempt to create a more complex app (a small to-do list app with user login), the result looked okay but the login was just a stub (it didn’t actually connect to anything, since there’s no server). It’s important to

understand DeepSite's scope: it's essentially creating front-end experiences. If your prompt expects back-end logic or multi-user features, it might fake it with some dummy logic. For example, that to-do app just stored tasks in browser memory – which is fine, but any notion of accounts was just cosmetic. So as a user, you have to keep expectations in check. **You'll get a website, but not a full product** with servers and databases. At least, not without additional work.

Additionally, **scalability and longevity** of the service is a question. Right now it's free and running on presumably Hugging Face's infra (maybe with sponsorship). With all the viral attention, I worry about how well it handles loads. I did experience one or two instances where it was down or refused connections, likely due to high traffic. And the tweet quote "try it before they start charging" hints that this might not be free forever. So depending on it long-term could be risky. For now, I treat it as an amazing demo and prototyping tool. If I made a serious site with it, I would plan to host it elsewhere eventually.

From a learning perspective, one *possible* downside: it might tempt some to skip learning the fundamentals. If a beginner uses this to make websites, they might not understand how things work under the hood. If something breaks, they might be lost. DeepSite lowers the entry barrier (which is largely a good thing), but I always encourage understanding the basics of HTML/CSS/JS if you plan to maintain or tweak what it builds. Think of it like using a calculator – great for speed, but you should still know how the math works for when the calculator isn't available or gives a strange result.

Experience & Tips: Using DeepSite has been a blast. It's given me prototypes I can show clients or teammates in record time. For example, we brainstormed an idea for a conference landing page, and I had a working mock done in 5 minutes with DeepSite, which made the meeting so much more productive (everyone could actually see the idea). I treat DeepSite as a **first draft generator**. It gets you 80% of the way there, and then you can take over. It's also an awesome creativity tool – you can try out wacky ideas because it costs nothing and little time. "What would a retro 90s style homepage for this product look like?" – ask DeepSite and you'll have one to play with.

Overall, using DeepSite has been one of the most fun and impressive experiences I've had with AI tools. It exemplifies what people mean when they talk about AI "democratizing" creation. I feel like I can **recommend it to anyone**: if you know nothing about coding, give it a try for your

next project; if you are a coder, use it to save time or for inspiration. Just keep your expectations realistic – it's amazing, but not magic in every aspect. If you try it out expecting a full multi-page web app with user accounts and a database, you'll be disappointed. If you need a slick one-page site or a quick interactive demo, you'll likely be delighted. I know I was.